The market for influence

Itay P. Fainmesser and Andrea Galeotti*

July 21, 2018

Abstract

Influencer marketing is the fast-growing practice in which marketers purchase product endorsements from influencers—who are individuals with many followers and strong reputations in niche markets. This paper develops a model of the market interactions between influencers, followers, and marketers. Influencers trade off the increased revenue they obtain from more paid endorsements with the negative impact that this has on their followers' engagement, which in turn affects the price marketers are willing to pay for their endorsement. Our analysis provides testable predictions on how the price that influencers receive depends on the size of their audience and how an improvement in the online search technology affects influencers' competition for followers and marketers. We show that, in equilibrium, over- and under-provision of paid endorsements coexist. We evaluate the strategic effects of recent, transparency-motivated policy interventions implemented by competition authorities in the US and Europe, requiring influencers to clearly indicate marketer-sponsored content.

*Fainmesser: The Johns Hopkins Carey Business School and The Department of Economics, The Johns Hopkins University, Baltimore, MD 21202 (email: Itay_Fainmesser@jhu.edu); Galeotti: Department of Economics, European University Institute and University of Essex, Florence, Italy (email: andrea.galeotti@eui.eu). For their insightful comments and their advice, we thank Nageeb Ali, Kfir Eliaz, Glenn Ellison, Jack Fanning, Ben Golub, Manuel Hermosilla, Matt Jackson, Wei Li, Jacopo Perego, Eduard Talamàs, Ruxian Wang, Leeat Yariv, participants of the Retreat on Networks, Information, and Social Economics, and seminar participants at Brown University, The Johns Hopkins University, London Business School, Paris School of Economics, and the Sixth European Meeting on Networks in Barcelona. We thank Simon Skipka and Jianhui Li for excellent research assistance. Andrea Galeotti is grateful to the European Research Council for support through ERC-consolidator grant (award no. 283454).

1 Introduction

Social media influencers are individuals who produce online content, often focused on one product category, and have a large number of followers engaging with and trusting their recommendations. Influencers create blogs and are active on Instagram, Twitter, and other online platforms. They reach audiences ranging from as few as 3000 people to more than 20 million and are most prominent in categories such as fashion & style, tech & gaming, and food & lifestyle. An influencer who reaches many potential customers and whose audience perceives his recommendations/advice to be authentic attracts the attention of marketers, who are willing to pay him to endorse brands and products.[1] This marketing practice, referred as to influencer marketing, is growing and becoming a sizeable component of the marketing budget of many brands.[2]

The fast growth of influencer marketing has led competition authorities in many countries to monitor the relationship between marketers and influencers. Recent interventions ask influencers to clearly indicate marketer–sponsored content in order to make influencers' recommendations more transparent.

We develop a model of the market for influence which captures the aspects that distinguish it from other media markets. We derive testable implications for the pricing of sponsored content, uncover inefficiencies in its provision, and evaluate the aforementioned policy interventions, and the potential impact of changes to the online search technology by which followers find influencers.

The market for online influence consists of followers, influencers, and marketers. An influencer features products and describes his own experience with respect to varieties offered in the market. In this way, an influencer offers advice to followers about a specific product category. Marketers may pay influencers to sponsor their product, with the hope of boosting their demand. Each influencer is thus a two-sided platform whose recommendations facilitate the matching between marketers and followers.[3] Since there are no major/substantial fixed

[1]This is different from display advertising, often provided by large platforms, in which authenticity does not play a role. For more on display advertising, ad-targeting, ad-skipping, and ad-blocking, see the references in Section 6.

[2]We refer the reader to recent articles in *The New York Times* ("Inside the Mating Rituals of Brands and Online Stars") and *Forbes* ("7 Predictions on the Future of Influencer Marketing"), and a recent 2018 report from eMarketer, both discussing trends in influencer marketing. These articles—representative of the discussion that can be found on influencer marketing among practitioners—depict a strong shift of marketing resources away from traditional advertising practice towards influencer marketing. eMarketer reported that, in 2017, globally, marketers spent $570 million on influencer marketing on Instagram and that, in a survey of beauty marketers worldwide, 66% reported allocating 10–50% of their marketing budget to influencer marketing. *Forbes*'s article reports that "as many as 39 percent of marketers are actively increasing their marketing budgets for 2018, with 86 percent of marketers relying on the strategy for at least some of their 2017 campaigns." *The New York Times*'s article focuses on YouTube influencers and adds that "Deals between big brands and viral online video performers, once an informal alternative to traditional celebrity sponsorships, are quickly maturing into a business estimated to reach $10 billion in 2020." While returns on investment of influencer marketing are difficult to measure, the Jonny Was brand reports a return of 226% on its RYPLpowered influencer marketing campaign (eMarketer 2018 influencer marketing roundup).

[3]A two-sided market is one in which the participants on each side care about the number of participants on the other side, so that there are bilateral network externalities. The two sides are intermediated by one or more platforms, which typically compete for business from both sides. Most applications treated in the literature (such as credit cards) are concerned with positive externalities (see Rochet and Tirole 2003, Armstrong 2004, and references

costs in creating, say, an account on Instagram and in engaging on a topic of interest, it is not surprising that, in contrast to traditional media markets, the market for influence consists of many small platforms. The proliferation of influencers also gives rise to frictions in the process of matching between followers and influencers. Finally, once influencers engage with marketers, followers will be less confident of the value of their recommendation, which may decrease engagement.

Our model formalizes these considerations. There are many influencers, each providing recommendations about a specific topic. They differ in their ability to create content. This heterogeneity is exogenous and it differentiates influencers vertically. It could be interpreted as heterogeneity in skills in writing, taking pictures, understanding market trends, and so on. It could also be interpreted as heterogeneity in "status" such as being a celebrity or not. We will refer to this vertical dimension as intrinsic ability. Each influencer produces a fixed amount of recommendations. An influencer chooses how many should be marketer–sponsored; the other–organic–recommendations are based on his genuine assessments.

Once a follower is matched to a particular influencer, her utility is the sum of the intrinsic quality of the influencer and the value that she obtains from his recommendations. We assume that each organic recommendation points the follower to a marketer selling a high–quality good and that each sponsored recommendation points the follower to a marketer selling, in expectation, lower–quality goods. Influencers who supply few sponsored posts therefore provide, on average, better product recommendations. We start with the model in which followers cannot tell if a recommendation is organic or sponsored; in other words, influencers' advice is opaque.

The heterogeneity in intrinsic ability across influencers, together with their choice of sponsored recommendations, leads to a ranking of influencers based on the utility they provide to followers. We assume that followers are more likely to be matched to highly ranked influencers, but there are frictions, which we parametrize with a single parameter $\alpha \geq 0$. When α converges to zero, the matching technology becomes purely random in the sense that the matching of followers to influencers is independent of utility. As α becomes greater, the matching is more and more effective in allocating followers to the highly ranked influencers. In the limit of α approaching infinity, the technology matches all followers to the best influencer around.[4]

There are different rationales for these search frictions. Notably, almost all online searches are mediated by search or curation algorithms and the technology can be imperfect.[5] Our

therein). An important analytical issue is equilibrium multiplicity due to expectations of consumers about where the consumers on the other side of the market are going (see Caillaud and Jullien 2001 and 2003). In our case, as in many models of traditional media (see Section 5 for a review of that literature), the externalities imposed by marketers are negative and therefore equilibrium multiplicity is not an issue.

[4]Formally, the matching function we use is a variant of the urn-ball matching function. Influencers are the "balls" but some balls are more prominent than others; that is, are more likely to be drawn. A ball's prominence depends positively on the utility the influencer provides to followers, relative to other influencers.

[5]The importance of search frictions in online markets gave rise to an active area of research in the intersection of marketing, management, computer science, and economics. We briefly review briefly this literature in Section 5. We refer the reader to Tadelis (2016) for a survey of feedback systems in online platforms and a discussion of the possible bias in feedback systems and reputation systems.

underlying assumption is that the matching technology tries to match followers to influencers who provide the highest quality services (measured in terms of utility provided to followers) and that the parameter α captures underlying imperfections in the technology. Search frictions could also be manifestations of an adverse selection problem that does not allow for clear ranking across influencers. In this case, a low α represents institutions that are not able to screen influencers with respect to the utility they provide to followers, whereas a high α describes institutions that efficiently screen influencers. Finally, high search costs for followers –or any form of followers' behavioral bias that prevents them from screening influencers– could also be modeled using low α.

For a marketer, the value of having an influencer endorsing its product depends on how many followers the influencer has and on the surplus that the marketer can obtain from each follower. We assume that there are many, perfectly competitive marketers who compete for sponsored content. Hence, the price a marketer pays to an influencer for a sponsored post is such that the marketer makes zero profit and reflects accurately the value of a recommendation.

In equilibrium, each influencer chooses the amount of sponsored recommendations to maximize his profit, taking as given the choice of other influencers and taking into account that the price per sponsored post is determined by the zero-profit conditions of marketers.[6] This simple economic model of influence leads to a basic trade-off that is at the core of all our results. When an influencer chooses the profit-maximizing level of sponsored posts he anticipates that, on the one hand, an increase in the fraction of sponsored posts will have a direct positive effect on his revenue but, on the other hand, such an increase will decrease followers' utilities and, therefore, reduce his number of followers and their willingness to pay for products he recommends. Ultimately, this second effect will lower the price that marketers are willing to pay for his recommendations.

Our first set of results characterizes the unique equilibrium of this model. We show that influencers with higher intrinsic ability supply more sponsored posts, have more followers, and receive a higher per-post price from marketers, but a lower per-reader-per-post price. The intuition for these results relies on the observation that the competition between influencers for followers is imperfect. A marginal increase in the influencer's intrinsic ability is therefore passed through to followers' utility only in part, whereas he extracts part of the remaining surplus by increasing his amount of sponsored content. The fact that the per-post-per-reader price declines in the intrinsic ability of an influencer reflects this imperfect competition.

Because the market for influence is decentralized, accurate data on the pricing of sponsored posts are unavailable. However, our predictions on how equilibrium prices depend on the number of followers are supported by suggestive evidence. TRIBE (www.tribegroup.co), a website connecting brands with influencers, publishes recommended per-post price ranges based on the number of followers an influencer has. This is summarized in Table 1. Per-post

[6] We assume that an influencer only cares about the profit from sponsoring content. That is, the profit to an influencer is the product of the price per post and the fraction of sponsored posts. The analysis extends to the case in which influencers are also intrinsically motivated; e.g., they also obtain some utility from the fact that they have followers.

Followers	Ballpark rate (AUD$)
3-10K	$75-150
10-25K	$150-220
25-50K	$220-350
50-100K	$350-500
100K+	$500+

Table 1 – TRIBE suggested price-per-post (https://www.tribegroup.co/faq)

prices are increasing in the number of followers, but the per-post-per follower price declines in the number of followers.[7] Fascinatingly, this is the opposite of the effect in traditional media (e.g., the price per viewer for a TV ad is the highest during the Super Bowl; see Chwe 2001), suggesting that our model captures an important and unique feature of the market for online influence.

More generally, our equilibrium analysis provides a micro-foundation for the so called "rise of the micro-influencers". Micro-influencers are influencers who reach a sizable number of followers, usually between 3000 and 10,000, and they seem able to remain credible in the eyes of their followers and generate engagement with their products' recommendation. This is in contrast with celebrities who have many more followers, yet often generate much lower engagement.[8] In our model, micro-influencers are influencers with low to intermediate intrinsic abilities. They obtain a sizable audience because the search frictions and their choice of fewer sponsored posts allows them to catch up with celebrities in the competition for followers. Posting fewer sponsored posts also implies that, for a follower, a micro-influencer's recommendation has more value than a celebrity's, because followers understand that celebrity is more tempted to post sponsored content. This pushes up the price for posting via micro-influencers.

We find that the equilibrium allocation of sponsored content, which supports this rise of micro-influencers, is inefficient, over providing sponsored content by high-intrinsic-ability influencers such as celebrities, and under providing sponsored content by low-intrinsic-ability influencers (that is, micro-influencers). A more efficient allocation of sponsored content will have only low-ability influencers post sponsored recommendations, making them even less appealing to followers and making it easier for the search technology to separate out the

[7] A similar pattern is captured in the following example from an article titled: "How Bloggers Make Money on Istagram" on harpersbazaar.com: "[...] if you have hundreds of thousands of followers you can make anywhere from $500 to $5,000 a post, but if you have upwards of 6 million followers, your fee can be $20,000 to $100,000 a shot."

[8] The following extract from eMarketer's "Influence Marketing Roundup" (August 2017) summarizes common views amongst practitioners on the importance of micro-influencers: "...some brands are seeing more success with micro-influencers that have up to 10,000 followers and middle influencers with up to 250,000 followers." "Working with celebrities has become significantly less effective," said Gil Eyal, founder of influencer marketing platform HYPR Brands."Their engagement rates are typically lower than similar influencers with a smaller audience." Why? "Consumers can see right through it. A post from a person with millions of followers about a brand they've never talked about before seems disingenuous," said Mallorie Rosenbluth, head of social media at food delivery service GrubHub." A 2018 eMarketer report finds that 59% of marketers who use influencer marketing incorporate micro-influencers (50-25,000 followers) in their campaigns, as opposed to only 44% who use macro-influencers (over 100,000 followers).

low-ability from the high-ability influencers.

Our second set of results clarifies the market implications of a change in the search technology's efficiency; that is, a change in α. Following our initial interpretation, we think of an increase in α as an improvement of the underlying online institution that allows followers to better screen the influencers. The key observation is that an increase in search-technology efficiency generates more competition between influencers because small differences in the utility they provide to followers now create large differences in the numbers of followers, which translate to large differences in per-post prices. We then show that, when competition for followers between influencers become fiercer, influencers reduce their supply of sponsored content, and the distribution of followers becomes more skewed towards high-ability influencers. Both forces lead to an increase of followers' aggregate surplus and total surplus.

Finally, we derive results on the effect of transparency regulations on the market for influence. Recently, competition and media authorities in different countries–such as the US Federal Trade Commission (FTC) in 2016, AGM (Italy's state competition authority) in 2017, and Landesmedienanstalten (Germany's state media authority) in 2017– have instructed influencers to clearly indicate sponsored content. This is in line with regulation forbidding covert ads in traditional media.[9] The direct effect on consumers of introducing transparency is positive: They can follow up organic recommendations and, if they have better outside options, they can ignore sponsored recommendations. However, transparency also affects the market interaction between influencers, followers, and marketers. Our model allows for the evaluation of the strategic effects of transparency regulations.

When transparency is introduced, the price that an influencer receives to endorse a brand becomes less elastic with respect to his choice of sponsored content. This is so for two reasons. First, when recommendations are transparent, an increase in the fraction of sponsored posts has less effect on the utility of followers, as they can now exercise their outside options instead of following recommendations they know to be sponsored. In turn, the influencer's audience– and therefore the per-post price– become less sensitive to the composition of sponsored versus organic content. Second, an increase in the fraction of sponsored posts no longer affects followers' assessments of the authenticity of sponsored recommendations, which are now known to be sponsored. As a result, followers' willingness to pay for recommended products and marketers' willingness to pay per-follower-per-sponsored-post become less sensitive to the composition of sponsored versus organic content.

