WORKING PAPER NO: 580

Influencer Marketing with Fake Followers

Abhinav Anand
Assistant Professor
Finance and Accounting
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 5600 76
abhinav.anand@iimb.ac.in

Souvik Dutta
Assistant Professor
Economics and Social Sciences
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 5600 76
souvik@iimb.ac.in

Prithwiraj Mukherjee
Assistant Professor
Marketing
Indian Institute of Management Bangalore
Bannerghatta Road, Bangalore – 5600 76
pmukherjee@iimb.ac.in

Year of Publication – March 2019

Influencer Marketing with Fake Followers

Abstract

Influencer marketing is a practice where an advertiser pays a popular social media user (influencer) in exchange for brand endorsement. We characterize the advertiser's optimal contract when the influencer can inflate her publicly displayed follower count by buying fake followers. We derive optimal contracts for two scenarios: (a) \pre sign-up" where a potential influencer is not yet on a given social media platform, but has a promise of a following and (b) \post sign-up" where the influencer is on social media and privately knows her true follower count.

The optimal contract stipulates a fixed payment equal to the influencer's outside option and a variable payment increasing in her follower count. In the pre sign-up scenario, the advertiser extracts all the surplus and the equilibrium features truthful display of the influencer's follower count. However in the post sign-up scenario, the advertiser must pay over and above the influencer's outside option; and needs to tolerate high levels of faking. Our results suggest that advertisers are better o_ hiring potential influencers with authentic, social media-independent mass appeal rather than the more common practice of hiring them based on merely their follower count.

Keywords: Digital marketing, social media, influencer marketing, fake followers, optimal control, contract theory.

Influencer marketing with fake followers*

Abhinav Anand†
Souvik Dutta‡
Prithwiraj Mukherjee§

March 28, 2019

Abstract

Influencer marketing is a practice where an advertiser pays a popular social media user (influencer) in exchange for brand endorsement. We characterize the advertiser's optimal contract when the influencer can inflate her publicly displayed follower count by buying fake followers. We derive optimal contracts for two scenarios: (a) "pre sign-up" where a potential influencer is not yet on a given social media platform, but has a promise of a following and (b) "post sign-up" where the influencer is on social media and privately knows her true follower count.

The optimal contract stipulates a fixed payment equal to the influencer's outside option and a variable payment increasing in her follower count. In the pre sign-up scenario, the advertiser extracts all the surplus and the equilibrium features truthful display of the influencer's follower count. However in the post sign-up scenario, the advertiser must pay over and above the influencer's outside option; and needs to tolerate high levels of faking. Our results suggest that advertisers are better off hiring potential influencers with authentic, social media-independent mass appeal rather than the more common practice of hiring them based on merely their follower count.

Keywords: Digital marketing, social media, influencer marketing, fake followers, optimal control, contract theory.

JEL Classification: D82, D86, M31, M37

---

*Author names are in alphabetical order with equal contribution. We thank Bruno Badia, Manaswini Bhalla, Pradeep Chintagunta, Tirthatanmoy Das, Sreelata Jonnalagedda, Ashish Kumar, Sanjog Misra and participants at the Chicago Booth-India Quantitative Marketing Conference (2018) at IIM Bangalore for their feedback and comments.

†Finance and Accounting, Indian Institute of Management Bangalore.

‡Economics and Social Sciences, Indian Institute of Management Bangalore.

§Marketing, Indian Institute of Management Bangalore.

"At Unilever, we believe influencers are an important way to reach consumers and grow our brands. Their power comes from a deep, authentic and direct connection with people, but certain practices like buying followers can easily undermine these relationships." — Keith Weed, Chief Marketing and Communications Officer, Unilever (Stewart, 2018)

1 Introduction

Advertisers often pay popular social media users known as *influencers* to endorse their products online. Many of these influencers have large numbers of self-selected followers who share their interests (travel, cooking etc.), looking up to them for advice in these domains. According to The Economist (2016), YouTube influencers with over 7 million followers command upto $300,000 per sponsored post, while the corresponding figures for Instagram, Facebook and Twitter are $150,000, $187,500 and $60,000 respectively, allowing social media followings to be monetized lucratively. Even influencers with less than 250,000 followers can make hundreds of dollars per sponsored post. Figure 1 shows some typical compensations for influencers on various platforms versus their follower counts. A Linqia (2018) survey across sectors including consumer packaged goods, food and beverage and retail in the US finds that 86% of marketers surveyed used some form of influencer marketing in 2017, and of them, 92% reported finding it effective. 39% of those surveyed planned to increase their influencer marketing budgets. Similar trends reported by eMarketer (2017) and IRI (2018) suggest that influencer marketing is growing.

Influencer marketing has led to the emergence of shady businesses called *click farms* which for a price, offer influencers fake followers, inflate the number of "likes" on their fan pages, and post spurious comments on their posts. Influencers use these services to fraudulently command higher fees from advertisers for promotional posts. A New York Times exposé (Confessore et al., 2018) finds that several personalities including social media influencers have bought fake followers from a click farm called Devumi.

Sway Ops, an influencer marketing agency estimates the total magnitude of influencer fraud to be about $1 billion (Pathak, 2017). They find that in a single day, of 118,007 comments sampled on #sponsored or #ad tagged Instagram posts, less than 18% were made by genuine users. Another study by the Points North Group finds that influencers hired by Ritz-Carlton have 78% fake followers (Neff, 2018). The corresponding numbers for Procter and Gamble's Pampers and Olay brands are 32% and 19% respectively. The quote at the beginning of this paper, by Unilever's Chief Marketing and Communications Officer at the Cannes film festival, indicates that marketers are acutely aware of the fake follower problem.

Paquet-Clouston et al. (2017) report that click farm clients pay an average of $49 for every 1,000 YouTube followers. The corresponding figures are $34 for Facebook, $16 for Instagram and $15 for Twitter. Average prices for 1,000 likes on these platforms are $50, $20, $14 and $15 respectively. Google searches corroborate these claims—hundreds of click farms and millions of fake followers are accessible to anyone with a credit card. A Buzzfeed investigation suggests increasing sophistication of fake follower bots; shell companies that calibrate bots based on the behavioral patterns of genuine users (and thus virtually impossible to detect) on their own apps may have siphoned off millions of dollars from advertisers (Silverman, 2018).