Since the equilibrium per-post price is less elastic with respect to the level of sponsored content under transparency, the marginal cost for an influencer to supply an additional sponsored post is lower. Therefore, the introduction of transparency increases the supply of sponsored posts by all influencers, decreases the aggregate consumers' utility and total surplus, and thus

[9] This is the FTC's rationale for its intervention, taken from its *Endorsements Guides*: "Say you're planning a vacation. You do some research and find a glowing review on someone's blog that a particular resort is the most luxurious place he has ever stayed. If you knew the hotel had paid the blogger hundreds of dollars to say great things about it or that the blogger had stayed there for several days for free, it could affect how much weight you'd give the blogger's endorsement. The blogger should, therefore, let his readers know about that relationship."

confounds the positive direct effect of the policy. We show that this strategic effect is not second-order. Transparency decreases followers' surplus and total surplus when the search technology is either sufficiently inefficient or sufficiently efficient. For intermediate regions, we have not been able to show analytically that transparency decreases followers' surplus and total surplus, but our numerical analysis suggests that the result extends.[10]

Our paper relates to the literatures on (a) advertising provision in traditional media, (b) curation algorithms and news aggregation in online platforms, and (c) advertisement and advice. We review the contribution of our paper with respect to existing work in Section 5. The next section develops the model, Section 3 characterizes the equilibrium and derives comparative statics, and Section 4 assess policy interventions. Proofs are relegated to the Appendix.

2 Model: The market for influence

The market for online influence consists of followers, online influencers, and marketers. Influencers create online content in the form of product recommendations. Followers consume this online content which influences their purchase decisions. Marketers pay influencers to sponsor their products, with the hope of boosting their demand. The influencers' segment is composed of many, relatively small platforms. As a consequence, the matching between followers and influencers is determined by a search technology that entails frictions. In what follows, we describe all these elements in detail and provide a parsimonious model to study the market for influence.

2.1 Influencers

There is a continuum of influencers of mass 1. Influencers are vertically differentiated; their abilities $\theta = (\theta_i)_{i\in[0,1]}$ are uniformly distributed on $[0, 1]$ and commonly known. Each influencer creates a fixed amount of recommendation content, normalized to 1. Recommendation content describes products and services and so provides, more or less explicitly, purchase recommendations to consumers. For simplicity, a unit of content corresponds to a unit of distinctive products' recommendations.

The content includes organic content and sponsored content. Organic content is the result of the influencer's experience with a specific product. It is a "genuine" influencer's assessment and generates no revenue for the influencer. Sponsored content is, in contrast, the result of an exchange between the influencer, who is paid to post the content, and the marketer, who tries

[10] The idea that transparency in economic interactions may trigger unwanted strategic effects is also studied in principal-agent models. In the canonical model of Holmstrom (1982), transparency always improves aggregate welfare. However, Prat (2005), and Dasgupta, Prat, and Verardo (2011) show that in a career-concern model, in which an agent type is her ability to correctly identify the optimal action, transparency over agents' actions may lead to conformity and suboptimal action choices. Those authors' motivation, analysis, and underlying forces are very different from ours.

to reach out to the influencer's followers. As we formalize in the next section, followers value organic recommendations more than sponsored recommendations.

The strategy of influencer i specifies the fraction of sponsored content s_i and the fraction of organic content $1 - s_i$. We denote by $\mathbf{s}$ the strategy profile of all influencers.

2.2 Followers

There is a unit mass of identical followers. The matching between followers and influencers is described in the next subsection. For the moment, consider a follower who is matched to influencer i who produces content s_i. We postulate that the follower's utility from this match is:

$$q_i(s_i) = \theta_i + [s_i\tau + (1 - s_i)](1 - \beta). \tag{1}$$

The follower benefits from the influencer's ability, θ_i; this summarizes the idea that highly able influencers express content in a way that is more enjoyable and useful to read. The second part of the expression reflects the value for the follower of the influencer's recommendations. The formulation models the following market structure: Each unit of organic content directs the follower to a marketer who sells a high-quality product. Each unit of sponsored content directs the follower to a marketer who has probability τ of selling a high-quality product and probability $1 - \tau$ of selling a low-quality product. The follower is willing to pay 1 for high-quality products and 0 for low-quality products. Hence, the expected valuation for a product reached via influencer i's recommendation is $s_i\tau + (1 - s_i)$. Finally, we assume that the consumer gets a fraction $1 - \beta$ of this expected surplus and that the marketer extracts the rest. Throughout the analysis, we maintain the assumption that $\tau < 1/2$; this assumption assures interiority of equilibrium and it is made only for expositional reasons. Section 2.6 discusses, in detail, the functional form of followers' utility.

2.3 Search technology

We assume that followers are more likely to be matched with influencers who provide greater utility. The degree of this correlation depends on the efficiency of the search technology. Our model aims to capture, in a static way, a reputation mechanism that aggregates followers' experiences and uses them to create high-quality matching. Formally, the fraction of followers matched to influencer i, given $\mathbf{s}$ and θ, is:

$$n_i(\mathbf{s}) = \frac{q_i(s_i)^\alpha}{q(\mathbf{s})},$$

where $q(\mathbf{s}) = \int_j [q_j(s_j)]^\alpha d_j$. The search technology is a variant of the classical urn-ball matching function, in which influencers take the role of balls. Here, in contrast with the standard urn-ball matching function, balls have a difference prominence and the prominence of influencer i depends on the utility he provides to followers.

The extent to which higher follower's utility of an influencer translates into higher prominence depends on $\alpha \geq 0$, which captures the "efficiency" of the search technology. When α is close to zero, the ranking of influencers with respect to $\{q_i\}$ is irrelevant for the matching. Hence the distribution of followers across influencers is uniform. This describes a situation in which influencers who provide better services to followers are not distinguishable from influencers who provides mediocre services. At the other extreme, when α tends to infinity, each follower is matched, with very high probability, with the highest-ranked influencer; that is, the influencer with higher q_i. There are different rationales for these search frictions and we refer the reader to our discussion in the Introduction and to the related literature, discussed in Section 5.

2.4 Marketers and influencers' payoffs

Marketers approach influencers to advertise their products. When a marketer approaches influencer i, the marketer observes his ability θ_i and his choice s_i. An agreement between a marketer and influencer i determines a price p_i that the marketer pays to the influencer and the influencer's commitment to create content concerning the product sold by the marketer. The profit to marketer m who contracts with influencer i for sponsored content at a price p_i is:

$$V_{m,i}(\mathbf{s}, p_i) = n_i(\mathbf{s})[s_i\tau + 1 - s_i]\beta - p_i.$$

The marketer's revenue from the sponsored post depends on the number of followers of influencer i's, $n_i(\mathbf{s})$, the willingness to pay of each follower who will meet the marketer via the influencer, $s_i\tau + (1 - s_i)$, and the bargaining power of the marketer *viz.* the followers, β.

The profit to influencer i with ability θ_i, given a price p_i per-sponsored post is:

$$\Pi_i(s_i, p_i) = p_i s_i.$$

One could extend influencers' objectives by introducing intrinsic motivation in providing content.[11]

2.5 Timing and definition of equilibrium

The decisions taken by market participants and the timing of these decisions is summarized as follows:

- First, influencers choose the level of sponsored and organic content out of a unit of content. These choices are summarized by a strategy profile $\mathbf{s}$, which specifies the supply of sponsored content s_i of each influencer i. The profile $\mathbf{s}$ is observed by all market participants.

[11] A natural extension is to consider that influencers directly care about the number of followers they have, so that $\Pi_i(\mathbf{s}, p_i) = p_i s_i + \omega n_i(\mathbf{s})$. As we will see, even when $\omega = 0$, influencers care about readership size because it increases the price they receive per sponsored post. The additional terms obtained when $\omega > 0$ does not affect the qualitative results and conclusions we present.

- Second, each influencer is approached by many, perfectly competitive marketers. Each marketer observes the influencer's ability and the supply of sponsored content and purchases sponsored posts at a market price. Because marketers are competitive, the price that influencer i receives per sponsored post is pinned down (in equilibrium) by a zero-profit condition on marketers' profits.
- Third, the market clears, sponsored and organic content is created, and the search technology matches followers to influencers. Followers want to buy a unit of products. A follower matched to influencer i purchases products following i's recommendations. The surplus division between marketers and followers follows the simple reduced form we describe above.

Definition 1. An equilibrium is a strategy profile $\mathbf{s}^*$ and a profile of prices $\mathbf{p}^*$ such that

1. The profile of price $\mathbf{p}^*$ is such that marketers obtain zero-profit; that is, for each influencer i, the price p_i^* is such that $V_{m,i}(\mathbf{s}^*, p_i^*) = 0$.
2. For every influencer i, the fraction of sponsored content s_i^* maximizes his profit (expression 2.4) given the supply of sponsored content $\mathbf{s}_{-i}^*$.

2.6 Discussion

We make a number of assumptions. We comment on these assumptions and elaborate on the role they play in the analysis that follows.

Observability. We assume that marketers and followers observe influencers' decisions $\mathbf{s}$ when they make their decisions. First, the idea behind the assumption that marketers observe s_i is that marketers react quickly to changes in the number of followers $n_i(\mathbf{s})$ due to a change in s_i. Consequently, the price that a marketer pays to influencer i depends on the realized number of followers of influencer i, which is consistent with empirical evidence.[12] Second, the idea behind the assumption that followers observe s_i is that followers quickly adjust their expected value of each influencer's recommendation when s_i changes. It aims at capturing an underlying (un-modelled) reputation mechanism: If an influencer deviates by creating more sponsored posts, as compared to what followers have conjectured, consumers will systematically experience a lower utility than expected when visiting the influencer. We are assuming that this information propagates more quickly to other consumers than the influencer can change the supply of sponsored posts.[13] In line with this assumption, anecdotal evidence and discussions amongst practitioners suggests that influencers balance sponsored content with organic content in order to avoid losing reputation and trust among followers.

We will comment further on this assumption when we examine the equilibrium conditions in Section 3. We emphasize that our results and insights hold if we relax the assumption that followers observe s_i, as long as marketers can pay influencers based on their realized number

[12] An equivalent model is one in which marketers do not observe s_i but marketers and influencers can write contracts that are contingent on the influencer's realized number of followers.

[13] Formally, we are assuming that this information propagates instantaneously, but this is just for convenience.

of followers. If both followers and marketers do not observe s_i, the model would predict that all influencers only produce sponsored content, which is counterfactual.

Follower's utility. The way in which we model the interaction between marketers and followers via influencers is in line with traditional models of the media market, such as Anderson and Coate (2005). The utility that a follower obtains when matched to influencer i is the sum of the influencer's ability, θ_i, and the value to the consumers of his recommendation, $[1 - s_i(1-\tau)](1-\beta)$, which is decreasing in the fraction of sponsored recommendations. For simplicity, we have assumed that the influencer's ability enters linearly in followers' utility and does not affect the value to the consumers of his recommendation. This latter assumption means that θ_i affects the interaction between followers and marketers only via the choice of sponsored content. From a formal point of view, it means that the cross partial derivative of q_i with respect to s_i and θ_i is zero. In the online Appendix, we extend our equilbrium analysis to the following case:

$$q_i(s_i) = v(\theta_i) + [s_i\tau + 1 - s_i](1-\beta)\mu(\theta_i),$$

where $\mu(\theta_i)$ and $v(\theta_i)$ are increasing function of θ_i. In this case, the value that consumers obtain by being matched to marketers via influencer i is increasing in the intrinsic ability of influencer i. This could reflect, for example, that high-θ influencers curate their content better. If we interpret θ as a parameter indicating status/celebrity, the function $\mu(\theta_i)$ indicates consumers' preference to buy products recommended by celebrities.

We also note that, in reality, it is possible for influencers to invest more or less effort in researching and curating recommendations. Curation effort decisions would enter our model though their effect on followers' utility. We get that

$$q_i(s_i, \nu_i) = \theta_i + [s_i\tau + (1-s_i)\nu_i](1-\beta),$$

where ν_i is the influencer's costly effort in curating organic recommendations. We note that the effect of an increase in ν_i have the opposite effect from an increase in s_i and that, qualitatively, our results extend to influencers' curation effort decisions with a change of sign.

Opaque content. We assume that influencers do not disclose whether a specific product's recommendation is sponsored or not. In this sense, the content produced by influencers is opaque for followers. As we discussed in the Introduction, competition authorities in different countries have recently introduced new legislation or emphasised the application of existing legislation to influencer marketing, with the attempt to increase market transparency. We study the effect of influencers' transparency in Section 4.

3 Equilibrium analysis

In equilibrium, the price that a marketer pays to influencer i with ability θ_i is such that marketers obtain zero profit. Formally:

$$V_{m,i}(\mathbf{s}, p_i) = n_i(\mathbf{s})[s_i\tau + 1 - s_i]\beta - p_i = 0,$$

which is captured by the following price correspondence:

$$p_i(\mathbf{s}) = n_i(\mathbf{s})[1 - s_i(1-\tau)]\beta. \tag{2}$$

Consider influencer i, who expects that all other influencers will follow strategy $\mathbf{s}_{-i}$, and that the price per sponsored post is determined by Expression 2. The profits to influencer i are:

$$\begin{aligned}\Pi_i(\mathbf{s}) &= p_i(\mathbf{s})s_i \\ &= n_i(\mathbf{s})[1 - s_i(1-\tau)]\beta s_i.\end{aligned}$$

The influencer will select $s_i \in [0,1]$ in order to:

$$\max_{s_i} n_i(\mathbf{s})[1 - s_i(1-\tau)]\beta s_i.$$

We have that:

$$\frac{\partial \Pi_i(\mathbf{s})}{\partial s_i} = \underbrace{n_i(\mathbf{s})[1 - s_i(1-\tau)]\beta}_{\uparrow \text{ Sponsored content's revenue}} + \overbrace{\underbrace{\frac{\partial n_i(\mathbf{s})}{\partial s_i}[1 - s_i(1-\tau)]s_i\beta}_{\downarrow \text{ Number of followers}} - \underbrace{n_i(\mathbf{s})s_i\beta(1-\tau)}_{\downarrow \text{ Recommendations' surplus}}}^{\downarrow \text{ Price per posted content}} \tag{3}$$

Influencer i's marginal profits can be decomposed into three terms, as described in Expression 3. When influencer i substitutes organic content with sponsored content, his revenue increases. This is the marginal benefit to the influencer of increasing s_i (first term of Expression 3). The marginal cost of increasing s_i is the decrease in the price per sponsored post. The remaining two terms describes this effect. First, an increase in s_i decreases the utility that followers receive from influencer i. His number of followers therefore declines, so the price marketers are willing to pay influencer i goes down (second term of Expression 3). Second, an increase in s_i decreases how much followers value influencer i's recommendation, so marketers' revenue from advertising via influencer i goes down. In turn, this pushes down the price p_i (last term of Expression 3).

We remark that if we relax the assumption that s_i is observed by marketers and followers, then p_i would be based on the conjecture that marketers and followers have about s_i. Hence, an increase in s_i from the conjectured level would not affect p_i. It follows that, for any possible equilibrium conjecture $s_i < 1$, the marginal costs for influencer i to increase s_i would be zero. In this case, the only equilibrium would be that influencers only produce sponsored content; i.e., $s_i = 1$ for all i. This would also imply that the price per reader that an influencer gets per sponsored post is constant across influencers; i.e., $p_i/n_i = \tau\beta$. Both these predictions are counterfactual. This reiterates the point we made in Section 2.6: the assumption that s_i is observed captures a form of reputation mechanism that is key to disciplining influencers to produce organic content.