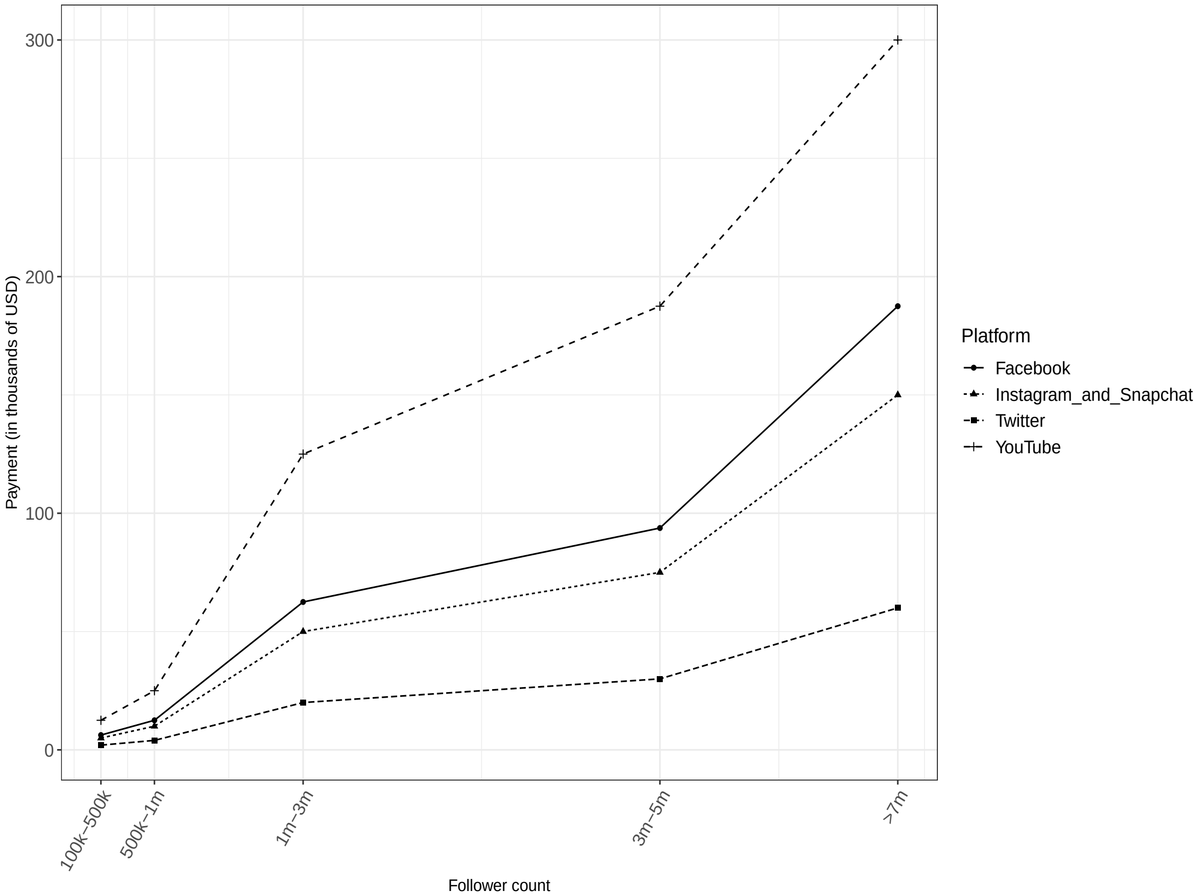

Figure 1: Typical compensation schemes for influencers versus follower counts on various platforms. Adapted from The Economist (2016)

In our work, we adapt the model of fraud in Crocker and Morgan (1998) to design an optimal contract between a risk-neutral advertiser and a risk-neutral influencer. The advertiser proposes to pay the influencer in exchange for brand endorsement in order to reach her followers on a social media platform. In order to appear more popular or influential (see section 2.1 for a link between popularity and influence) and thereby command more payment, the influencer can buy costly fake followers. The advertiser, as a result, can only observe the publicly displayed follower count consisting of both real and fake followers. We assume that costs to the influencer of inflating her follower count are convex, i.e., rising progressively as the number of fake followers increases. Additionally, we assume that the advertiser cannot reliably verify the true number of followers due to the increasing sophistication of click farms and fake bots' online behavior (as noted in Silverman, 2018).

Optimal contract design in such a setting features an intrinsic tradeoff. A contract that is sensitive to the observed number of followers ensures efficiency but generates perverse incentives for an influencer to inflate her true follower count. On the other hand, a payment scheme that is unresponsive to the observed follower count mitigates the problem of falsification but is inefficient, as it may fail to provide incentives for those with high follower counts to accept the contract. Hence the optimal contract must find a balance between these two opposing forces.

We show that with information asymmetry when the influencer has private knowledge of her true follower count, under the optimal contract, it is worthwhile for the influencer to buy fake followers, and the advertiser needs to pay her a variable rate increasing in her follower count over and above her outside option value. We refer to this case as the "post sign-up" scenario. Thus the post sign-up influencer is able to extract informational rents from the advertiser. The latter observes high levels of faking from (almost) all types of influencers but tolerates such behavior since it provides an important tool to ordinally rank, and hence distinguish between different types of privately informed influencers. In other words, the ability of faking lets the influencer signal her true type credibly. As misrepresentation costs are convex and increasing in the degree of misrepresentation, each influencer type faces different costs of displaying a given follower count. The advertiser is able to exploit this variation in faking costs to distinguish influencers based on their true follower count.

We also design an optimal contract for a potential influencer who does not yet have a social media account but may be incentivized to open one given the expectation of a certain follower count. We refer to this case as the "pre sign-up" scenario. In this case there are no informational rents that the influencer can extract and the optimal contract stipulates that the advertiser pay a fixed sum equal to the influencer's outside option value. Under such a contract, there is no incentive for the influencer to resort to any falsification.

From the perspective of the advertiser the pre sign-up case is better than the post sign-up case. This is so because the advertiser is able to extract the full surplus in the former, and can limit the payment to the influencer at her outside option value. The influencer is able to extract informational rents in the post sign-up scenario because of the information asymmetry and the advertiser has to tolerate it as this provides a screening mechanism for the influencer to report her true type credibly. Therefore, for advertisers to reach their audience most cost-effectively, it is better to induce potential social media influencers to sign-up, then restrict their payment to the monetary value associated with their outside option; and finally extract the benefits of their brand endorsements.

The rest of the paper is structured as follows: in section 2, we discuss previous work that is related to our modeling assumptions and methodology. We then outline our model in section 3, which leads to the framing and solution of an optimal control problem which we describe and solve in section 4. We outline three propositions, whose analytical proofs are presented in appendices at the end of the paper. Finally, we discuss the properties, managerial implications and some scope for extension of our model in section 5.

2 Background

We briefly outline three streams of literature germane to our problem: influencer marketing, models of economic fraud and applications of contract theory to marketing. Because some of these domains are large, we outline only a few studies in each, to motivate and contextualize our own work. Section 2.1 outlines the role of follower count in influencer marketing. A large follower count could lead to more effective endorsements for the advertiser which could incentivize influencers to inflate their follower count fraudulently. In section 2.2 we discuss related models of economic fraud in online and offline businesses, and finally in section 2.3, we motivate our choice of methodology of contract theory.

2.1 Influencer marketing: the role of follower count

Influencer marketing is a promotional method that is somewhere between advertising and word of mouth.[1] It is advertising in the sense that firms pay influencers to endorse their products, but choose this channel due to influencers' (endorsers') connect with their followers, as well as their reach. Just as endorsers in advertising, influencers can be experts (tech bloggers, cooks etc.), celebrities (musicians, actors etc.) or even lay endorsers[2] (Tellis, 2003, chapter 11). Possible mechanisms of influence could be via source (influencer) credibility (Hovland and Weiss, 1951) and source attractiveness theories (McGuire, 1985; McCracken, 1989).[3] According to the source credibility theory, an endorsement message is more trustworthy when it comes from a source perceived to be an "expert" in a given domain. In a qualitative study of Instagram users, Djafarova and Rushworth (2017) find that non-traditional celebrities may be perceived as more credible than traditional celebrities, and have a powerful impact on consumers' fashion choices. An experimental study by Jin and Phua (2014) finds that the perceived credibility of a Twitter user is positively related to her number of followers. Alongside source credibility is source attractiveness (familiarity, likeability and similarity) (Tellis, 2003, chapter 11). Experiments by De Veirman et al. (2017) demonstrate that Instagram influencers' perceived likeability is positively related to their follower counts, possibly driven by perceived popularity.

From a more straightforward perspective, a social media user with a larger follower count represents an advertising medium with a larger reach. While the reach of traditional media like TV, radio, hoardings and newspapers can only be approximately estimated, it is reasonable to expect a naive advertiser unaware of fake followers, to place a high amount of trust in seemingly accurate measures like follower counts, number of likes on a sponsored posts, retweets and impressions offered by web analytics tools. Agent-based studies of targeted new product seeding like Libai et al. (2013) demonstrate that targeting hubs with high numbers of acquaintances speeds up the adoption process, while Yoganarasimhan (2012), in an empirical analysis, suggests that the popularity of a social media message over and above an individual's immediate neighborhood is driven by her follower count.

While follower count is not the only means of determining social influence (see Kannan and Li (2017) for a comprehensive review), it certainly is a popular metric used by digital marketers today to identify influentials on social networks. The above discussions shed some light on why advertisers pay more to influencers with higher follower counts, in turn generating incentives for influencers to boost their own followings via unethical means like buying fake followers. Given that the practice of buying fake followers is highly prevalent today, our work sheds more light into the economics of this fraud. We show that under the optimal contract between an advertiser and an exisiting social media influencer, it is not possible to guarantee efficiency and elicit truthfulness simultaneously. In fact, we show that those with high true follower counts buy more fake followers to credibly signal their type to

[1] This has led to ethical conundrums surrounding undisclosed influencer promotions. The US Federal Trade Commission has taken cognizance of the possibility of consumers being misled by influencer marketing, and has mandated that sponsored posts must now clearly mention the relationship between the influencer and brand, usually as a hashtag such as #sponsored or #ad. Platforms like Instagram have implemented algorithms to automatically detect and tag paid influencer posts.

[2] For example, Loki the Wolf Dog, an Instagram account owned by Denver-based blogger Kelly Lund, has over 2 million followers and has done influencer campaigns for companies like Toyota and Mercedes-Benz.

[3] There is also the powerful meaning transfer theory (McCracken, 1989), but we do not find any studies relating it to a social media user's follower count.

the advertiser. We demonstrate that this problem may be mitigated if the advertsier scouts and signs up “diamonds in the dust,” i.e., potential influencers who are not yet on a social media platform, but show promise of a large following.

2.2 Economic fraud

Online businesses are prone to several kinds of marketing fraud. Advertisers paying per click frequently encounter click fraud where click farms simulate genuine clicks. Wilbur and Zhu (2009) in an important, related study investigate the problem of click fraud in search advertisement in a game theoretic setting. Their results suggest that usage of a neutral third party to audit click fraud detection can benefit the search advertising industry. Another form of online fraud consists of fake reviews, when businesses post either fake positive reviews for themselves or fake negative reviews for their competitors. Lappas et al. (2016) demonstrate how even a few fake reviews can significantly boost hotels’ visibility. Luca and Zervas (2016) find that the prevalence of suspicious restaurant reviews on Yelp has grown over time. They find that restaurants with weaker reputations tend to engage more in online review fraud when faced with increasing competition.

A model of insurance and sharecropping fraud where agents involve in costly falsification is developed in a contract theoretic setting in Crocker and Morgan (1998). Their model yields results that have been extended to other fraud scenarios, like misreporting of earnings by CEOs (Crocker and Slemrod, 2007; Sun, 2014), many types of insurance fraud (Crocker and Tennyson, 2002; Dionne et al., 2009; Doherty and Smetters, 2005), and in designing optimal product return policies (Crocker and Letizia, 2014).

Employee theft in retail can also be modelled analytically. Mishra and Prasad (2006) demonstrate that a complete elimination of theft may be economically infeasible and derive an optimal frequency of random inspections to minimize losses due to theft by retail employees. With this paper, we contribute to the literature on economic fraud by modeling the emerging phenomenon of fraud in influencer marketing and demonstrating that eliminating fraud may be impossible even under optimal contracts.

2.3 Contract theory

In contract theory a principal wishes to hire an agent and designs an optimal contract which maximizes its profit while respecting the agent’s participation and incentive compatibility constraints.[4] Contract theoretic approaches have been used in marketing scenarios such as designing warranties and extended service contracts (Padmanabhan and Rao, 1993) and delegation of pricing decisions to salespersons (Bhardwaj, 2001; Mishra and Prasad, 2004, 2005) to name a few. Other noteworthy applications of contract theory in marketing include explaining product development incentives (Simester and Zhang, 2010) and to explain how internal lobbying by salespersons for lower prices can elicit truthful information about market demand (Simester and Zhang, 2014). Our paper incorporates contract theory and optimal control theory in the digital marketing literature, illustrating the economics of influencer marketing fraud.

[4] See Bolton and Dewatripont (2005) for a comprehensive exposition.

Term	*Description*
n	True number of followers of influencer
$f(n), F(n)$	Probability density and cumulative distribution function of n
$[n_L, n_H]$	Support of the probability density function f
$u(n)$	Influencer's displayed number of followers (true + fake)
$v_1(n)$	Variable payment to the influencer
v_2	Fixed payment to the influencer
$A(n)$	Advertiser's revenue function from reaching n followers
$c(u(n) - n)$	Cost of displaying $u(n)$ followers when true followers are n
$\Pi(\cdot)$	Advertiser's payoff
$Y(\cdot)$	Influencer's payoff
$\bar{Y}$	Influencer's outside option value
$\mathbb{H}$	Hamiltonian
$\lambda(n)$	Co-state variable
μ	Lagrange multiplier

Table 1: Summary of notation used in our model

3 The model

We adapt the model of sharecropping fraud in Crocker and Morgan (1998) which examines the design of compensation schemes while taking into account potential falsification of claims when verification by the principal is not possible. We consider a risk-neutral advertiser who wishes to reach the followers of a risk neutral influencer. The advertiser proposes to pay the influencer for brand endorsement via social media posts. Table 1 provides a short description of all notation used in our model.

The influencer privately knows her own number of followers n, and can inflate her follower count by buying costly fake followers. The advertiser observes only the publicly displayed follower count and not the true number of followers. However, the advertiser is aware that the true number of followers of the influencer are distributed in $[n_L, n_H]$ according to the probability density $f(n)$.[5]

The advertiser's profit is denoted by $\Pi(v_1; v_2; u)$ and the influencer's payoff is denoted by $Y(v_1; v_2; u; n)$. Under the optimal contract between the two, the equilibrium outcome is characterized by a 3-tuple: $\{v_1, v_2, u\}$ where $v_1(n)$ is a variable payment depending on the influencer's follower count; v_2 is a fixed payment; and $u(n)$ is the function used by the influencer to inflate her true follower count n.