Note that each influencer has negligible effect on the aggregate distribution of the utility that influencers provide to followers. That is, the choice of influencer i does not alter $q(\mathbf{s}) = \int_j [q_j(s_j)]^\alpha dj$. Influencers with the same ability therefore face the same maximization problem. Since their payoffs are strictly concave in the amount of sponsored content, in equilibrium, influencers with the same ability will adopt the same strategy.[14] In a symmetric profile $\mathbf{s}$, we denote by $s(\theta)$, $q(\theta)$, and $n(\theta)$ the corresponding $(s_i, q_i(s_i), n_i(\mathbf{s}))$ for each influencer i with $\theta_i = \theta$. Developing Expression 3, we obtain that the equilibrium conditions for a symmetric and interior equilibrium read:

$$\beta n(\theta) \left[\underbrace{[1 - s(\theta)(1-\tau)]}_{\uparrow \text{ Sponsored content's revenue}} - \overbrace{\left(\underbrace{s(\theta)\alpha(1-\tau)\frac{(q(\theta)-\theta)}{q(\theta)}}_{\downarrow \text{ Number of followers}} + \underbrace{s(\theta)(1-\tau)}_{\downarrow \text{ Recommendations' surplus}} \right)}^{\downarrow \text{ price per-posted content}} \right] = 0$$

We obtain the following characterization:

Proposition 1. *There exists a unique equilibrium, which is interior. In equilibrium, every influencer with ability θ selects $s^*(\theta)$ that satisfies:*

$$q^*(\theta)[1 - 2s^*(\theta)(1-\tau)] = \alpha(1-\tau)(1-\beta)[1 - s^*(\theta)(1-\tau)]s^*(\theta), \tag{4}$$

and receives from marketers a price per sponsored content equal to $p^*(\theta) = n^*(\theta)[1 - s^*(\theta)(1 - \tau)]\beta$.

Figure 1 illustrates equilibrium condition 4. The LHS is a decreasing function of $s(\theta)$ and is positive for all $s(\theta) \leq \frac{1}{2(1-\tau)}$; here, the assumption that $\tau < 1/2$ is sufficient to assure that equilibrium decisions are interior. The RHS is increasing in $s(\theta)$ and equals 0 at $s(\theta) = 0$.

We can then use the equilibrium characterization to provide the following predictions:

Proposition 2. *In equilibrium, the following properties hold:*

1. *Influencers with higher abilities select higher levels of sponsored content, generate higher utility for followers, and have more followers, i.e., $s^*(\theta)$, $q^*(\theta)$, and $n^*(\theta)$ are all increasing in θ.*
2. *Influencers with higher abilities command higher prices per sponsored post, but lower prices per follower per sponsored post; i.e., $p^*(\theta)$ is increasing in θ, but $p^*(\theta)/n^*(\theta)$ is decreasing in θ.*

[14] This result extends to a discrete version of our model. In fact, influencers' actions are strategic complements, which is sufficient to imply that all equilibria must be symmetric.

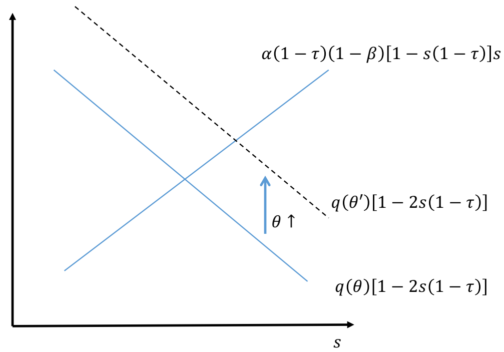

Figure 1 – Illustration of equilibrium condition; $\theta' > \theta$.

Since the matching technology favors influencers who provide higher utility, influencers with higher θ have a comparative advantage in attracting followers. This means that the elasticity of the number of followers, with respect to sponsored content, decreases in the influencer's ability. So, the fraction of sponsored posts increases in θ. More formally, notice that the ability parameter θ is a shift parameter of the LHS of equilibrium condition 4: higher θ shifts the LHS of condition 4 rightward, but the RHS does not change (see also Figure 1). Although more able influencers post more sponsored content, they still provide a greater utility to followers; i.e., $q^*(\theta)$ is still increasing in θ. This, then, implies that equilibrium readership size is positively correlated with true influencer's ability—see Figure 2a.

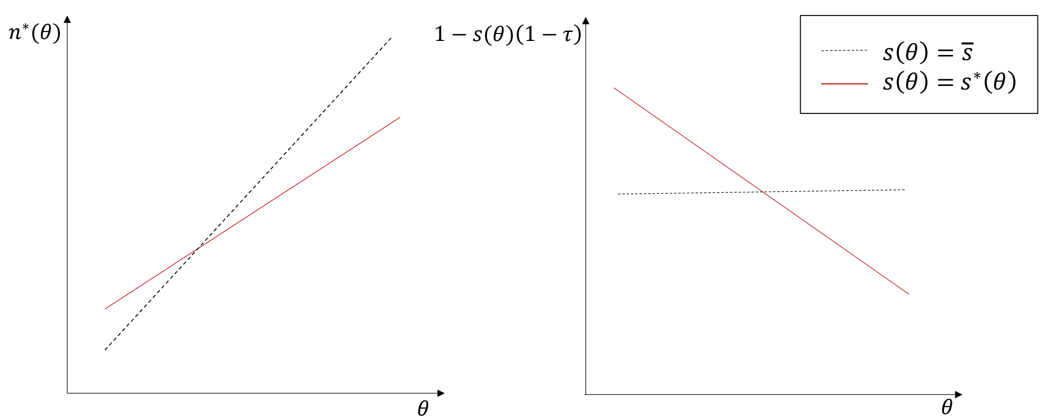

(a) Equilibrium readership size is positively correlated with influencer's ability, but less so than if all influencers have the same fraction of sponsored content.

(b) High-ability influencers supply more sponsored content. Their recommendations are therefore less valuable to followers

Figure 2 – The rise of the micro-influencers

To understand the conclusion about the price, note that, in equilibrium, the price received by an influencer with ability θ is equal to:

$$p^*(\theta) = n^*(\theta)[1 - s^*(\theta)(1-\tau)]\beta.$$

On the one hand, a marketer is willing to pay a higher price to more able influencers because they have larger readerships; on the other hand, a high-ability influencer supplies more sponsored content, so his followers value his recommendation less (see Figure 2b). Overall, price is increasing in θ, but the price per follower declines with the influencer's ability.

Another way to understand this is by looking at the pass-through to followers' utility for a marginal increase in influencer's ability. Influencers compete for followers but this competition is not perfect. A marginal increase in θ is therefore passed through to followers only in part; the other part is extracted by influencers who increase the amount of sponsored content. When α goes to infinity, the market for influence becomes competitive, influencers' level of sponsored content becomes negligible (we will state this formally in Section 3.2) and, as a consequence, a marginal increase in productivity translates into the same marginal increase in followers' utility.

Proposition 2 supports the idea that the key players in the market for influence are the micro-influencers. These are the influencers who reach many followers and yet are able to be credible and, therefore, to engage their followers. They obtain a sizeable number of followers because the search frictions and their choice of less sponsored content allows them to catch up with celebrities in the competition to attract followers. In addition, the value for a follower of following a micro-influencer's recommendation is higher than the value of following a celebrity's recommendation, as followers understand that a celebrity is more tempted to sponsor posts. This pushes up the price for posting via micro-influencers. As we discussed in the Introduction, all of this is consistent with existing evidence.

3.1 Inefficiencies

To explore equilibrium inefficiencies, we compare equilibrium outcomes with the outcomes generated by allocating sponsored content across influencers with the goal of maximizing total surplus, $s^{FB}(\theta)$, while maintaining the assumption that influencers, marketers, and followers interact according to the market rules.[15] We do not view this exercise as normative. Indeed, it is difficult to think how a planner could systematically control the supply of sponsored content of each influencer. The aim of the analysis is to illustrate natural distortions that emerge in the market for influence relative to a clear efficient benchmark.

For a symmetric allocation of sponsored content $s(\theta)$ and corresponding readership $n(\theta)$, the expression for total surplus is:

$$TS = \int n(\theta)[\theta + 1 - s(\theta)(1 - \tau)]d\theta.$$

Proposition 3. *Consider a planner choosing the allocation of sponsored content $s(\theta)$ to maximize total surplus. There exists $\widetilde{\theta}$ such that for every $\theta < \widetilde{\theta}$, $s^{FB}(\theta) = 1$ and for every $\theta > \widetilde{\theta}$, $s^{FB}(\theta) = 0$. Furthermore, $\widetilde{\theta}$ is increasing in α, tends to 0 as α tends to 0, and tends to 1 as*

[15] In the online Appendix, we characterize the optimal allocation of sponsored content across influencers for two alternative objectives: followers' welfare and influencers' surplus.

α *tends to infinity.*

An obvious inefficiency in the market for influence is that the supply of sponsored content substitutes for the more valuable organic content. There is also a technological inefficiency: because of frictions, followers are not always matched to the influencers that provide the highest utility to them. This technological inefficiency is reinforced by influencers' equilibrium choices. The fact that highly able influencers post more sponsored content implies that influencers become less heterogeneous with respect to the utility they offer to their followers. As demonstrated in Figure 2(a), this makes the assignment of followers to influencers even less efficient.

Proposition 3 shows that the best way to alleviate these inefficiencies is to tolerate positive levels of sponsored content and use it in a way that amplifies the heterogeneity across influencers with respect to their utility to followers. Specifically, the planner allocates sponsored content to low-ability influencers and only organic content to high-ability influencers. In this way, high-ability influencers provide a much greater utility to followers and low-ability influencers a much lower utility to followers. By creating these strong asymmetries, followers are directed to high-ability influencers and away from low-ability influencers. This effect is illustrated in Figure 3.

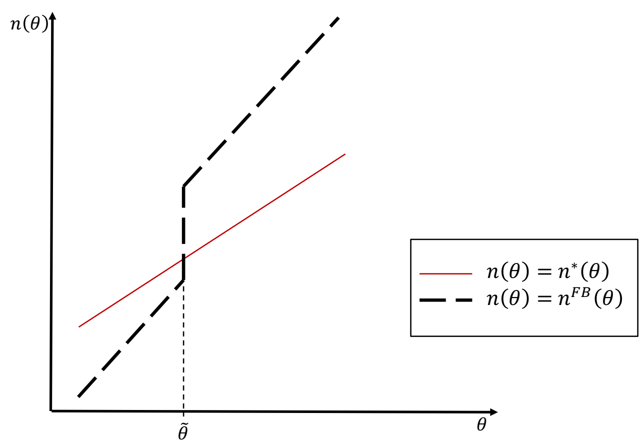

Figure 3 – Efficient readership distribution *viz.* equilibrium readership distribution

When the matching technology is very inefficient, creating these strong asymmetries does not improve the matching outcome substantially, so the planner will tolerate only a small amount of sponsored content; i.e., $\tilde{\theta}$ is low when α is low. In contrast, when the matching technology is highly efficient, creating such asymmetries is cost-effective because most followers will be assigned to influencers with only organic content; i.e., $\tilde{\theta}$ is high when α is high.

3.2 The role of the search technology

We now study the implications for the market for influence of a change in the efficiency of the search technology; i.e., a change in α. As we discussed in the Introduction and in Section

2.3, an increase in α can be thought of as an improvement of the underlying online institution that allows followers to better screen the influencers.

Proposition 4. *An increase in the efficiency of the search technology decreases the sponsored content for each influencer, thus increasing the utility that each influencer provides to his followers; i.e., $s^*(\theta)$ decreases and $q^*(\theta)$ increases for all θ. Furthermore, as the search technology becomes random ($\alpha \to 0$), the fraction of sponsored content of each influencer converges to $1/[2(1-\tau)]$, whereas as the search technology becomes perfectly efficient ($\alpha \to \infty$), influencers create only organic content; i.e., $s^*(\theta)$ converges to 0 for all θ.*

Proposition 4 points out that the efficiency of the search technology determines the level of competition. An increase in efficiency creates more competition amongst influencers. As α increases, an influencer can gain more followers by reducing sponsored content, thus offering greater utility to followers. As his number of followers goes up, the influencer can demand a higher rate per sponsored post. These extreme market structures emerge when α converges to 0 or to infinity.

When α tends to 0, the distribution of followers across influencers does not depend on content choices. It is as if each influencer has a loyal base of followers. In this case, each influencer acts as a monopolist and the only cost of supplying more sponsored content is that followers will value the influencer's recommendation less. At the other extreme, when α converges to infinity, an influencer can get all followers by offering the highest utility. Influencers have, therefore, strong incentives to undercut each other's level of sponsored content. As influencers compete á la Bertrand for followers, in equilibrium, they produce only organic content and the equilibrium outcome becomes efficient.

We next consider the effect of an increase in α on followers' aggregate welfare, influencers' surplus, and total surplus. The expressions for followers' welfare and influencers' welfare are, respectively:

$$W_F = \int n^*(\theta)q^*(\theta)d\theta \text{ and } W_I = \frac{\beta}{1-\beta}\int n^*(\theta)[q^*(\theta)-\theta]s^*(\theta)d\theta.$$

Proposition 4 says that an increase in search-technology efficiency improves the utility that a follower obtains when matched to an influencer. However, it also affects the distribution of followers across influencers and, since influencers provide different utility to followers, this affects the different measures of welfare. The following result, also illustrated in Figure 4, determines how the distribution of followers across influencers changes.

Proposition 5. *Consider an increase in the efficiency of the search technology. There exists a $\overline{\theta} \in (0,1)$ such that the equilibrium number of followers of influencers with ability $\theta > \overline{\theta}$ increases and the equilibrium number of followers of influencers with ability $\theta < \overline{\theta}$ decreases.*

It is useful to think of $n^*(\cdot)$ as a density function for θ. Then, Proposition 5 implies the following corollary:

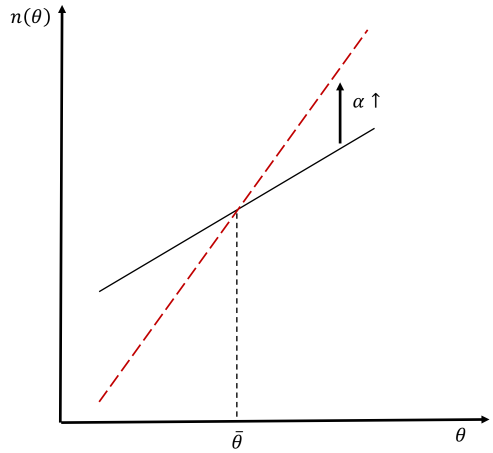

Figure 4 – The effect of an increase in the search-technology efficiency on the readership redistribution.

Corollary 1. *An increase in the search-technology efficiency leads to a first-order stochastic dominance (FOSD) shift in the distribution of followers across influencers with ability* $\theta \in [0, 1]$.

When α increases, there are both a direct and a strategic effect. The direct effect is that, since influencers with higher θ have higher $q^*(\theta)$, they will increase their number of followers at the expense of low-θ influencers. The indirect effect comes from the fact that each influencer decreases the amount of sponsored content, which changes the slope of $q^*(\theta)$. When we start from a high α, the highest-ability influencers compete fiercely for followers. An increase in α, then, leads high-θ influencers to reduce their level of sponsored content the most. Hence, the function $q^*(\theta)$ becomes steeper, which implies that the search technology will direct followers more often to highly able influencers. In this case, the indirect effect complements the direct effect. However, when α is low to begin with, an increase in α leads low-θ influencers to decrease sponsored content the most. In this case, the profile $q^*(\theta)$ becomes flatter, which means that the search technology is, in essence, more noisy. This effect crowds out, at least in part, the effect of an increase in α. Proposition 5 points out that this strategic effect is always smaller than the direct effect.

Aggregating the effects captured in Corollary 1 and Proposition 4, we get the following result:

Proposition 6. *An increase in the search-technology efficiency increases aggregate followers' equilibrium welfare and total equilibrium surplus.*

Because an increase in α leads to an FOSD shift in the distribution of followers (Corollary 1) and because $q^*(\theta)$ is increasing in θ, followers' welfare increases. In addition, when α increases, influencers lower $s^*(\theta)$, so $q^*(\theta)$ increases for all θ (Proposition 4), which reinforces the former effect so that, overall, followers' welfare increase. The same intuition and logic applies to total welfare.

The effect of an increase in α on influencers' aggregate profits is ambiguous. Note that

$[q^*(\theta) - \theta]s^*(\theta) = (1 - s^*(\theta)(1 - \tau))\, s^*(\theta)\, (1 - \beta)$ is an increasing function in θ. So, the first-order stochastic shift of $n^*(\cdot)$, due to an increase in α, leads to an increase in W_I. On the other hand, an increase in α decreases $[q^*(\theta) - \theta]s^*(\theta)$ for every θ, which decreases W_I. This ambiguity is reflected in the fact that influencers with higher θ are better off when α increases, but influencers with low θ are worse off.

Proposition 7. *There exists $\bar{\bar{\theta}} \in (0, 1]$ such that for every $\theta \leq \bar{\bar{\theta}}$, the profit for an influencer with ability θ is decreasing in α and for every $\theta > \bar{\bar{\theta}}$, the profit for an influencer with ability θ is increasing in α.*

4 Content transparency

This section extends the benchmark model to study the effect of content transparency. We refer to the Introduction for examples of several competition authorities that have instructed influencers to clearly indicate sponsored content.

A rationale for enforcing transparency in the content published by influencers is to prevent the bundling of sponsored and organic content, thus allowing followers to ignore sponsored recommendations. With transparent content, followers recognize that the value of a sponsored recommendation is $(1 - \beta)\tau$; a follower will ignore such recommendation when she has a more valuable outside option. To introduce this effect, we assume that followers are heterogeneous with respect to an outside option c, which is distributed according to some distribution F in the support $\left[0, \frac{1}{2}(1 - \beta)\right]$. Only followers with an outside option $c < \tau(1 - \beta)$ will follow sponsored recommendations; followers will always follow organic recommendations (as their value is $1 \times (1 - \beta)$).

Remark 1. The restriction that $c \leq \frac{1}{2}(1 - \beta)$ implies that, in our benchmark model (without transparency), followers never exercise their outside option. Hence, we can understand the implication of transparency by comparing the equilibrium outcomes under transparency an in our benchmark model. We denote by $\hat{x}$ a variable of interest in the model with transparency and by x the same variable prior to the intervention. ‖

Let $\gamma \in [0, 1]$ be the probability that $c < \tau(1 - \beta)$. Let $C = \frac{1}{1-\beta} \int_{\tau(1-\beta)}^{\frac{1}{2}(1-\beta)} c dF$ noting that $C \in [\tau, 1/2]$. The expected utility of a follower matching with influencer i (prior to the realization of the outside option) is:

$$\hat{q}_i(s_i) = \theta_i + (1 - s_i)(1 - \beta) + s_i[\tau\gamma + (1 - \gamma)C](1 - \beta).$$

For the same level of sponsored content s_i, the utility a follower obtains from influencer i under transparency, $\hat{q}_i(s_i)$, is larger than the utility she obtains prior to the intervention (Expression 1). This reflects the *direct benefit* of transparency: it allows the follower to substitute sponsored recommendations with more profitable alternatives.

The profit, to a marketer, from advertising a product via influencer i is:

$$\hat{V}_{m,i}(\hat{\mathbf{s}}, \hat{p}_i) = \gamma \hat{n}_i(\hat{\mathbf{s}})\tau\beta - \hat{p}_i.$$

For the same level of sponsored content and for the same price per post, the profit to a marketer from advertising the product via influencer i is lowered by the introduction of transparency. This is because, under transparency, only a fraction γ of followers will follow the sponsored recommendation and each now believes that the product has an expected value of τ. Perfect competition among marketers for influencer i implies that, in equilibrium, $\hat{V}_{m,i}(\hat{\mathbf{s}}, \hat{p}_i) = 0$. The equilibrium price under transparency and, for comparison, the equilibrium price prior to the intervention are:

$$\underbrace{\hat{p}_i(\hat{\mathbf{s}}) = \hat{n}_i(\hat{\mathbf{s}})\gamma\tau\beta}_{\text{price with transparency}} \quad \text{and} \quad \underbrace{p_i(\mathbf{s}) = n_i(\mathbf{s})[s_i\tau + 1 - s_i]\beta}_{\text{price prior to intervention}}$$

The price of influencer i is, under transparency, less sensitive to the level of sponsored content, relative to the pre-intervention case. There are two important effects. First, when content is transparent, an increase in the fraction of sponsored content affects the price marketers are willing to pay only via the change in the influencer's number of followers and not through followers' willingness to pay for the marketers' products. Second, under transparency, the utility a follower obtains from an influencer is less sensitive to an increase in the amount of sponsored content provided by the influencer because a follower with good enough outside options will not follow sponsored recommendations. This, in turn, implies that the influencer's number of followers decreases less when sponsored content increases. We get the following result:

Proposition 8. *Under transparency, there is a unique and symmetric equilibrium. The equilibrium level of sponsored posts is increasing in θ and equals:*

$$\hat{s}(\theta) = \min\left\{\frac{\theta + 1 - \beta}{(1-\beta)(\alpha+1)(1-\gamma\tau-(1-\gamma)C)}, 1\right\}. \tag{5}$$

Furthermore, the introduction of transparency increases the equilibrium level of sponsored content for each influencer; that is, $\hat{s}(\theta) > s^*(\theta)$ for all θ.

The overall implication of the transparency policy is that *the price that influencer i obtains to sponsor content becomes less elastic with respect to s_i*. This is the key determinant of the strategic effect of transparency. In fact, recall that the profit to influencer i with ability θ_i, when he expects that all other influencers follow strategy $\hat{\mathbf{s}}_{-i}$ is:

$$\hat{\Pi}_i(\hat{\mathbf{s}}, \hat{p}_i) = \hat{p}_i(\hat{\mathbf{s}})\hat{s}_i.$$

As transparency decreases the elasticity of $\hat{p}_i(\hat{\mathbf{s}})$ with respect to the amount of sponsored content $\hat{s}_i$, the intervention will increase the amount of sponsored content each influencer posts.

Figure 5 illustrates the equilibrium profile of sponsored content under transparency. There are three regions. When the search technology is inefficient, $(1+\alpha)(1-\gamma\tau-(1-\gamma)C) \leq 1$, competition amongst influencers is weak, so they all provide only sponsored content (Figure 5(a)). When the efficiency of the search technology is intermediate, $(1+\alpha)(1-\gamma\tau-(1-$

$\gamma)C) \in \left(1, \frac{2-\beta}{1-\beta}\right)$, competition amongst low-ability influencers becomes strong, which leads to a threshold configuration: there exists a $\breve{\theta} \in (0,1)$ so that influencers with ability $\theta < \breve{\theta}$ choose organic and non-organic content, whereas influencers with $\theta \geq \breve{\theta}$ only select sponsored content (Figure 5(b)). When the search technology is sufficiently efficient, the equilibrium profile of sponsored content is interior (Figure 5(c)).

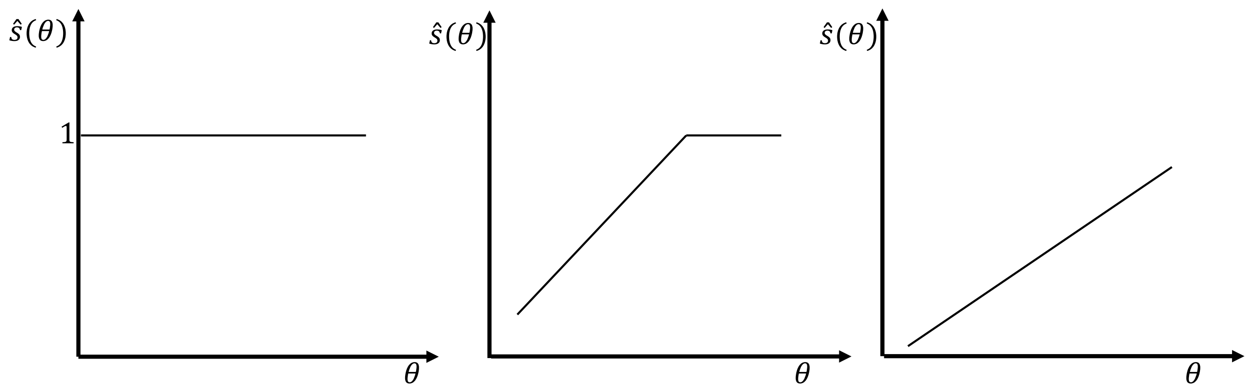

(a) Small α (b) Medium α (c) Large α

Figure 5 – Equilibrium sponsored content under transparency

4.1 The welfare effect of content transparency

Our interest lies in determining the effect of the policy on welfare. We have noticed that the policy has a direct effect of increasing the utility of followers matched to influencer i; that is, $\hat{q}_i(\hat{s}_i) > q_i(s_i)$ when $\hat{s}_i = s_i$. Yet introducing the policy also has a strategic effect: each influencer chooses a higher level of sponsored content (see Proposition 8). This strategic effect alters followers' utility in two distinctive ways. First, it decreases the utility a follower obtains when matched to influencer i, as now the influencer supplies more sponsored content; this effect confounds the direct positive effect. Second, the assignment of followers to influencers also changes. It becomes more efficient when the profile $\hat{q}(\theta)$ is steeper than the same profile prior the intervention, $q^*(\theta)$, and it becomes less efficient when the profile $\hat{q}(\theta)$ is flatter than the same profile prior the intervention, $q^*(\theta)$.

The following proposition pins down these two strategic effects. Let $\breve{\theta} \in [0,1]$ be the equilibrium threshold under transparency, such that high-ability influencers post only sponsored content and low-ability influencers post at least some organic content.[16]

Proposition 9. *Relative to the case without transparency, the introduction of transparency implies that:*

1. *A follower's utility from a match to an influencer with ability θ decreases; that is, $q^*(\theta) > \hat{q}(\theta)$ for all θ.*

[16]Formally, the influencer $\theta = \breve{\theta}$ so that $\hat{s}(\breve{\theta}) = 1$ and $\hat{s}(\theta) < 1$ for all $\theta < \breve{\theta}$; if such $\breve{\theta}$ does not exist, then set $\breve{\theta} = 1$.

2. *The decrease in a follower's utility from a match to an influencer with ability θ, $\Delta(\theta) = q^*(\theta) - \hat{q}(\theta)$, increases in θ for all $\theta \leq \breve{\theta}$, whereas it decreases in θ for all $\theta \geq \breve{\theta}$.*

The first part of the proposition implies that a follower matched to an influencer i is worse off under transparency. Hence, the equilibrium response of influencers, which consists of increasing the supply of sponsored post, erodes all the direct benefits associated with content transparency. The second part of the proposition is more subtle. When the search technology is inefficient (α has a small value), we know that $\breve{\theta} = 0$. In this case, the profile $\hat{q}(\theta)$ is steeper than $q^*(\theta)$. This implies that the distribution of readership under transparency *first order stochastic dominate* (FOSD) the distribution of readership prior the intervention; that is, the introduction of transparency improves the efficiency of the assignment between followers and influencers. When α is high, we know that $\breve{\theta} = 1$ and so, in this case, $\hat{q}(\theta)$ is flatter than $q^*(\theta)$. In this case, transparency implies that the assignment of followers to influencers becomes less efficient. For an intermediate α, we have that $\breve{\theta} \in (0, 1)$, so $q^*(\theta)$ is steeper than $\hat{q}(\theta)$ for $\theta < \breve{\theta}$ and shallower for $\theta > \breve{\theta}$. In this case, transparency decreases readership of influencers with θ around $\breve{\theta}$ while increasing readership of influencers remote from $\breve{\theta}$. This effect is demonstrated in Figure 6.

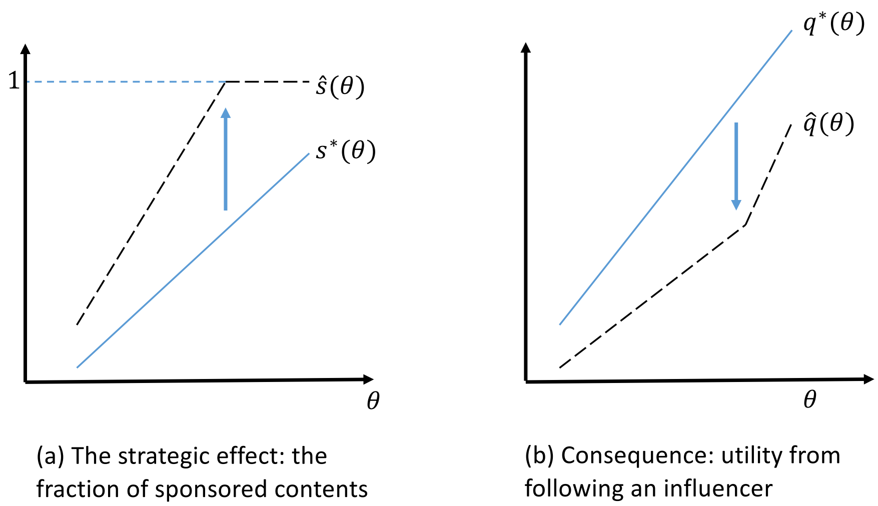

Figure 6 – The effect of transparency

These considerations provide insights on how followers' welfare and total surplus are affected by the introduction of the policy. In particular, part 1 of Proposition 9 implies that, regardless of the level of α and keeping the distribution of readership constant across influencers, followers' welfare and total surplus decrease due to transparency. Next, when α is sufficiently low, the assignment of followers to influencers is, both with and without transparency, (roughly) uniform. Hence, changes in the ranking across influencers do not affect welfare much and we can therefore conclude that, for low level α the introduction of transparency hurts both followers' and aggregate surplus. When α is large, because followers' welfare and total surplus both increase with FOSD shifts in the readership distribution and because, by part 2 of

Proposition 9, when α is large, the distribution of readership before transparency FOSD the one after transparency, transparency hurts both followers' and aggregate surplus.