While in practice we expect to observe an indirect mechanism where the advertiser's payment to the influencer is conditioned on the post-falsification follower count u, in order to characterize the solution we focus on the corresponding direct mechanism. In other words, although the variable payment depends on influencer's publicly displayed, possibly inflated number of followers, i.e., $v_1(u)$, the revelation principle (Myerson, 1979) guarantees that the same equilibrium outcome can be achieved under an incentive-compatible direct mechanism where the influencer receives variable payment $v_1(n)$.

We note that for the advertiser, the decision variable is not the underlying follower count n of the influencer. This is because we assume that neither the advertiser nor the influencer can exercise control on the number of followers that the influencer possesses; and that the advertiser can only vary

[5]While n is a natural number we assume hereon for the sake of mathematical convenience that n is a non-negative real number.

the compensation scheme $(v_1(n), v_2)$ in response to the display function $u(n)$ of the influencer.[6]

In order for the equilibrium to be incentive compatible, it must be that at the optimal v_1^*, v_2^*, u^*, there is no incentive for the influencer to not act according to her own type. This happens only if:

$$Y(v_1^*(n); v_2^*; u^*(n); n) \geq Y(v_1^*(\bar{n}); v_2^*; u^*(\bar{n}); n) \qquad \forall \bar{n} \neq n \in [n_L, n_H]$$

For brevity we denote the optimal value function $Y(v_1^*; v_2^*; u^*; n) \equiv Y^*(n)$ and note that since $Y^*(\cdot)$ is optimal, its derivative with respect to the arguments v_1, v_2, u must be 0:

$$\left.\frac{\partial Y^*}{\partial v_1}\right|_{v_1=v_1^*} = \left.\frac{\partial Y^*}{\partial v_2}\right|_{v_2=v_2^*} = \left.\frac{\partial Y^*}{\partial u}\right|_{u=u^*} = 0$$

Using the envelope theorem, by means of the total derivative, we establish the dependence of the optimal value function Y^* on the parameter n by:

$$\frac{dY^*}{dn} = \frac{\partial Y^*}{\partial v_1^*}\frac{dv_1^*}{dn} + \frac{\partial Y^*}{\partial v_2^*}\frac{dv_2^*}{dn} + \frac{\partial Y^*}{\partial u^*}\frac{du^*}{dn} + \frac{\partial Y^*}{\partial n} \cdot 1$$

This leads to the standard envelope condition:

$$\frac{dY^*}{dn} = \frac{\partial Y^*}{\partial n}$$

3.1 The optimization program

The advertiser wishes to maximize its expected profit:[7]

$$\max_{v_1, v_2, u} \left(\int_{n_L}^{n_H} \Pi(v_1; v_2; u) f(n) dn \right)$$

subject to the incentive compatibility constraint:

$$\frac{dY}{dn} = \frac{\partial Y}{\partial n}$$

and the participation constraint which in general could be of the following two types: pre sign-up or post sign-up.

3.1.1 Pre sign-up participation constraint

Suppose a potential influencer has not yet signed up on social media. In order to ensure that it is worthwhile for her to participate by signing up and endorsing the advertiser's product, it must be that the ex-ante expected payoff from participation is more than her outside option $\bar{Y}$:

$$\int_{n_L}^{n_H} Y(v_1; v_2; u; n) f(n) dn \geq \bar{Y}$$

[6] Consequently, risk neutrality of payoff functions in this setting implies that the profit function of the advertiser is linear not in the follower count n but in the functions v_1, v_2, u.

[7] We re-emphasize that the maximization program is expressed with respect to the functions v_1, v_2, u and not with respect to n.

3.1.2 Post sign-up participation constraint

In this case the influencer is already on social media and privately knows her true number of followers. In order for her to find participation worthwhile, it must be that her realized payoff from n followers is higher than her outside option $\bar{Y}$.[8] Hence the post sign-up participation constraint is:

$$Y(v_1; v_2; u; n) \geq \bar{Y}$$

4 The optimal control problem

Optimization programs featuring integrals in objective functions and derivatives in the constraints can be solved by setting up an optimal control problem and finding the stationary points of the associated Hamiltonian.

4.1 Pre sign-up optimal control

The pre sign-up optimal control problem is:

$$\max_{v_1, v_2, u} \left(\int_{n_L}^{n_H} \Pi(v_1; v_2; u) f(n) dn \right) : \tag{1}$$

$$\frac{dY}{dn} = \frac{\partial Y}{\partial n} \tag{2}$$

$$\int_{n_L}^{n_H} Y(v_1; v_2; u; n) f(n) dn \geq \bar{Y} \tag{3}$$

The expected profit function (1) under the incentive compatibility constraint (2) and pre sign-up constraint (3) can be combined into the following Hamiltonian:

$$\mathbb{H} = \Pi(v_1; v_2; u) f(n) + \lambda(n) Y_n + \mu Y(v_1; v_2; u; n) f(n) \tag{4}$$

In the above Hamiltonian formulation, $Y(\cdot)$, the influencer's payoff function is the state variable with its equation of motion represented by condition (2). The control variable is $u(\cdot)$; $\lambda(n)$ is the co-state variable corresponding to the incentive compatibility contraint (2); and μ is the Lagrangian multiplier associated with the pre sign-up participation contraint (3). The necessary first order conditions are obtained from the Pontryagin maximum principle as below:

Pontryagin conditions

1. Optimality condition:

$$\max_{u} \mathbb{H} \qquad \forall n \in [n_L, n_H]$$

[8]While in general, the post sign-up outside option could be different from the pre sign-up outside option, we can, without loss of generality take $\bar{Y}$ to be the maximum of the two and interpret it to be a measure of the influencer's opportunity cost.

2. Equation of motion for state:

$$\frac{dY}{dn} = \frac{\partial \mathbb{H}}{\partial \lambda} = Y_n$$

3. Equation of motion for costate:

$$\frac{d\lambda}{dn} = -\frac{\partial \mathbb{H}}{\partial Y}$$

4. Transversality condition for state:

$$\lambda(n_H) = 0$$

Together, the four conditions stated above must be necessarily true at the optimal. Proposition 1 characterizes the solution further.

Proposition 1. *The necessary conditions which characterize the solution of the optimal control problem with the pre sign-up participation constraint are as follows:*

$$f \cdot \left(\Pi_u - \Pi_{v_1} \frac{Y_u}{Y_{v_1}} \right) + \lambda \cdot \left(Y_{u,n} - Y_{v_1,n} \frac{Y_u}{Y_{v_1}} \right) = 0 \tag{5}$$

$$\dot{\lambda} = \frac{d\lambda}{dn} = -f \cdot \frac{\Pi_{v_1}}{Y_{v_1}} - \lambda \cdot \frac{Y_{v_1,n}}{Y_{v_1}} - \mu f \tag{6}$$

$$\int_{n_L}^{n_H} \left(\Pi_{v_2} + \mu \cdot Y_{v_2} \right) f(n) dn = 0 \tag{7}$$

Proof. See appendix A. □

4.1.1 Pre sign-up optimal contract

So far we have not assumed much about the payoff functions of the advertiser or the influencer apart from their risk neutrality. We now discuss the payoff functions of the advertiser and the influencer.