Proposition 10. *There exists a* $0 < \underline{\alpha} < \overline{\alpha}$ *so that if* $\alpha < \underline{\alpha}$ *or* $\alpha > \overline{\alpha}$*, then the introduction of transparency decreases followers' welfare and total surplus.*[17]

We have not been able to prove that the introduction of transparency decreases followers' welfare and total surplus for intermediate levels of α. In that range, the distribution of followers across influencers under transparency *second order stochastic dominate* (SOSD) the distribution of readership before the intervention. This shift may lead to more efficient matches between followers and highly able influencers. We have solved numerically the effect of transparency on followers' and total surplus for the regions of parameters in which we do not have an analytical proof.[18] The numerical analysis suggests that the negative effect of transparency also applies to intermediate levels of α. Figure 7 plots total surplus and followers' surplus before and after the intervention for different values of α. The shapes of the relevant welfare functions before and after the intervention are representative of what we find in the numerical analysis.

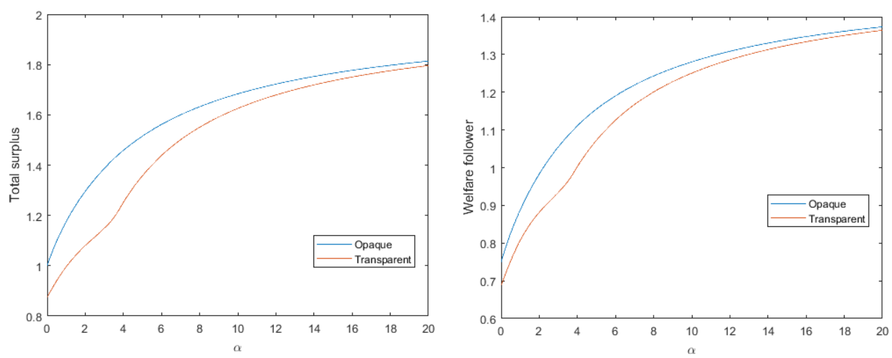

Figure 7 – The welfare effect of transparency, $\beta = 1/2$, $\tau = 1/3$, $C = 1/2$, $\gamma = 1/2$

5 Related literature

Traditional media. There is an established literature on content and advertising provision in traditional media. Anderson and Coate (2005) study competition in broadcasting, Peitz and Valletti (2007) compare pay-TV and free-to-air media platforms, and Wilbur (2005) provides an empirical model of television advertising and estimates viewers' aversion to advertising; see Anderson and Gabszewicz (2006) for a survey. More recent work, such as Ambrus et al.

[17]Moreover, we know that $\overline{\alpha}$ is smaller than the minimum level of α where the equilibrium under transparency is interior; that is, $\overline{\alpha}$ is smaller than the α that solves $(1+\alpha)(1-\gamma\tau-(1-\gamma)C) = \frac{2-\beta}{1-\beta}$.

[18]We note that $\overline{\alpha}$ is bounded from above by a value that is uniquely determined by γ, C, τ and β (see Footnote 17). Our numerical analysis consisted of an extensive search in the admissible space of (γ, C, τ, β). For each point selected in this set, we checked the effect of transparency for α ranging from 0 to the upper bound of α. The code for this analysis is available upon request from the authors.

(2016) and Athey et al. (2016), studies media competition in advertising markets with multi-homing users. In this literature, content (programming) has entertainment value and does not include organic recommendations. Advertising is, therefore, never considered authentic and is modelled using a nuisance cost function.[19] In contrast, an important channel in our model is that followers are interested in influencers' recommendations and sponsored content may or may not be "hidden" amongst organic recommendations.

A second significant difference is that traditional media markets are concentrated and therefore the literature has focused on oligopolistic competition; that is, there are only a few platforms matching the two sides of the market. Online influencers, in contrast, have low entry costs leading to an environment in which (a) influencers are abundant and (b) search frictions are important in shaping the competition amongst them.

Online frictions. The importance of search frictions on online markets has stimulated research at the intersection of management science, computer science, and economics. Search frictions in social media lead many platforms to develop curation algorithms to help populate consumers' feeds. Curation algorithms are, in essence, selection and ranking algorithms that help users (followers) search for the most relevant content (influencers). Early papers study how to better design such algorithms (see also Shardanand and Maes 1995 and Linden et al. 2003), whereas the more recent literature studies their effect on the content produced (see Latzer et al. 2016 for a survey). For example, Berman and Katona (2016) consider the impact of three curation algorithms on the quality of content created by producers. Su et al. (2016) analyze Twitter's "Who to follow" system that gives users suggestions for which other users to follow. Our comparative statics, with respect to search-technology efficiency, aim at capturing this technological innovation and studying its interaction with market forces.

In related literature on news aggregators, Dellarocas et al. (2013) and Roos et al. (2015) find that one effect of news aggregators is increased competition amongst content creators' websites. We find a parallel of that effect when we analyse the effects of improvements in the search and matching technologies. Athey and Mobius (2012), Chiou and Tucker (2015), and Calzada and Gil (2017) find empirical support for the hypothesis that news aggregators serve as a complement to news websites and that they are especially beneficial to niche content providers. In the market we study, all content providers (influencers) are niche and our analysis shows that an improvement in the search technology, akin to a better aggregator, may benefit the high-quality content providers, but hurt the low-quality ones.

Advertising and advice. Perhaps most related to our work is literature, dating back to Brin and Page (1998), on the conflict between advertising and advice on the Internet. Brin and Page (1998) focus on search engines and highlight the difficulty of having unbiased advertising-funded search engines. Mitchell (2017) formalizes this idea in the context of a dynamic relationship between an influencer and a follower. The influencer chooses a mix of organic and

[19] Even when advertising is not considered a nuisance, its benefits are not studied in the context of recommendation authenticity: Gabszewicz, Laussel, and Sonnac (2001) present a model of newspaper competition in which newspaper readers do not find advertisements a nuisance because ads can be ignored in a written medium. Rysman (2004) studies a model of the market for Yellow Pages directories in which readers like advertisements.

sponsored posts to maximize advertising income and intrinsic utility from being influential, whereas the follower forms expectations with respect to the authenticity of the influencer's recommendations and seeks to maximize surplus from product purchases. We take a reduced-form approach to model the long-term relationship between a follower and an influencer and instead focus on analyzing the equilibrium in a market with many influencers, followers, and marketers.

The work on the conflict between advertising and advice on the Internet is predated by an extensive literature that points out the limitations of free advice, often paid for by commissions and kickbacks, as in the case of physicians' recommendations of drugs and treatments and of advice given by financial intermediaries (see Inderst and Ottaviani 2012 and references therein). Our paper focuses on the competition between intermediaries (influencers), whereas the aforementioned literature focuses on (a) competition between marketers to be recommended by an advisor (e.g., Inderst and Ottaviani 2012) or (b) the direct relationship between a powerful advisor and a marketer (e.g., Fulghieri et al. 2014). These modeling choices reflect differences between the underlying markets, including (a) the number and accessibility of advisors/influencers, (b) the way consumers choose to take advice from a financial advisor or physician as opposed to an Instagramer, and (c) the resulting market power of advisors/influencers.

6 Conclusion

We develop a model of market interactions between influencers, followers, and marketers. It provides testable predictions on the joint distribution of price per sponsored post and numbers of followers and detects a novel source of inefficiency in this market. We then study how an improvement in the technology that matches followers to influencers affects these market outcomes. Finally, we use the model to assess how recent competition and media authorities' interventions in these markets affect market interactions and outcomes.

A layer of the market, from which we abstracted, involves the interactions of influencers, followers, and marketers with the platforms hosting them. Influencers are two-sided platforms bringing together followers and marketers. However, influencers and followers are also hosted by a third party. For concreteness, consider the platform Instagram. It does not charge influencers and followers and does not get a cut of the fee that influencers receive from marketers. Rather, its business model is to obtain revenue from display advertisements that marketers place directly on the platform.[20] Hence, the relationship between Instagram and its clients is complex. On the one hand, Instagram competes with the influencers it hosts for attracting

[20] There is a large body of work on display advertising. Evans (2009) provides an early study of the market structure of the online advertising industry, focusing on display advertising and the large platforms that provide it. Bergemann and Bonatti (2011) provide a more comprehensive analysis of the market for offline and online ads, taking into account online markets' greater ability to target audiences, and Goldfarb and Tucker (2011a, 2011b) find that targeting online display advertisements is highly effective. More recent work on ads targeting includes Deng and Mela (2018) and references within. Followup studies consider the effect of ad skipping and ad-blocking (see also Kumar 2018 and Tuchman et al. 2018).

advertising revenue from marketers. On the other hand, the attractiveness of Instagram for marketers is related to the presence of influencers and followers, while the attractiveness of influencers for marketers depend on the quality of the Instagram platform. We believe that our basic model could be extended to these under-studied interactions.

References

Ambrus, Attila, Emilio Calvano, and Markus Reisinger (2016). Either or Both Competition: A Two-sided Theory of Advertising with Overlapping Viewerships. *American Economic Journal: microeconomics*, 8(3), 189-222.

Anderson, Simon P. and Stephen Coate (2005). Market Provision of Broadcasting: A Welfare Analysis. *The Review of Economic Studies*, 72(4), 947-972.

Anderson, Simon P. and Jean J. Gabszewicz (2006). The Media and Advertising: A Tale of Two-Sided Markets, Editor(s): Victor A. Ginsburg, David Throsby, Handbook of the Economics of Art and Culture, Elsevier, Volume 1, 2006, Chapter 18, 567-614.

Armstrong, Mark (2006). Competition in Two-Sided Markets. *The Rand Journal of Economics*, 37(3), 668-691.

Athey Susan and Markus Mobius (2012). The impact of news aggregators on internet news consumption: The case of localization. Mimeo.

Athey Susan, Emilio Calvano, and Joshua S. Gans (2018). The Impact of Consumer Multi-homing on Advertising Markets and Media Competition. *Management Science*, 64(4), 1574-1590.

Bergemann, Dirk, and Alessandro Bonatti (2011). Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics*, 42(3), 417-443.

Berman, Ron and Katona, Zsolt (2016). The Impact of Curation Algorithms on Social Network Content Quality and Structure. NET Institute Working Paper No. 16-08.

Brin, Sergey and Lawrence Page (1998). The anatomy of a large scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1/7), 107-117.

Caillaud, Bernard and Jullien Bruno (2001). Competing Cybermediaries. *European Economic Review*, 45, 797-808.

Caillaud, Bernard and Jullien Bruno (2003). Chicken and Egg: Competition among Internet Service Providers. *The Rand Journal of Economics*, 34, 309-328.

Calzada, Joan and Ricard Gil (2017). What Do News Aggregators Do? Evidence from Google News in Spain and Germany. Working paper.

Chiou, Lesley and Catherine Tucker (2015). Content aggregation by platforms: The case of the news media. NBER Working Paper.

Chwe, Michael S.-Y. (2001). Rational Ritual: Culture, Coordination, and Common Knowledge. Princeton, NJ: Princeton University Press.

Dasgupta, Amil, Andrea Prat, and Michela Verardo (2011). Institutional Trade Persistence and Long-Term Equity Returns. *Journal of Finance*, 66(2), 635-653.

Dellarocas, Chrysanthos, Zsolt Katona, and William Rand (2013). Media, aggregators, and the link economy: Strategic hyperlink formation in content networks. *Management Science*, 59(10), 2360-2379.

Deng, Yiting and Carl F. Mela (2018). TV Viewing and Advertising Targeting. *Journal of Marketing Research*, 55(1), 99-118.

Evans, David S. (2008). The Economics of the Online Advertising Industry. *Review of Network Economics*, 7(3), 359-391.

Fulghieri, Paolo, Ganter Strobl, and Han Xia (2014). The Economics of Solicited and Unsolicited Credit Ratings. *The Review of Financial Studies*, 27(2), 484-518.

Gabszewicz, Jean J., Dider Laussel, and Nathalie Sonnac. (2001). Press Advertising and the Ascent of the 'Pensée Unique.' *European Economic Review*, 45(4-6), 641-651.

Goldfarb, Avi and Catherine E. Tucker (2011). Online Display Advertising: Targeting and Obtrusiveness. *Marketing Science*, 30(3), 389-404.

Goldfarb, Avi and Catherine E. Tucker (2011). Privacy Regulation and Online Advertising. *Management Science*, 57(1), 57-71.

Holmstrom, Bengt (1979). Moral Hazard and Observability. *Bell Journal of Economics*, 10, 74-91.

Rajeev Kumar (2018). Individual Incentive Based Reward Policy for an Ad Exchange to Mitigate the Threat of Ad Blocking. *International Journal of Information Systems and Social Change*, 9(3), 23-35.

Latzer, Michael, Katharina Hollnbuchner, Natascha Just, and Florian Saurwein (2016). The economics of algorithmic selection on the internet. In Handbook on the Economics of the Internet, Chapter 19, 395-425. Edward Elgar Publishing.

Linden, Greg, Brent Smith, and Jeremy York (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing*, 7(1), 76-80.

Mitchel Matthew (2017). Free (Ad)vice. Working paper.

Peitz, Martin, Tommaso M. Valletti (2008). Content and advertising in the media: Pay-tv versus free-to-air. *International Journal of Industrial Organization*, 26(4), 949-965.

Prat, Andrea (2005). The Wrong Kind of Transparency. *American Economic Review*, 95(3), 862-877.

Rochet, Jean-Charles, and Jean Tirole (2003). Platform Competition in Two-Sided Markets. *Journal of the European Economic Association*, 1(4), 990-1029.

Inderst, Roman, and Marco Ottaviani (2012). Competition through Commissions and Kickbacks. *American Economic Review*, 102(2), 780-809.

Roos, Jason M.T., Carl F. Mela, and Ron Shachar (2015). The effect of links and excerpts on internet news consumption. SSRN working paper 2678938.

Rysman, Marc (2004). Competition between Networks: A Study of the Market for Yellow Pages. *The Review of Economic Studies*, 71(2), 483-512.

Shardanand, Upendra and Pattie Maes (1995). Social information filtering: algorithms for automating "word of mouth". *Proceedings of the SIGCHI conference on Human factors in computing systems*, 210-217. ACM Press/Addison-Wesley Publishing Co.

Su, Jessica, Aneesh Sharma, and Sharad Goel (2016). The effect of recommendations on network structure. *Proceedings of the 25th International Conference on World Wide Web*, 1157-1167. International World Wide Web Conferences Steering Committee.

Tuchman, Anna E., Harikesh S. Nair, Pedro M. Gardete (2018). Television ad-skipping, consumption complementarities and the consumer demand for advertising. *Quantitative Marketing and Economics*, 16(2), 111-174.

Wilbur, Kenneth C. (2008). A Two-Sided, Empirical Model of Television Advertising and Viewing Markets. *Marketing Science*, 27(3), 356-378.