Advertiser's payoff

The advertiser must pay the influencer a sum of (dollars, say) $v_1(n) + v_2$. On the other hand by exposing n followers of the influencer to its endorsement, it earns a sum of $A(n)$, where $A(\cdot)$ is an exogenous revenue function which according to the advertiser captures the benefits (in dollar terms) of reaching out to n social media followers. For example, $A(n)$ could be based on the advertiser's past experience and managerial judgment.[9]

The advertiser's profit function is the difference between the revenue from reaching the influencer's n followers and the payment made to the influencer.

$$\Pi(v_1; v_2; u) = A(n) - v_1(n) - v_2$$

We note that $A(n)$ is exogenous to v_1, v_2, u, thus converting the profit maximization problem to cost

[9] We note that it may not be feasible to have an exact formulation of $A(n)$ due to possibly unreliable metrics (Lewis and Rao, 2015; Sridhar et al., 2017).

minimization for the advertiser:

$$\max_{\{v_1,v_2,u\}} \Pi(v_1;v_2;u) = \max_{\{v_1,v_2,u\}} (A(n) - v_1(n) - v_2) = A(n) + \max_{\{v_1,v_2,u\}} (-v_1(n) - v_2)$$

We leave the details of the revenue function $A(n)$ for the section 5.2. However, we note that the admissible class of revenue functions is dictated by the consideration that for the advertiser to find it worthwhile to propose a contract for hiring an influencer, the revenue must be sufficiently high.

Influencer's payoff

For the influencer, payoff is gained due to payments from the advertiser but there is a cost of inflating follower count. We assume that the cost function varies with the degree of misrepresentation: $c(u-n)$. We assume that $c \geq 0$; $c(0) = 0$; $c'(0) = 0$ and $c'' > 0$. This means that costs are at least zero, no inflation entails no cost and the costs of fraud rise progressively higher as the extent of cover-up increases. Thus the payoff of the influencer is given by:

$$Y(v_1;v_2;u;n) = v_1(n) + v_2 - c(u(n) - n)$$

Proposition 2 characterizes the pre sign-up optimal contract.

Proposition 2. *The pre sign-up optimal contract is characterized by the following 3-tuple:*

$$u(n) = n \tag{8}$$

$$v_1(n) = 0 \tag{9}$$

$$v_2 = \bar{Y} \tag{10}$$

Proof. See appendix B. □

The optimal contract stipulates that the influencer be paid a fixed amount equal to her outside option value. The variable payment to be given to the influencer is 0. Thus faced with such a reward schedule, the influencer has no incentive to inflate her follower count and hence she reports her true number of followers $u(n) = n$ and receives fixed payment $\bar{Y}$.

4.2 Post sign-up optimal control

The following Hamiltonian $\mathbb{H}$ captures the post-sign up optimal control problem:

$$\mathbb{H} = \Pi(v_1;v_2;u)f(n) + \lambda(n)Y_n + \mu(Y(v_1;v_2;u;n) - \bar{Y})$$

As before, the necessary first order conditions are obtained from the Pontryagin maximum principle. These yield the following two conditions from proposition 1:

$$f \cdot \left(\Pi_u - \Pi_{v_1} \frac{Y_u}{Y_{v_1}} \right) + \lambda \cdot \left(Y_{u,n} - Y_{v_1,n} \frac{Y_u}{Y_{v_1}} \right) = 0$$
$$\dot{\lambda} = \frac{d\lambda}{dn} = -f \cdot \frac{\Pi_{v_1}}{Y_{v_1}} - \lambda \cdot \frac{Y_{v_1,n}}{Y_{v_1}} - \mu f$$

Proposition 3 characterizes the post sign-up optimal contract.

Proposition 3. *The post sign-up optimal contract is characterized by the following:*

$$\frac{c'(u-n)}{c''(u-n)} = \frac{F(n)}{f(n)} \tag{11}$$
$$u(n_L) = n_L \text{ and } u(n) > n \quad \forall n \in (n_L, n_H] \tag{12}$$
$$v_2 = \bar{Y} \text{ and } v_1(n) = c(u-n) + \int_{n_L}^{n} c'(u(t) - t)dt \tag{13}$$

Proof. See appendix C. □

Post-sign up optimal conditions stipulate that the advertiser pay a fixed sum equaling the influencer's outside option value; and a variable sum that increases in n according to equation (13). Moreover as equation (12) suggests, the influencer overstates her number of followers.

Except for the influencer at the lowest end of the interval whose number of followers are $n = n_L$, *all* other types resort to buying fake followers. The advertiser is aware of such high levels of faking yet tolerates it since such a behavior guarantees efficiency and helps the advertiser ordinally rank and thus distinguish high type influencers from low types. Indeed, in the absence of such faking, influencers cannot signal their true type credibly. A payment mechanism that ignores that widespread faking is necessary for efficiency will not be incentive compatible and will leave the advertiser no way to detect influencers with high number of true followers from those with low true follower counts.

4.2.1 Implementability and sufficiency

In our setup, when $Y_{v_1} > 0$, implementability requires that:

$$\frac{\partial}{\partial n} \left(\frac{Y_u}{Y_{v_1}} \right) \cdot \frac{du}{dn} < 0$$

The first term is essentially the Spence-Mirrlees single-crossing condition. Since $Y_u = -c'$ and $Y_{v_1} = 1$, the first term reduces to $-c'' < 0$. This implies that $u' > 0$ is necessary for implementability. For a quadratic cost function of buying fake followers, this is satisfied if $\frac{d}{dn}\big(F(n)/f(n)\big) > 0$.

For sufficiency, it is enough that $Y_n > 0$ which in turn guarantees that the post sign-up participation constraint binds only at $n = n_L$ leading to $Y(n_L) = \bar{Y}$; and $Y(n) > \bar{Y} \quad \forall n > n_L$. This is ensured if $c' > 0$.

4.3 Illustration: uniform distribution with quadratic costs of faking

Consider the class of quadratic cost functions:

$$c(u-n) = \alpha \cdot (u-n)^2 + \beta \cdot (u-n) + \gamma$$

From the conditions imposed on this cost function: $c(0) = 0, c'(0) = 0, c'' > 0$, the only admissible quadratic cost functions are:

$$c(u-n) = \alpha \cdot (u-n)^2, \qquad \alpha > 0$$

From proposition 3, equation (11) yields:

$$u(n) = n + \frac{F(n)}{f(n)}$$

For uniformly distributed $n \sim U[n_L, n_H]$

$$u(n) = 2n - n_L$$

The variable payment is then:

$$v_1(n) = \alpha(u-n)^2 + \int_{n_L}^{n} c'(u(t)-t)dt$$
$$v_1(n) = \alpha((2n-n_L)-n)^2 + 2\alpha\int_{0}^{n-n_L} zdz$$

leading to the following optimal variable payment schedule:

$$v_1(n) = 2\alpha(n-n_L)^2 \tag{14}$$

Thus the inflation function is affine and increases with n; and the variable payment assumes a quadratic functional form. Additionally, the cost of faking assumes the following form:

$$c(u-n) = \alpha \cdot (n-n_L)^2$$

Hence the ratio of the variable payment to faking cost is:

$$\frac{v_1(n)}{c(u-n)} = 2$$

5 Discussion

Given the analytical results in the previous section, we highlight some important observations.