Appendix

Proof of Proposition 1. Recall that

$$\frac{\partial \Pi_i(\mathbf{s})}{\partial s_i} = \frac{\beta q_i(s_i)^{\alpha-1}}{q(\mathbf{s})}\left[q_i(s_i)(1-s_i(1-\tau)) - s_i\left[\alpha(1-\tau)(1-\beta)(1-s_i(1-\tau)) + (1-\tau)q_i(s_i)\right]\right]$$

and therefore $\frac{\partial \Pi_i(\mathbf{s})}{\partial s_i} \geq 0$ if, and only if,

$$q_i(s_i)(1-2s_i(1-\tau)) \geq s_i\alpha(1-\tau)(1-\beta)(1-s_i(1-\tau)).$$

Define $\hat{s} = \frac{1}{2(1-\tau)}$ and note that $\hat{s} \leq 1$ because $\tau \leq 1/2$. Note also that the LHS of the above inequality equals $\theta + (1-\beta)$ at $s_i = 0$ and 0 at $s_i = \hat{s}$, and it is decreasing in s_i. The RHS is 0 at $s_i = 0$, it equals $\alpha(1-\beta)/4$ at $s_i = \hat{s}$ and it is increasing in $s_i \in [0, \hat{s}]$. So, there is a unique solution s_i and $s_i \in (0, \hat{s})$.[21] So, influencer i with ability θ_i chooses s_i so that

$$q_i(s_i)(1-2s_i(1-\tau)) = s_i\alpha(1-\tau)(1-\beta)(1-s_i(1-\tau)).$$

Since there is a unique solution to this equation, influencers with the same θ will choose the same strategy. Hence, the equilibrium is symmetric. To conclude the proof of Proposition 1 we notice that the equilibrium price to influencers with ability θ is derived by the following zero profit condition $n(\mathbf{s})[s(\theta)\tau + b(\theta)]\beta - p(\theta) = 0$.

Proof of Proposition 2. First, to see that $s^*(\theta)$ is increasing in θ, consider equilibrium condition 4 and note that the LHS shifts up as θ increases. The RHS is independent of θ. Second, to see that $q^*(\theta)$ is increasing in θ, rewrite equilibrium condition 4 as follows

$$q^*(\theta) = \alpha(1-\tau)(1-\beta)\frac{[1-s^*(\theta)(1-\tau)]s^*(\theta)}{1-2s^*(\theta)(1-\tau)}. \tag{6}$$

The LHS is $q^*(\theta)$, and, holding $s^*(\theta)$ fixed, it shifts up when θ increases, whereas the RHS is independent of θ. Third, since $q^*(\theta)$ is increasing in θ, it follows immediately that $n^*(\theta)$ is increasing in θ. Fourth, $p^*(\theta)/n^*(\theta) = \beta[1-s^*(\theta)(1-\tau)]$ and since $s^*(\theta)$ is increasing in θ it follows that $p^*(\theta)/n^*(\theta)$ is decreasing in θ.

We conclude by showing that $p^*(\theta) = \beta n^*(\theta)[1-s^*(\theta)(1-\tau)]$ is increasing in θ. To see this note that the claim is true whenever $q^*(\theta)[1-s^*(\theta)(1-\tau)]$ is increasing in θ, and this follows by inspection of equilibrium condition 6: multiply the LHS and RHS of 6 by $[1-s^*(\theta)(1-\tau)]$ and then note that the LHS is increasing in θ and it is decreasing in $s^*(\theta)$ and that the RHS is independent of θ and it is increasing in $s^*(\theta)$. This concludes the proof of proposition 2.

Proof of Proposition 3. Recall that total surplus reads:

$$\begin{aligned} TS &= W^r + W^b + W^f = \int n(\theta)[\theta + s(\theta)\tau + b(\theta)]d\theta \\ &= \frac{\int q(\theta)^\alpha[\theta + 1 - s(\theta)(1-\tau)]d\theta}{\int_0^1 q(\theta)^\alpha d\theta}. \end{aligned}$$

[21] The case in which $\tau > 1/2$ will lead to a similar characterization. The only difference is that influencers with sufficiently high θ will only select sponsored content. We restrict the analysis to $\tau \leq 1/2$ so that we do not take into account the possibility of a corner solution for some influencers, and this makes the analysis easy to present.

Next, note that TS increases in $s(\theta')$ if

$$\frac{\frac{\partial\left(q(\theta')^{\alpha}[\theta'+1-s(\theta')(1-\tau)]\right)}{\partial s(\theta')}}{\int q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta} > \frac{\frac{\partial q(\theta')^{\alpha}}{\partial s(\theta')}}{\int_0^1 q(\theta)^{\alpha}d\theta}$$

or

$$\frac{\frac{1}{1-\beta}\left(\frac{1+\alpha}{\alpha}q(\theta')-\beta\theta'\right)}{\int q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta} < \frac{1}{\int_0^1 q(\theta)^{\alpha}d\theta}$$

and decreases otherwise. Note further that if the inequality holds for θ' for some $s(\theta')$ then it holds for any larger $s(\theta')$, and if the reverse inequality holds for θ' for some $s(\theta')$ then it holds for any smaller $s(\theta')$. Thus, for any θ' the planner will choose $s(\theta') \in \{0,1\}$, which maps to $q(\theta') \in \{\theta'+\tau(1-\beta), \theta'+(1-\beta)\}$.

To be more specific, the planner will choose $s(\theta')=0$ (equivalently, $q(\theta')=\theta'+(1-\beta)$) if

$$\frac{q(\theta')^{\alpha}[\theta'+1-s(\theta')(1-\tau)]|_{s(\theta')=0} - q(\theta')^{\alpha}[\theta'+1-s(\theta')(1-\tau)]|_{s(\theta')=1}}{\int_0^1 q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta} > \frac{q(\theta')^{\alpha}|_{s(\theta')=0} - q(\theta')^{\alpha}|_{s(\theta')=1}}{\int_0^1 q(\theta)^{\alpha}d\theta}$$

and choose $s(\theta')=1$ (equivalently, $q(\theta')=\theta'+\tau(1-\beta)$) otherwise. The condition simplifies to

$$\frac{(\theta'+(1-\beta))^{\alpha}[\theta'+1] - (\theta'+\tau(1-\beta))^{\alpha}[\theta'+1-(1-\tau)]}{\int_0^1 q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta} > \frac{(\theta'+(1-\beta))^{\alpha} - (\theta'+\tau(1-\beta))^{\alpha}}{\int_0^1 q(\theta)^{\alpha}d\theta}$$

or

$$\frac{(\theta'+(1-\beta))^{\alpha}[\theta'+1] - (\theta'+\tau(1-\beta))^{\alpha}[\theta'+1-(1-\tau)]}{(\theta'+(1-\beta))^{\alpha} - (\theta'+\tau(1-\beta))^{\alpha}} > \frac{\int_0^1 q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta}{\int_0^1 q(\theta)^{\alpha}d\theta} \tag{7}$$

where the RHS is independent of θ' and the LHS increases in θ'. To see that the LHS increases in θ' we write it as follows

$$1+\theta'+(1-\tau)\frac{(\theta'+\tau(1-\beta))^{\alpha}}{(\theta'+(1-\beta))^{\alpha} - (\theta'+\tau(1-\beta))^{\alpha}}$$

which is increasing in θ' if the following is decreasing in θ'

$$\frac{(\theta'+(1-\beta))^{\alpha} - (\theta'+\tau(1-\beta))^{\alpha}}{(\theta'+\tau(1-\beta))^{\alpha}}$$

or if

$$\frac{(\theta'+(1-\beta))}{(\theta'+\tau(1-\beta))}$$

is decreasing in θ', which is always the case because $\tau<1$. This completes the proof that there is a threshold $\widetilde{\theta}$ as required.

We next show how the threshold $\widetilde{\theta}$ changes with α. To show that the threshold $\widetilde{\theta}$ is increasing in α, it is sufficient to show that inequality 7 is easier to satisfy when α decreases. That will work if for example, the LHS decreases in α and the RHS increases in α. That the RHS increases in α is immediate because $q(\theta)^{\alpha}$ and $[\theta+1-s(\theta)(1-\tau)]$ both increase in θ. We next wan to show that

$$\frac{(\theta'+(1-\beta))^{\alpha}[\theta'+1] - (\theta'+\tau(1-\beta))^{\alpha}[\theta'+1-(1-\tau)]}{(\theta'+(1-\beta))^{\alpha} - (\theta'+\tau(1-\beta))^{\alpha}}$$

decreases in α.

$$\frac{\partial\left(\frac{(\theta+(1-\beta))^{\alpha}(\theta+1)-(\theta+\tau(1-\beta))^{\alpha}(\theta+1-(1-\tau))}{(\theta+(1-\beta))^{\alpha}-(\theta+\tau(1-\beta))^{\alpha}}\right)}{\partial\alpha}=$$

$$=(1-\tau)(\theta+\tau-\beta\tau)^{\alpha}(\theta-\beta+1)^{\alpha}\frac{\ln(\theta+\tau-\beta\tau)-\ln(\theta-\beta+1)}{\left((\theta-\beta+1)^{\alpha}-(\theta+\tau-\beta\tau)^{\alpha}\right)^{2}}$$

which is negative if and only if $\ln(\theta+\tau-\beta\tau)-\ln(\theta-\beta+1)<0$, or $(1-\beta)\tau<1-\beta$, which always holds.

Next, we show that if $\alpha\to 0$ then inequality 7 always holds. The inequality can be rewritten as follows

$$\frac{(\theta'+(1-\beta))^{\alpha}[\theta'+1]-(\theta'+\tau(1-\beta))^{\alpha}[\theta'+1-(1-\tau)]}{\int_0^1 q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta}>\frac{(\theta'+(1-\beta))^{\alpha}-(\theta'+\tau(1-\beta))^{\alpha}}{\int_0^1 q(\theta)^{\alpha}d\theta}$$

and substituting $\alpha=0$ we get

$$\frac{(1-\tau)}{\int_0^1[\theta+1-s(\theta)(1-\tau)]d\theta}>0,$$

which always holds.

Next, we show that if $\alpha\to\infty$ then inequality 7 holds only when $\theta'=1$. We begin by showing that inequality 7 holds for $\theta'=1$ regardless of α. When $\theta'=1$ the inequality becomes:

$$\frac{2(2-\beta))^{\alpha}-[1+\tau](1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha}-(1+\tau(1-\beta))^{\alpha}}>\frac{\int_0^1 q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta}{\int_0^1 q(\theta)^{\alpha}d\theta}$$

note that

$$\frac{\int_0^1 q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta}{\int_0^1 q(\theta)^{\alpha}d\theta}<\frac{\int_0^1 2q(\theta)^{\alpha}d\theta}{\int_0^1 q(\theta)^{\alpha}d\theta}=2$$

and

$$\frac{2(2-\beta)^{\alpha}-[1+\tau](1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha}-(1+\tau(1-\beta))^{\alpha}}>\frac{2(2-\beta)^{\alpha}-2(1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha}-(1+\tau(1-\beta))^{\alpha}}=2.$$

Finally, we show that

$$\lim_{\alpha\to\infty}\frac{\int_0^1 q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta}{\int_0^1 q(\theta)^{\alpha}d\theta}=2$$

and

$$\lim_{\alpha\to\infty}\frac{2(2-\beta)^{\alpha}-[1+\tau](1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha}-(1+\tau(1-\beta))^{\alpha}}=2.$$

Hence, if $\alpha\to\infty$ the inequality 7 holds only when $\theta'=1$.

To see the first limit

$$\begin{aligned}\lim_{\alpha\to\infty}\frac{\int_0^1 q(\theta)^{\alpha}[\theta+1-s(\theta)(1-\tau)]d\theta}{\int_0^1 q(\theta)^{\alpha}d\theta} &= \lim_{\alpha\to\infty}\frac{\ln(q(1))\,q(1)^{\alpha}[1+1]-\ln(q(0))\,q(0)^{\alpha}[1-(1-\tau)]}{\ln(q(1))\,q(1)^{\alpha}-\ln(q(0))\,q(0)^{\alpha}}\\ &= \lim_{\alpha\to\infty}\frac{2\ln(q(1))\,q(1)^{\alpha}-\tau\ln(q(0))\,q(0)^{\alpha}}{\ln(q(1))\,q(1)^{\alpha}-\ln(q(0))\,q(0)^{\alpha}}\\ &= 2\end{aligned}$$

where the last equlity holds because $s(0) = 1$ and $q(0) < 1$. We now prove the second limit

$$\begin{aligned}
\lim_{\alpha\to\infty} \frac{2(2-\beta)^{\alpha} - (1+\tau)(1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha} - (1+\tau(1-\beta))^{\alpha}} &= \lim_{\alpha\to\infty}\left(1 + \frac{(2-\beta)^{\alpha} - \tau(1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha} - (1+\tau(1-\beta))^{\alpha}}\right) \\
&= \lim_{\alpha\to\infty}\left(1 + \frac{1}{\frac{(2-\beta)^{\alpha} - (1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha} - \tau(1+\tau(1-\beta))^{\alpha}}}\right) \\
&= \lim_{\alpha\to\infty}\left(1 + \frac{1}{1 + \frac{-(1-\tau)(1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha} - \tau(1+\tau(1-\beta))^{\alpha}}}\right) \\
&= 1 + \lim_{\alpha\to\infty}\left(\frac{1}{1 + \frac{-(1-\tau)(1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha} - \tau(1+\tau(1-\beta))^{\alpha}}}\right) \\
&= 2
\end{aligned}$$

where the last inequality holds because

$$\lim_{\alpha\to\infty} \frac{-(1-\tau)(1+\tau(1-\beta))^{\alpha}}{(2-\beta)^{\alpha} - \tau(1+\tau(1-\beta))^{\alpha}} = 0$$

Proof of proposition 4. Consider equilibrium condition 4. The RHS increases in α (specifically, an increase in α rotates leftward the RHS), whereas the LHS does not change with α. Recalling that the LHS increases in $s^*(\theta)$ and the RHS decreases in $s^*(\theta)$, we get that $s(\theta)$ declines in α for all θ. Furthermore, as α goes to 0, the RHS goes to zero for all θ and so $s(\theta)$ goes to $\frac{1}{2(1-\tau)}$ for all θ. As α goes to ∞, the RHS goes to ∞ unless $s(\theta)$ goes to 0 for all θ. Next, by definition $q(\theta) = \theta + [1 - s(\theta)(1-\tau)](1-\beta)$; hence, an increase in α decreases $s(\theta)$ and therefore $q(\theta)$ increases for all θ.

Proof of proposition 5. Since $q^*(\theta)$ increases in θ, it follows that $n^*(\theta)$ is increasing in θ. Since total readership is fixed, it is sufficient to show that in equilibrium

$$\frac{\frac{dq^*(\theta)^{\alpha}}{d\alpha}}{q^*(\theta)^{\alpha}}$$

increases in θ, or equivalently, to prove that

$$\frac{\frac{dq^*(\theta)}{d\alpha}}{q^*(\theta)}$$

increases in θ (to see why this is equivalent, note that for any well behaved function f, $\frac{\frac{df(x)^{\alpha}}{dx}}{f(x)^{\alpha}} = \frac{\alpha f(x)^{\alpha-1}\frac{df(x)}{dx}}{f(x)^{\alpha}} = \alpha\frac{\frac{df(x)}{dx}}{f(x)}$).