5.1 The extent of fraud

In the illustrative example in section 4.3, the optimal display function for the influencer takes the form:

$$u(n) = 2n - n_L$$

This leads to an interesting observation. Consider the influencer of the highest type $n = n_H$. This influencer has the maximum follower count n_H but will display the following:

$$u(n_H) = 2n_H - n_L = n_H + (n_H - n_L) > n_H$$

In other words, the number of followers displayed is more than the maximum possible! Clearly, the advertiser, who knows that the true number of followers cannot be more than n_H will not believe this overstated follower count and will conclude correctly that the influencer is faking. Moreover, such behavior is not limited to the influencer of the highest type. Except for the influencer with $n = n_L$ followers, all types overreport their true follower count; and in particular, all influencers with number of followers $n > (n_L + n_H)/2$ display a follower count strictly higher than the maximum possible count n_H. In fact, from the display function $u(n) = 2n - n_L$ it is clear that the displayed follower count increases in n—i.e., an influencer of higher type (more true followers) fakes more than another of lower type (fewer true followers). Why then should the advertiser tolerate such egregious faking?[10]

The reason why the advertiser allows such obvious faking is related to the efficiency of the optimal contract since such a behavior helps it to ordinally rank, and thereby distinguish between privately informed influencers of different types. In fact, in the absence of such faking, an influencer cannot signal her true underlying type credibly. Since for a given displayed follower count, the costs of faking are different for different types, the advertiser can harness this variation to screen an influencer with high underlying true follower count from one with lower true following. However, the efficiency of the optimal contract comes at a price since the advertiser needs to pay a variable rate to the influencer which is over and above her outside option value.

5.2 The advertiser's revenue function

An implicit assumption in our model (indeed, in any general contract theoretic model) is that the participation constraint of the advertiser is satisfied. If this were not so, it will not be worthwhile for the advertiser to draw up a contract in the first place. We have assumed that the revenue function of the advertiser from reaching out to n followers of the social media influencer is $A(n)$. However, this term in the advertiser's profit function is not subject to maximization since the arguments over which it maximizes its profits are the functions v_1, v_2 and u. In other words, we assume that the revenue function of the advertiser is a fixed, exogenous function of n.

In order for the (implicit) participation contraint of the advertiser to be satisfied, its profit at the optimal payment schedule must be positive. This condition imposes constraints on the class of admissible revenue functions. In particular, at the optimal v_1^*, v_2^*, u^*:

[10]While this observation may seem puzzling, we note that such results have also been noted in Maggi and Rodriguez-Clare (1995) and Crocker and Morgan (1998).

$$\Pi(v_1^*, v_2^*, u^*) = A(n) - v_1^*(n) - v_2^* \geq 0$$

The above condition necessitates, for example, that if the payment rate to the influencer rises quadratically in n—as is the case in the illustration in section 4.3, equation (14)—the growth rate of $A(n)$ must not be sub-quadratic. More specifically, one can further express constraints on the revenue function in terms of the primitives of our model by noting that $v_2^* = \bar{Y}$ and $v_1^*(n) = c(u^* - n) + \int_{n_L}^{n} c'(u^*(t) - t)dt$.

Hence while the revenue function is fixed and exogenous, it must belong to a certain class of functions with growth rate at least as high as that of the variable payment. In other words, the monetary rewards of reaching out to n followers of the influencer must be sufficiently high for the advertiser in order to offer a contract to the influencer.

5.3 Implications for the advertiser

From the perspective of the advertiser, it is better to find potentially influential or popular social media influencers before they have signed up on a given platform. If the advertiser is able to do so successfully, it will extract all the surplus and limit the payment of the influencer to a fixed sum equalling her outside option value. Even further, the advertiser will be able to elicit truthfulness from the influencer who will not resort to falsifying her subsequent follower count. In such a scenario, there is no informational advantage that the influencer can exploit and by setting the optimal contract wherein there is no variable payment made out to the influencer, the advertiser reaps the benefit of making lower payments while at the same time eliminating fraud.

From a practical viewpoint this situation is not dissimilar to scouting for football players (say). When scouting agents are able to spot "diamonds in the dust"—young players who show great potential before they have been discovered by the sporting world at large—they are able to reap the maximum benefits for their team. Indeed, once such players and their potential value becomes well known, in order to lure them to change teams, vast sums of payment may be needed.

We acknowledge however, that finding such potentially high value social media influencers may not be simple. In particular, our paper offers no special insight into how to spot potential influencers with authentic mass-appeal before they become popular. Our model only suggests that if the advertiser can spot such influencers relatively early, they can extract all the surplus; and if they miss the boat and need to draw from a pool of existing social media influencers, they need to pay much more and tolerate high levels of faking.

5.4 Concluding remarks

While our work addresses an important concern facing online advertisers today, there are many related problems about which our model sheds no light. For example, our model does not take into account the possibility that buying fake followers could attract more genuine followers. We also do not account for innately honest influencers who will never buy fake followers irrespective of incentives to do so, in the manner of Mishra and Prasad (2006). Additionally, the more complex problem wherein the number of true followers is endogenous and optimal contracts incentivize influencers to increase their popularity are also beyond the scope of our model.

Our approach of modeling the fake follower problem in a contract-theoretic manner offers some important insights. We demonstrate that the influencer can be curtailed from purchasing fake followers in a pre sign-up, but not post sign-up setting. In the latter case, we show that influencers with higher genuine follower counts will buy more fake followers, and the optimal contract must account for this. Finally we offer motivations for advertisers to scout for potential influencers with authentic popularity beyond the narrow confines of social media space, since by identifying their true potential before it becomes commonly known, they can reach out to their target audience most cost-effectively.