Next, recall that in equilibrium

$$q^*(\theta) = \alpha(1-\tau)(1-\beta)\frac{[1 - s^*(\theta)(1-\tau)]s^*(\theta)}{1 - 2s^*(\theta)(1-\tau)}.$$

With some abuse of notation we denote $s = s^*(\theta)$ for the remainder of the proof when the dependence on θ is clear from the context. We get that

$$\begin{aligned}\frac{dq^*(\theta)}{d\alpha} &= (1-\tau)(1-\beta)\left(\frac{(1-s(1-\tau))\,s}{1-2s(1-\tau)}+\alpha\frac{\partial\left(\frac{(1-s)(1-\tau))s}{1-2s(1-\tau)}\right)}{\partial s}\frac{\partial s}{\partial\alpha}\right)\\ &= (1-\tau)(1-\beta)\left(\frac{(1-s(1-\tau))\,s}{1-2s(1-\tau)}+\alpha\frac{\left(2s^2\tau^2-4s^2\tau+2^2+2s\tau-2s+1\right)}{(2s\tau-2s+1)^2}\frac{\partial s}{\partial\alpha}\right)\end{aligned}$$

and

$$\begin{aligned}\frac{\frac{dq(\theta)}{d\alpha}}{q(\theta)} &= \frac{(1-\tau)(1-\beta)\left(\frac{(1-s(1-\tau))s}{1-2s(1-\tau)}+\alpha\frac{\left(2s^2\tau^2-4s^2\tau+2s^2+2s\tau-2s+1\right)}{(2s\tau-2s+1)^2}\frac{\partial s}{\partial\alpha}\right)}{\alpha(1-\tau)(1-\beta)\frac{[1-s(\theta)(1-\tau)]s(\theta)}{1-2s(\theta)(1-\tau)}}\\ &= \frac{\left((1-s(1-\tau))\,s+\alpha\frac{\left(2s^2\tau^2-4s^2\tau+2s^2+2s\tau-2s+1\right)}{(1-2s(1-\tau))}\frac{\partial s}{\partial\alpha}\right)}{\alpha\,(1-s(1-\tau))\,s}\\ &= \frac{1}{\alpha}+\frac{1}{(1-s(1-\tau))}\frac{\frac{\partial s}{\partial\alpha}}{s}+\frac{2s\,(1-\tau)^2}{(1-2s(1-\tau))\,(1-s(1-\tau))}\frac{\partial s}{\partial\alpha}\end{aligned}$$

We note that $\frac{2s(1-\tau)^2}{(1-2s(1-\tau))(1-s(1-\tau))}$ and $\frac{1}{(1-s(1-\tau))}$ increase in s and therefore in θ. Therefore, to prove that $\frac{\frac{dq(\theta)}{d\alpha}}{q(\theta)}$ increases in θ it is sufficient to show that $\frac{\frac{\partial s}{\partial\alpha}}{s}$ increases in θ.

To move forward we rewrite the equilibrium condition for s as follows:

$$(\theta+(s\tau+1-s)\,(1-\beta))\,(1-2s(1-\tau))=\alpha(1-\tau)(1-\beta)\,(1-s(1-\tau))\,s$$

and apply the implicit function theorem to get

$$\frac{\frac{\partial s}{\partial\alpha}}{s}=-\,(1-\beta)\,\frac{1-s\,(1-\tau)}{2\theta-4s+\alpha-3\beta-2s\alpha+4s\beta+4s\tau-\alpha\beta+2s\alpha\beta+2s\alpha\tau-4s\beta\tau-2s\alpha\beta\tau+3}$$

which is increasing in θ if

$$\frac{1-s\,(1-\tau)}{2\theta-4s+\alpha-3\beta-2s\alpha+4s\beta+4s\tau-\alpha\beta+2s\alpha\beta+2s\alpha\tau-4s\beta\tau-2s\alpha\beta\tau+3}$$

decreases in θ. We know that $1-s\,(1-\tau)$ decreases in θ, and therefore it is sufficient to show that

$$2\theta-4s+\alpha-3\beta-2s\alpha+4s\beta+4s\tau-\alpha\beta+2s\alpha\beta+2s\alpha\tau-4s\beta\tau-2s\alpha\beta\tau+3$$

increases in θ. This is true because $-4s+\alpha-3\beta-2s\alpha+4s\beta+4s\tau-\alpha\beta+2s\alpha\beta+2s\alpha\tau-4s\beta\tau-2s\alpha\beta\tau+3$ is increasing in s which increases in θ.

Proof of proposition 6. An increase in α leads to a first order stochastic shifts $n(\theta)$ (see Corollary 1). This observation together with the observation that $q(\theta)$ is increasing in θ (see Proposition 2), implies that W^r increases in α, keeping the function $q(\theta)$ constant. Furthermore, in view of proposition 4 we know that when α increases $q(\theta)$ increases for all θ. Hence, W^r increases even further. These two observations are easily adapted to total surplus.

Proof of proposition 7. We first note that

$$\Pi^*(\theta')=\frac{q^*(\theta')^\alpha[1-s^*(\theta')\,(1-\tau)]\beta s^*(\theta')}{\int_0^1 q^*(\theta)^\alpha d\theta}$$

Next, we show that

$$\Psi = \frac{\frac{d(q^*(\theta)^\alpha[1-s^*(\theta)(1-\tau)]\beta s^*(\theta))}{d\alpha}}{q^*(\theta)^\alpha[1-s^*(\theta)\,(1-\tau)]\beta s^*(\theta)}$$

is increasing in θ, which will imply that there is a threshold $\overline{\overline{\theta}} \in [0,1]$ such that for every $\theta \leq \overline{\overline{\theta}}$, the utility of a blogger with ability θ is decreasing in α, and for every $\theta > \overline{\overline{\theta}}$, the utility of a blogger with ability θ is increasing in α. To that end, with some abuse of notation we denote $s = s^*(\theta)$ and $q = q^*(\theta)$ and note that

$$\begin{aligned}\frac{d\left(q^\alpha[1-s\,(1-\tau)]\beta s\right)}{d\alpha} &= \beta\left(\alpha q^{\alpha-1}\frac{dq}{d\alpha}[1-s\,(1-\tau)]s - q^\alpha\frac{\partial s}{\partial\alpha}\,(1-\tau)\,s + q^\alpha[1-s\,(1-\tau)]\frac{\partial s}{\partial\alpha}\right)\\ &= \beta q^{\alpha-1}\left(\alpha\frac{dq}{d\alpha}[1-s\,(1-\tau)]s - q\frac{\partial s}{\partial\alpha}\,(1-\tau)\,s + q[1-s\,(1-\tau)]\frac{\partial s}{\partial\alpha}\right)\end{aligned}$$

and therefore

$$\begin{aligned}\Psi &= \frac{\alpha\frac{dq}{d\alpha}[1-s\,(1-\tau)]s + q[1-s\,(1-\tau)]\frac{\partial s}{\partial\alpha} - q\frac{\partial s}{\partial\alpha}\,(1-\tau)\,s}{q[1-s\,(1-\tau)]s}\\ &= \frac{\alpha\frac{dq}{d\alpha}[1-s\,(1-\tau)]s + q\frac{\partial s}{\partial\alpha}\,(1-2s\,(1-\tau))}{q[1-s\,(1-\tau)]s}\\ &= \alpha\frac{\frac{dq}{d\alpha}}{q} + \frac{\frac{\partial s}{\partial\alpha}\,(1-2s\,(1-\tau))}{[1-s\,(1-\tau)]s}.\end{aligned}$$

Next recall that in proposition 5 and its proof we showed that $\frac{\frac{dq(\theta)}{d\alpha}}{q(\theta)}$ and $\frac{\frac{\partial s}{\partial\alpha}}{s}$ increase in θ. Therefore, to show that Ψ increase in θ it is sufficient to show that $\frac{(1-2s(1-\tau))}{[1-s(1-\tau)]}$ increases in θ, which always holds because $\frac{(1-2s(1-\tau))}{[1-s(1-\tau)]}$ decreases in s.

To prove that the threshold is strictly positive we note that from proposition 5 we know that there exists $\overline{\theta} \in (0,1)$ such that $n^*(\theta)$ is increasing in α if $\theta > \overline{\theta}$ and decreasing otherwise. It is then sufficient to note that $[s(\theta)\tau + b(\theta)]\beta s(\theta)$ is increasing in s and thus decreasing in α.

Proof of Proposition 8. We start with the first part of the proposition. The equilibrium price to influencer i is

$$\hat{p}_i(\hat{\mathbf{s}}) = \gamma\hat{n}_i(\hat{\mathbf{s}})\tau\beta. \tag{8}$$

The profits to influencer i, by choosing $\hat{s}_i$, are

$$\hat{\Pi}_i(\hat{\mathbf{s}}) \quad = \quad \hat{p}_i(\hat{\mathbf{s}})\hat{s}_i = \gamma\hat{n}_i(\hat{\mathbf{s}})\tau\beta\hat{s}_i,$$

Influencer i selects $\hat{s}_i$ in order to

$$\max_{\hat{s}_i} \gamma\tau\beta\hat{n}_i(\theta,\hat{\mathbf{s}})\hat{s}_i$$

We have that

$$\frac{\partial\hat{\Pi}_i(\hat{\mathbf{s}},\hat{p}_i)}{\partial\hat{s}_i} = \gamma\tau\beta\hat{n}_i(\hat{\mathbf{s}}) + \gamma\tau\beta\frac{\partial\hat{n}_i(\hat{\mathbf{s}})}{\partial\hat{s}_i}\hat{s}_i. \tag{9}$$

In an interior equilibrium, influencer i with ability $\theta_i = \theta$ will select $\hat{s}_i = \hat{s}(\theta)$ so that $\frac{\partial\hat{\Pi}_i(\hat{\mathbf{s}})}{\partial\hat{s}_i}|_{\hat{s}_i=\hat{s}(\theta)} = 0$. Developing expression 3 for a symmetric an interior equilibrium we have that

$$\frac{\partial\hat{\Pi}_i(\hat{\mathbf{s}})}{\partial\hat{s}_i}|_{\hat{s}_i=\hat{s}(\theta)} = 0$$

if and only if

$$\theta_i + (1 - \hat{s}_i (1 - \gamma\tau - (1-\gamma) C)) (1-\beta) - \alpha \hat{s}_i (1 - \gamma\tau - (1-\gamma) C) (1-\beta) = 0. \tag{10}$$

Hence

$$\hat{s}(\theta) = \min\left\{ \frac{\theta + 1 - \beta}{(1-\beta)(\alpha+1)(1-\gamma\tau-(1-\gamma)C)}, 1 \right\}.$$

We now turn to the second part of the Proposition: We prove that $s^*(\theta) < \hat{s}(\theta)$ for all θ. Absent transparency $s^*(\theta) \in (0,1)$ for all θ. Hence, if for a specific θ, transparency leads to $\hat{s}(\theta) = 1$, then the claim holds for influencers with ability θ. Suppose, next, that after the policy $\hat{s}(\theta) \in (0,1)$ for some θ. We know that

$$\hat{s}(\theta) = \frac{\theta + 1 - \beta}{(1-\beta)(\alpha+1)(1-\gamma\tau-(1-\gamma)C)}.$$

Furthermore, from the FOC above,

$$\hat{q}(\theta) - \alpha\hat{s}(\theta)(1-\beta)(1-\gamma\tau-(1-\gamma)C) = 0,$$

and since $C \geq \tau$,

$$\hat{q}(\theta) - \alpha\hat{s}(\theta)(1-\beta)(1-\gamma\tau-(1-\gamma)C) > \hat{q}(\theta) - \alpha\hat{s}(\theta)(1-\beta)(1-\tau).$$

Hence

$$\hat{q}(\theta) - \alpha\hat{s}(\theta)(1-\beta)(1-\tau) < 0$$

Now, take the FOC prior intervention for the same θ influencer, we have that

$$[q(\theta) - \alpha s(\theta)(1-\beta)(1-\tau)][1 - s(\theta)(1-\tau)] - s(\theta)q(\theta)(1-\tau) = 0$$

but if we evaluate this at the post intervention $\hat{s}$ and so $\hat{q}(\theta)$ we see that the first term is negative and therefore the all expression is negative. Concavity of the objective function implies that the $s(\theta)$ prior intervention must be lower than the one post intervention.

Proof of Proposition 9. Recall that $\breve{\theta} \in [0,1]$. We first show that both of the proposition's claims are true for all $\theta \leq \breve{\theta}$, then we show that the claims are true for all $\theta \geq \breve{\theta}$.

Step 1. Consider a $\theta \leq \breve{\theta}$.

Step 1.a. We first derive an explicit expression $\hat{q}(\theta)$. Recall that

$$\hat{q}(\theta) = \theta + 1 - \beta - \hat{s}(\theta)(1-\beta)[1 - \tau\gamma - (1-\gamma)C]$$

and interiority implies that

$$\hat{s}(\theta) = \frac{\theta + 1 - \beta}{(1-\beta)(\alpha+1)(1-\gamma\tau-(1-\gamma)C)}$$

and so

$$\hat{q}(\theta) = \theta + 1 - \beta - \frac{\theta + 1 - \beta}{(\alpha+1)}$$

or

$$\hat{q}(\theta) = [\theta + 1 - \beta]\frac{\alpha}{\alpha+1}$$

Step 1.b. Recall that

$$q(\theta) = \theta + (1-\beta)[s(\theta)\tau + 1 - s(\theta)].$$

or

$$q(\theta) = \theta + (1-\beta) - (1-\beta)s(\theta)(1-\tau).$$

Define $\Delta(\theta) = q^*(\theta) - \hat{q}(\theta)$ and note that

$$\Delta(\theta) = \frac{\theta + 1 - \beta}{1+\alpha} - (1-\beta)(1-\tau)s^*(\theta).$$

Step 1.c. We show that $\Delta(\theta)$ is increasing in θ (and so this prove the second part of the proposition for all $\theta \leq \breve{\theta}$). To see this note that

$$\frac{d\Delta(\theta)}{d\theta} = \frac{1}{1+\alpha} - (1-\beta)(1-\tau)\frac{ds^*(\theta)}{d\theta} > 0$$

if and only if

$$\frac{ds^*(\theta)}{d\theta} < \frac{1}{(1+\alpha)(1-\beta)(1-\tau)}$$

To show that the above inequality holds, we return to the unregulated market FOC

$$q^*(\theta) - \alpha(1-\tau)(1-\beta)\frac{s^*(\theta)[1 - s^*(\theta)(1-\tau)]}{[1 - 2s^*(\theta)(1-\tau)]} = 0$$

with some rearranging we get

$$\frac{ds^*(\theta)}{d\theta} = \frac{1}{(1-\beta)(1-\tau)(1+\alpha) + \frac{\alpha(1-\tau)(1-\beta)}{[1-2s^*(\theta)(1-\tau)]^2}\left[2s^*(\theta)(1-\tau)[1 - s^*(\theta)(1-\tau)]\right]}$$

To complete the proof of this step, note that the second term of the denominator of $\frac{ds^*(\theta)}{d\theta}$ is positive because prior to intervention $s^*(\theta) < 1/(2(1-\tau))$. Hence, $\frac{ds^*(\theta)}{d\theta} < \frac{1}{(1+\alpha)(1-\beta)(1-\tau)}$ as required.

Step 1.d. We now conclude and show that $\Delta(\theta) > 0$ for all $\theta \leq \breve{\theta}$. Since $\Delta(\theta)$ is increasing in θ for all $\theta \leq \breve{\theta}$, we just need to show that $\Delta(0) > 0$. To see this note that using the FOC for $s^*(\theta)$ and specializing it for $s^*(0)$ we obtain that

$$s^*(0) = \frac{1}{(2+\alpha)(1-\tau)}$$

and so

$$\Delta(0) = \frac{1-\beta}{1+\alpha} - (1-\beta)(1-\tau)s^*(0) = \frac{1-\beta}{1+\alpha} - \frac{(1-\beta)}{(2+\alpha)} > 0$$

This concludes the proof that $q(\theta) > \hat{q}(\theta)$ for all $\theta \leq \breve{\theta}$.