Appendices

A Proof of Proposition 1

We recall the first order Pontryagin conditions:

1. Optimality condition:

$$\max_u \mathbb{H} \qquad \forall n \in [n_L, n_H]$$

2. Equation of motion for state:

$$\frac{dY}{dn} = \frac{\partial \mathbb{H}}{\partial \lambda} = Y_n$$

3. Equation of motion for costate:

$$\frac{d\lambda}{dn} = -\frac{\partial \mathbb{H}}{\partial Y}$$

4. Transversality condition for state:

$$\lambda(n_H) = 0$$

A.1 Optimality condition

Since the control function is $u(\cdot)$ the derivative of $\mathbb{H}$ with respect to $u(\cdot)$ must be 0, yielding:

$$\frac{d\Pi}{du} \cdot f + \lambda \cdot \frac{dY_n}{du} + \mu \cdot f \frac{dY}{du} = 0 \tag{15}$$

We note that:

$$\frac{d\Pi}{du} = \Pi_{v_1} \frac{\partial v_1}{\partial u} + \Pi_u$$

and

$$\frac{dY_n}{du} = \frac{\partial Y_n}{\partial v_1} \cdot \frac{\partial v_1}{\partial u} + \frac{\partial Y_n}{\partial u} \cdot 1$$

and

$$\frac{dY}{du} = \frac{\partial Y}{\partial v_1} \cdot \frac{\partial v_1}{\partial u} + \frac{\partial Y}{\partial u} \cdot 1$$

Reinserting the above in (15):

$$f \cdot \left(\Pi_{v_1} \cdot \frac{\partial v_1}{\partial u} + \Pi_u\right) + \lambda \cdot \left(Y_{v_1,n} \cdot \frac{\partial v_1}{\partial u} + Y_{u,n}\right) + f \cdot \mu \cdot \left(Y_{v_1} \cdot \frac{\partial v_1}{\partial u} + Y_u\right) = 0$$

Additionally, at the optimal, since $\frac{dY}{dn} = \frac{\partial Y}{\partial n}$

$$\frac{\partial Y}{\partial v_1} \cdot \frac{dv_1}{dn} + \frac{\partial Y}{\partial u} \cdot \frac{du}{dn} = 0$$

leading to:

$$\frac{\partial v_1}{\partial u} = -\frac{Y_u}{Y_{v_1}}$$

We now express the optimality condition for the pre sign-up optimal control problem as:

$$f \cdot \left(\Pi_u - \Pi_{v_1} \frac{Y_u}{Y_{v_1}}\right) + \lambda \cdot \left(Y_{u,n} - Y_{v_1,n} \frac{Y_u}{Y_{v_1}}\right) = 0 \tag{16}$$

A.2 Equation of motion for costate

$$\frac{d\lambda}{dn} = -\frac{\partial \mathbb{H}}{\partial Y}$$

Consider the right hand side:

$$\frac{\partial \mathbb{H}}{\partial Y} = \frac{\partial \mathbb{H}}{\partial v_1} \cdot \frac{\partial v_1}{\partial Y} + \frac{\partial \mathbb{H}}{\partial v_2} \cdot \frac{\partial v_2}{\partial Y} + \frac{\partial \mathbb{H}}{\partial u} \cdot \frac{\partial u}{\partial Y}$$

Since v_2 does not vary with Y its partial derivative is 0. Additionally, from the optimal condition the last term is 0 as well. Thus

$$\frac{\partial \mathbb{H}}{\partial Y} = \frac{\partial v_1}{\partial Y} \cdot \frac{\partial \mathbb{H}}{\partial v_1}$$

Expanding $\frac{\partial \mathbb{H}}{\partial v_1}$ and substituting it back in the equation of motion for costate, we get:[11]

$$\dot{\lambda} = \frac{d\lambda}{dn} = -f \cdot \frac{\Pi_{v_1}}{Y_{v_1}} - \lambda \cdot \frac{Y_{v_1,n}}{Y_{v_1}} - \mu f \tag{17}$$

A.3 Optimal fixed payment

By definition, at the optimal $v_2 = v_2^*$ is fixed and hence if we consider a modified optimal control formulation without the state equation as constraint, v_2^* remains optimal:

$$v_2^* = \arg\max_{v_2} \left(\int_{n_L}^{n_H} \Pi(v_1; v_2; u) f(n) dn\right)$$
$$\text{subject to}$$
$$\int_{n_L}^{n_H} Y(v_1; v_2; u; n) f(n) dn \geq \bar{Y}$$

The first order condition for optimal v_2 yields:

[11] Assuming $\frac{\partial Y}{\partial v_1} \neq 0$.

$$\frac{d}{dv_2}\left(\int_{n_L}^{n_H} \Pi(v_1,v_2,u)\cdot f(n)dn + \mu \cdot \left(\int_{n_L}^{n_H} Y(v_1,v_2,u,n)f(n)dn - \bar{Y}\right)\right) = 0$$

This may be simplified via the Liebniz Integration Rule: $\frac{d}{dx}\int_a^b f(x,t)dt = \int_a^b \frac{\partial}{\partial x} f(x,t)dt$ to yield the third equation:

$$\int_{n_L}^{n_H} (\Pi_{v_2} + \mu \cdot Y_{v_2})\, f(n)dn = 0 \tag{18}$$

B Proof of Proposition 2

B.1 The optimal inflation function

The three necessary conditions (5), (6) and (7) characterize the pre sign-up optimal contract. For our functional specifications, we get $\mu = 1$ from (7). Using this in (6), we obtain

$$\dot{\lambda} = \frac{d\lambda}{dn} = -f\frac{-1}{1} - \lambda \cdot 0 - 1 \cdot f = 0$$

This implies that $\lambda(n) = \lambda$, a fixed constant. From the transversality condition $\lambda(n_H) = 0$, we get $\lambda = 0$. Reinserting $\lambda = 0$ in (5),

$$f \cdot Y_u = 0 \Rightarrow Y_u = 0$$

Evaluating Y_u:

$$Y_u = \frac{\partial}{\partial u}(v_1(n) + v_2 - c(u-n)) = 0$$
$$-c' = 0$$

This necessitates that $c(u-n)$ is constant. Further, since $c(0) = 0$, we get $c(u-n) = c(0)$ which yields $u(n) = n$. Moreover, since $\mu = 1$ this implies that the inequality (3) binds. Using $Y = v_1(n) + v_2 - 0$ and from the incentive compatibility constraint:

$$\int_{n_L}^{n_H} (v_1(n) + v_2) f(n)dn = \bar{Y}$$

we obtain $v_1(n) = 0$ and $v_2 = \bar{Y}$.

C Proof of Proposition 3

We note that the post sign-up participation constraint is slack and hence the corresponding multiplier $\mu = 0$. Using this in equation (6), we obtain:

$$\frac{d\lambda}{dn} = f(n)$$

which along with the transversality condition $\lambda(n_L) = 0$ yields:[12]

$$\lambda(n) = F(n)$$

where $F(n) = \int_{n_L}^{n} f(t)dt$. Using this value of the costate variable λ in equation (5) we obtain,

$$\frac{c'(u-n)}{c''(u-n)} = \frac{F(n)}{f(n)} \tag{19}$$

At the realization $n = n_L$, since $F(n_L) = 0$, it must be that $c'(u(n_L) - n_L) = 0 = c'(0)$, which immediately implies that $u(n_L) = n_L$. Additionally for any $n \neq n_L$, since the right hand side is positive, it implies that $c'(u(n) - n) > c'(0)$ which yields:

$$u(n) > n \quad \forall n \in (n_L, n_H] \text{ and } u(n_L) = n_L \tag{20}$$

Finally, noting that

$$\int_{\bar{Y}}^{Y} dY = \int_{n_L}^{n} \frac{dY}{dn} dn = \int_{n_L}^{n} \frac{\partial Y}{\partial n} dn$$

$$Y = \bar{Y} + \int_{n_L}^{n} c'(u(t) - t)dt \tag{21}$$

$$v_1(n) + v_2 = \bar{Y} + c(u(n) - n) + \int_{n_L}^{n} c'(u(t) - t)dt \tag{22}$$

In particular, equation (22) suggests that the advertiser pay the agent a fixed payment of $v_2 = \bar{Y}$ and a variable payment equaling $v_1(n) = c(u(n) - n) + \int_{n_L}^{n} c'(u(t) - t)dt$.

[12]This transversality condition is different from the pre-sign up case and features a vertical line boundary condition corresponding to the initial state $Y(n_L) \geq \bar{Y}$ (Chiang, 1992, chapter 3).

References

Bhardwaj, P. (2001). Delegating pricing decisions. *Marketing Science 20*(2), 143–169.

Bolton, P. and M. Dewatripont (2005). *Contract theory.* MIT press.

Chiang, A. C. (1992). *Elements of dynamic optimization.* McGraw-Hill.

Confessore, N., G. Dance, R. Harris, and M. Hansen (2018). The Follower Factory. *New York Times*.

Crocker, K. J. and P. Letizia (2014). Optimal policies for recovering the value of consumer returns. *Production and Operations Management 23*(10), 1667–1680.

Crocker, K. J. and J. Morgan (1998). Is honesty the best policy? Curtailing insurance fraud through optimal incentive contracts. *Journal of Political Economy 106*(2), 355–375.

Crocker, K. J. and J. Slemrod (2007). The economics of earnings manipulation and managerial compensation. *RAND Journal of Economics 38*(3), 698–713.

Crocker, K. J. and S. Tennyson (2002). Insurance fraud and optimal claims settlement strategies. *Journal of Law and Economics 45*(2), 469–507.

De Veirman, M., V. Cauberghe, and L. Hudders (2017). Marketing through Instagram influencers: the impact of number of followers and product divergence on brand attitude. *International Journal of Advertising 36*(5), 798–828.

Dionne, G., F. Giuliano, and P. Picard (2009). Optimal auditing with scoring: Theory and application to insurance fraud. *Management Science 55*(1), 58–70.

Djafarova, E. and C. Rushworth (2017). Exploring the credibility of online celebrities' Instagram profiles in influencing the purchase decisions of young female users. *Computers in Human Behavior 68*, 1–7.

Doherty, N. and K. Smetters (2005). Moral hazard in reinsurance markets. *Journal of Risk and Insurance 72*(3), 375–391.

eMarketer (2017). Attitudes Toward Influencer Marketing Among US Agency and Brand Marketers, Nov 2017 (% of respondents). Technical report.

Hovland, C. I. and W. Weiss (1951). The influence of source credibility on communication effectiveness. *Public Opinion Quarterly 15*(4), 635–650.

IRI (2018). BzzAgent and IRI Find Everyday Influencer Marketing Programs Drive the Highest Return on Ad Spend. Technical report.

Jin, S.-A. A. and J. Phua (2014). Following celebrities' tweets about brands: The impact of twitter-based electronic word-of-mouth on consumers' source credibility perception, buying intention, and social identification with celebrities. *Journal of Advertising 43*(2), 181–195.

Kannan, P. K. and H. A. Li (2017). Digital marketing: A framework, review and research agenda. *International Journal of Research in Marketing 34*(1), 22–45.

Lappas, T., G. Sabnis, and G. Valkanas (2016). The impact of fake reviews on online visibility: A vulnerability assessment of the hotel industry. *Information Systems Research 27*(4), 940–961.

Lewis, R. A. and J. M. Rao (2015). The unfavorable economics of measuring the returns to advertising. *Quarterly Journal of Economics 130*(4), 1941–1973.

Libai, B., E. Muller, and R. Peres (2013). Decomposing the Value of Word-of-Mouth Seeding Programs: Acceleration Versus Expansion. *Journal of Marketing Research 50*(2), 161–176.

Linqia (2018). The state of influencer marketing. Technical report, Linqia.

Luca, M. and G. Zervas (2016). Fake it till you make it: Reputation, competition, and Yelp review fraud. *Management Science 62*(12), 3412–3427.

Maggi, G. and A. Rodriguez-Clare (1995). Costly distortion of information in agency problems. *RAND Journal of Economics*, 675–689.

McCracken, G. (1989). Who is the celebrity endorser? Cultural foundations of the endorsement process. *Journal of Consumer Research 16*(3), 310–321.

McGuire, W. J. (1985). Attitudes and attitude change. *The handbook of social psychology*, 233–346.

Mishra, B. K. and A. Prasad (2004). Centralized pricing versus delegating pricing to the salesforce under information asymmetry. *Marketing Science 23*(1), 21–27.

Mishra, B. K. and A. Prasad (2005). Delegating pricing decisions in competitive markets with symmetric and asymmetric information. *Marketing Science 24*(3), 490–497.

Mishra, B. K. and A. Prasad (2006). Minimizing retail shrinkage due to employee theft. *International Journal of Retail & Distribution Management 34*(11), 817–832.

Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica 47*, 61–73.

Neff, J. (2018). Study of influencer spenders finds big names, lots of fake followers. *AdAge India*.

Padmanabhan, V. and R. C. Rao (1993). Warranty policy and extended service contracts: Theory and an application to automobiles. *Marketing Science 12*(3), 230–247.

Paquet-Clouston, M., O. Bilodeau, and D. Décary-Hétu (2017). Can We Trust Social Media Data?: Social Network Manipulation by an IoT Botnet. In *Proceedings of the 8th International Conference on Social Media & Society*, #SMSociety17, New York, NY, USA, pp. 15:1–15:9. ACM.

Pathak, S. (2017). Cheatsheet: What you need to know about influencer fraud. *Digiday*.

Silverman, C. (2018). Apps Installed On Millions Of Android Phones Tracked User Behavior To Execute A Multimillion-Dollar Ad Fraud Scheme. *BuzzFeed*.

Simester, D. and J. Zhang (2010). Why are bad products so hard to kill? *Management Science 56*(7), 1161–1179.

Simester, D. and J. Zhang (2014). Why do salespeople spend so much time lobbying for low prices? *Marketing Science 33*(6), 796–808.

Sridhar, S., P. A. Naik, and A. Kelkar (2017). Metrics unreliability and marketing overspending. *International Journal of Research in Marketing 34*(4), 761–779.

Stewart, R. (2018). Unilever's Keith Weed calls for 'urgent action' to tackle influencer fraud. *The Drum*.

Sun, B. (2014, April). Executive compensation and earnings management under moral hazard. *Journal of Economic Dynamics and Control 41*, 276–290.

Tellis, G. J. (2003). *Effective advertising: Understanding when, how, and why advertising works*. Sage.

The Economist (2016). Celebrities' endorsement earnings on social media - Daily chart.

Wilbur, K. C. and Y. Zhu (2009). Click fraud. *Marketing Science 28*(2), 293–308.

Yoganarasimhan, H. (2012, March). Impact of social network structure on content propagation: A study using YouTube data. *Quantitative Marketing and Economics 10*(1), 111–150.