Step 2. Consider a $\theta \geq \breve{\theta}$.

Step 2.a. We first derive an explicit expression $\hat{q}(\theta)$. Recall that $\hat{s}(\theta) = 1$ and so

$$\hat{q}(\theta) = \theta + (1-\beta)[\tau\gamma + (1-\gamma)C]$$

Step 2.b. Recall that

$$q(\theta) = \theta + (1-\beta)[s(\theta)\tau + 1 - s(\theta)].$$

Define $\Delta(\theta) = q^*(\theta) - \hat{q}(\theta)$ and note that

$$\Delta(\theta) = (1-\beta)[s^*(\theta)\tau + 1 - s^*(\theta) - \tau\gamma - (1-\gamma)C]$$

Step 2.c. Since $s^*(\theta)$ is increasing in θ, it follows that $\Delta(\theta)$ is decreasing in θ, for all $\theta \geq \breve{\theta}$.

Step 2.d. We now conclude and show that $\Delta(\theta) > 0$ for all $\theta \geq \breve{\theta}$. Since $\Delta(\theta)$ is decreasing in θ for all $\theta \leq \breve{\theta}$, we just need to show that $\Delta(1) > 0$. To see this note that $s^*(1) \leq \frac{1}{2(1-\tau)}$ and so

$$\Delta(1) = (1-\beta)[1 - s^*(1)(1-\tau) - \tau\gamma - (1-\gamma)C] \geq (1-\beta)[\frac{1}{2} - \tau\gamma - (1-\gamma)C] \geq 0$$

where the last inequality follows because, since $\tau < 1/2$ and $C \leq 1/2$, then $\tau\gamma + (1-\gamma)C \leq 1/2$. This concludes the proof of Proposition 9.

Proof of Proposition 10. Note that followers' welfare prior and post policy read:

$$W_F = \int n^*(\theta)q^*(\theta)d\theta \text{ and } \hat{W}_F = \int \hat{n}(\theta)\hat{q}(\theta)d\theta$$

We first claim that when $\alpha = 0$ the introduction of transparency decreases followers' welfare and total surplus. To see that, it is sufficient to note that when $\alpha = 0$, for all θ, $n^*(\theta) = \hat{n}(\theta)$ and $q^*(\theta) < \hat{q}(\theta) = 1$. By continuity, the result holds for all $\alpha < \overline{\alpha}$ for some $\overline{\alpha} > 0$.

Next note that the assumption $(1+\alpha)(1-\gamma\tau-(1-\gamma)C) \geq \frac{2-\beta}{1-\beta}$ together with proposition 9 imply that $\Delta(\theta) = q^*(\theta) - \hat{q}(\theta)$ is increasing in θ for all θ. Hence, $\hat{q}(\theta)$ is flatter than $q(\theta)$ and so the distribution of readership $n^*(\cdot)$ FOSD the distribution of readership post policy $\hat{n}(\cdot)$. Hence, since $q^*(\theta)$ is increasing in θ,we obtain that

$$W_F = \int n^*(\theta)q^*(\theta)d\theta > \int \hat{n}(\theta)q^*(\theta)d\theta.$$

We now use proposition 9 (i.e., $\Delta(\theta) = q^*(\theta) - \hat{q}(\theta) > 0$ for all θ), to conclude that

$$W_F = \int n^*(\theta)q^*(\theta)d\theta > \int \hat{n}(\theta)q^*(\theta)d\theta > \int \hat{n}(\theta)\hat{q}(\theta)d\theta = \hat{W}_F.$$

We now turn to total surplus. Recall that

$$TS = \int n^*(\theta)[\theta + 1 - s^*(\theta) + s^*(\theta)\tau]$$

and

$$\hat{TS} = \int \hat{n}(\theta)[\theta + 1 - \hat{s}(\theta) + \hat{s}(\theta)(\tau\gamma + (1-\gamma)C]$$

It is immediate to check that $q^*(\theta) > \hat{q}(\theta)$ for all θ implies that

$$\theta + 1 - s^*(\theta) + s^*(\theta)\tau > \theta + 1 - \hat{s}(\theta) + \hat{s}(\theta)(\tau\gamma + (1-\gamma)C$$

for all θ. And so, replicating the same steps for the readers' welfare, we obtain:

$$\begin{aligned} TS &= \int n^*(\theta)[\theta + 1 - s^*(\theta) + s^*(\theta)\tau] > \int \hat{n}(\theta)[\theta + 1 - s^*(\theta) + s^*(\theta)\tau] \\ &> \int \hat{n}(\theta)[\theta + 1 - \hat{s}(\theta) + \hat{s}(\theta)(\tau\gamma + (1-\gamma)C] = \hat{TS}. \end{aligned}$$

On-line appendix: not for publication

Generalization of followers' utility. Consider the model introduced in Section 2, with the difference that the utility to a consumer who is matched to influencer i is now

$$q_i(s_i) = v(\theta_i) + [1 - s_i(1-\tau)](1-\beta)\mu(\theta_i)$$

In the model of Section 2 we had $v(\theta) = \theta$ and $\mu(\theta) = 1$, for all θ. We assume that $v(\theta)$ and $\mu(\theta)$ are both increasing functions of θ. We further assume that

Assumption 1. *The function $v(\theta)$ grows faster than $\mu(\theta)$, i.e., $\frac{v(\theta)}{\mu(\theta)}$ is increasing in θ.*

Note that

$$\Pi_i(\mathbf{s}) = s_i n_i(\mathbf{s})[1 - s_i(1-\tau)]\beta\mu(\theta_i)$$

Taking the FOC and imposing interiority, we obtain:

$$1 - 2s_i(1-\tau) = s_i\alpha(1-\tau)\frac{q_i - v(\theta_i)}{q_i}. \tag{11}$$

We obtain that $s_i = s(\theta)$ for all $\theta_i = \theta$ and that the solution is unique and $s(\theta) < \frac{1}{2(1-\tau)}$: Indeed the LHS is decreasing in s_i and is zero at $s_i = \frac{1}{2(1-\tau)}$, and the RHS is 0 at $s_i = 0$ and it is increasing in s_i for $s_i \leq \frac{1}{2(1-\tau)}$. This is the analogous result to Proposition 1. In the rest of this Appendix we derive analogous results to Proposition 2.

Fact 1. We show that $s(\theta)$ is increasing in θ. For this it is sufficient to show that the RHS of 11 is decreasing in θ (note that the LHS is independent of θ). To show this, we need to show that

$$\frac{q(\theta) - v(\theta)}{q(\theta)}$$

is decreasing in θ. Taking the derivative of this ratio with respect to θ we get

$$\frac{1}{q^2(\theta)}\left[[1-s(\theta)(1-\tau)](1-\beta)\mu'(\theta)q(\theta) - [1-s(\theta)(1-\tau)](1-\beta)\mu(\theta_i)[v'(\theta) + [1-s(\theta)(1-\tau)](1-\beta)\mu'(\theta)]\right]$$

or

$$\frac{[1-s(\theta)(1-\tau)](1-\beta)}{q^2(\theta)}\left[[\mu'(\theta)q(\theta) - \mu(\theta)(1-s(\theta)(1-\tau))(1-\beta)\mu'(\theta)] - \mu(\theta)v'(\theta)\right]$$

or

$$\frac{[1-s(\theta)(1-\tau)](1-\beta)}{q^2(\theta)}\left[v(\theta)\mu'(\theta)] - \mu(\theta_i)v'(\theta)\right] < 0$$

where the last inequality follows from our assumption that the function $v(\theta)$ grows faster than $\mu(\theta)$.

Fact 2. We now show that $q(\theta)$ and $n(\theta)$ increases with θ. Note that in equilibrium

$$q(\theta) = \alpha(1-\tau)\frac{s(\theta)(q(\theta) - v(\theta)}{1 - 2s(\theta)(1-\tau)}$$

or

$$q(\theta) = \alpha(1-\tau)\mu(\theta)(1-\beta)\frac{s(\theta)[1 - s(\theta)(1-\tau)]}{1 - 2s(\theta)(1-\tau)}$$

Note that the first term of the LHS is increasing in θ and that the second term of the LHS is increasing in $s(\theta)$ and since $s(\theta)$ is increasing in θ, we obtain the result that $q(\theta)$ increases with θ. Since $q(\theta)$ is increasing in θ is follows that $n(\theta)$ is increasing in θ.

Fact 3. We finally show that the price-per-post-per-view decreases in θ. Recall that

$$\frac{p(\theta)}{n(\theta)} = (1 - s(\theta)(1-\tau))\,\beta\mu(\theta)$$

and therefore to show that $\frac{p(\theta)}{n(\theta)}$ decreases in θ it is sufficient to show that

$$\xi \triangleq (1 - s(\theta)(1-\tau))\,\mu(\theta)(1-\beta)$$

decreases in θ. Next, we can rewrite the FOC 11 as

$$\frac{1-2s(\theta)(1-\tau)}{s(\theta)} = \alpha(1-\tau)\frac{\xi}{v+\xi}.$$

Solving for ξ we get

$$\xi = -v\frac{1-2s(\theta)+2s(\theta)\tau}{-2s(\theta)-sv\alpha+2s(\theta)\tau+s(\theta)\alpha\tau+1}.$$

We then note that $v(\theta)$ increases in θ, and that $\frac{1-2s(\theta)+2s(\theta)\tau}{-2s(\theta)-s(\theta)\alpha+2s(\theta)\tau+s(\theta)\alpha\tau+1}$ increases in $s(\theta)$ and therefore in θ. As a result, ξ decreases in θ as required.

Other results on social planner analysis. Note that

$$W_F = \int n(\theta)q(\theta)d\theta$$

$$W_I = \frac{\beta}{1-\beta}\int n(\theta)[q(\theta)-\theta]s(\theta)d\theta$$

$$W_M = \int n(\theta)[1-s(\theta)(1-\tau)]\beta(1-s(\theta))d\theta = \frac{\beta}{1-\beta}\int n(\theta)[q(\theta)-\theta](1-s(\theta))d\theta$$

Proposition 11. *Consider a social planner:*

1. *If the planner wishes to maximize followers' welfare then: there exists $\widehat{\theta}$ such that for every $\theta < \widehat{\theta}$ the planner sets $s(\theta) = 1$ and for every $\theta > \widehat{\theta}$ the planner sets $s(\theta) = 0$.*
2. *If the planner wishes to maximize influencers' aggregate profit then the planner sets $s(\theta) = \frac{1}{2(1-\tau)}$ for all θ.*
3. *If the planner wishes to maximize marketers' profits then the planner sets $s(\theta) = 0$ for all θ.*

Proof of Proposition 11. We first consider followers' aggregate welfare. Consider any function $s(\theta)$, and a type θ'. We would like to know whether W_F increases or decreases in $s(\theta')$. It might be useful for this to write followers' welfare as follows

$$W_F = \frac{\int_0^1 q(\theta)^{\alpha+1}d\theta}{\int_0^1 q(\theta)^{\alpha}d\theta}$$

and recall that $q(\theta)$ decreases in $s(\theta)$. Next, note that W_F increases in $q(\theta')$ if

$$\frac{\frac{\partial q(\theta')^{\alpha+1}}{\partial q(\theta')}}{\int_0^1 q(\theta)^{\alpha+1}d\theta} > \frac{\frac{\partial q(\theta')^{\alpha}}{\partial q(\theta')}}{\int_0^1 q(\theta)^{\alpha}d\theta}$$

or

$$\frac{(\alpha+1)\,q(\theta')}{\int_0^1 q(\theta)^{\alpha+1}d\theta} > \frac{\alpha}{\int_0^1 q(\theta)^{\alpha}d\theta}$$

and decreases otherwise. Note further that if the inequality holds for θ' for some $q(\theta')$ then it holds for any larger $q(\theta')$, and if the reverse inequality holds for θ' for some $q(\theta')$ then it holds for any smaller $q(\theta')$. Thus, for any θ' the planner will choose $s(\theta') \in \{0,1\}$. Equivalently, for each θ' the planner will choose $q(\theta') \in \{\theta' + \tau(1-\beta), \theta' + (1-\beta)\}$.

To be more specific, the planner will choose $s(\theta') = 0$ (equivalently, $q(\theta') = \theta' + (1-\beta)$) if

$$\frac{q(\theta')^{\alpha+1}|_{s(\theta')=0} - q(\theta')^{\alpha+1}|_{s(\theta')=1}}{\int_0^1 q(\theta)^{\alpha+1} d\theta} > \frac{q(\theta')^{\alpha}|_{s(\theta')=0} - q(\theta')^{\alpha}|_{s(\theta')=1}}{\int_0^1 q(\theta)^{\alpha} d\theta}$$

and choose $s(\theta') = 1$ (equivalently, $q(\theta') = \theta' + \tau(1-\beta)$) otherwise. The condition simplifies to

$$\frac{(\theta' + (1-\beta))^{\alpha+1} - (\theta' + \tau(1-\beta))^{\alpha+1}}{\int_0^1 q(\theta)^{\alpha+1} d\theta} > \frac{(\theta' + (1-\beta))^{\alpha} - (\theta' + \tau(1-\beta))^{\alpha}}{\int_0^1 q(\theta)^{\alpha} d\theta}$$

or

$$\frac{(\theta' + (1-\beta))^{\alpha+1} - (\theta' + \tau(1-\beta))^{\alpha+1}}{(\theta' + (1-\beta))^{\alpha} - (\theta' + \tau(1-\beta))^{\alpha}} > \frac{\int_0^1 q(\theta)^{\alpha+1} d\theta}{\int_0^1 q(\theta)^{\alpha} d\theta}$$

where the RHS is independent of θ' and the LHS increases in θ'. This completes the proof of the first part of the proposition.

Consider next the case where the planner cares about influencers' aggregate profits. Note that,

$$W_I = \beta \frac{\int_0^1 q(\theta)^{\alpha} (1 - s(\theta)(1-\tau)) s(\theta) d\theta}{\int_0^1 q(\theta)^{\alpha} d\theta}$$

We now evaluate an increase in $s(\theta')$ in two steps. We first hold $q(\theta')$ fixed and evaluate the effect of the increase in $s(\theta')$ (namely, the sign of $\partial \left(\beta \frac{\int_0^1 q(\theta)^{\alpha}(1-s(\theta)(1-\tau))s(\theta)d\theta}{\int_0^1 q(\theta)^{\alpha} d\theta} \right) / \partial s(\theta')$) and then evaluate the effect through the decrease in $q(\theta')$.

For the first part, note that $(1 - s(\theta)(1-\tau)) s(\theta)$ is increasing and then decreasing in $s(\theta)$ for any θ (with the argmax being at $s = \frac{1}{2(1-\tau)}$ independent of θ, which is a value greater or equal to equilibrium $s(\theta)$).

For the second part it is sufficient to note that, setting $s(\theta)$ to equal some s' for all θ, $W_I = \beta(1 - s'(1-\tau))s'$ which is independent of θ. This complete the proof of the second part of the proposition.

We conclude with the case in which the planner wishes to maximize the aggregate profits of marketers. In our model, only marketers with products that appear in organic posts make profit, and each marketer's per-post-per-follower profit decreases in the fraction of sponsored content of the influencer who recommends its product. Therefore, to maximize marketers' profits the planned will set $s(\theta) = 0$ for all θ. This concludes the proof of the Proposition.