Free Ad(vice): Internet Influencers and Disclosure Regulation

Matthew Mitchell*

January 30, 2020

Abstract

Consumers rely on intermediaries ("influencers") such as social media recommendations to provide information about products. The advice may be mixed with endorsement in a way that is unobservable to the follower, creating a trade-off for influencers between the best advice and the most revenue. This paper models the dynamic relationship between an influencer and a follower. The relationship evolves between periods of less and more revenue. The model can provide insight into policies such as the Federal Trade Commission's mandatory disclosure rules. An opt-in policy may be superior: it deregulates influencers who are reaping the rewards of past good advice.

*University of Toronto, http://matthewmitchelltoronto.weebly.com. Thanks to Editor Kathryn Spier, two anonymous referees, Guy Arie, Heski Bar-Isaac, Karen Bernhardt-Walther, Rahul Deb, Hanming Fang, April Franco, Joshua Gans, Tanjim Hossain, Ginger Jin, Aaron Kolb, Elliot Lipnowski, Gerard Llobet, Alessandro Pavan, Xianwen Shi, and participants at many conferences and seminars.

1 Introduction

In many markets in which product differentiation is huge, consumers rely on intermediaries to provide information about options.[1] The internet has both increased the scope of product differentiation,[2] necessitating more search, and lowered the cost of providing advice through blogs and social media. Potential consumers often receive advice without directly paying the source of the advice. The world has more and more free advice.

Frequently the advice is supported through sponsors. Blogs often provide product reviews that seamlessly include paid endorsements. Twitter users provide sponsored recommendations to followers; in the U.S., FTC regulations suggest that the sponsorship should be disclosed, but anecdotal evidence suggests that it often is not.[3] That differentiates this form of advertising from typical media advertisements or paid endorsements, which are largely transparent and explicitly separate from content. Websites such as cnn.com include sponsored content alongside links to news content in a way that blurs these lines, often making the sponsored content difficult to separate from related stories on other parts of the site. Facebook chooses trending topics in a way that can steer users to different products or sponsors. An *influencer* mixes advice with various messages from sponsors in order to earn income from the advice. The small size of each piece of advice makes transferring money in exchange for advice prohibitive; Google alone does more than one trillion searches per year, and celebrities have millions of followers. The influencer's reward for providing good advice is to maintain followers.

This paper models the dynamic relationship between an influencer and a follower in a similar manner to the recent literature on dynamic contracting without monetary payments, especially the repeated trust model of delegation in organizations in Li et al. [2015] and the model of investment financing in DeMarzo and Fishman [2007]. The model is based on a tension between good advice and advertisement. The influencer faces a trade-off. On the one hand, it seeks to monetize the advice it gives, possibly by biasing advice

[1] An estimate of the number of products on Amazon alone is over 300 million.

[2] The Amazon estimate is approximately 2000 times the number of products at a Walmart supercenter (http://corporate.walmart.com/_news_/news-archive/2005/01/07/our-retail-divisions)

[3] http://bits.blogs.nytimes.com/2013/06/09/disruptions-celebrities-product-plugs-on-social-media-draw-scrutiny/?_r=0

toward advertisers. On the other hand, it needs to maintain good advice on average, or following will not be valuable to followers. The goal of the paper is to provide a model of this tension in the dynamic relationship and to use it to study policy issues surrounding the influencer market.

In the United States (via the FTC), as well as in other jurisdictions, one of the most important policy questions involves disclosure regulations. Policy guidance in the model is different from that in a standard model in which advice is directly paid for rather than motivated by the possibility of future attention. The model highlights an important consequence of regulation: regulating ads, even if it reduces the current temptation to bias advice, also may reduce the future reward for currently providing good advice by reducing the reward from those future ads. This dynamic effect may make disclosure regulation less effective. The model suggests alternative regulations that are superior, including opt-in disclosure regulation.

In the model, the relationship alternates between periods in which the influencer monetizes the opportunity to advise and periods win which it gives unbiased advice, in a "reap and sow" cycle. A sufficiently long period without good advice causes the relationship to break up permanently. The relationship can be summarized by its expected duration. For low values of the duration variable, good advice is given over advertisement opportunities, and every piece of good advice discretely improves the situation for the influencer. This is the sow period. The duration falls if good advice does not arrive. When the duration grows large enough, the follower can no longer offer enough of an increase in duration to incentivize good advice, and the influencer reaps the value of the past good advice by using the advertisement technology. Although the pattern in the model is extreme, it suggests a natural tendency for experienced, successful influencers to have more ability to bias information toward sponsors, compared to new or struggling influencers.

The optimal contract is solved by first positing that the duration of the relationship going forward is a sufficient statistic for the contract following any history. This variable is sufficient because the influencer's payoff is increasing in this duration: the longer the pair will be together, the more the influencer can earn. The follower's relationship is not monotonic. On the one hand, the relationship generates value for the two parties in total, so a longer relationship generates more value. However, the *share* of that value going to the follower declines with the length of the relationship. The follower, therefore, faces a cost of rewarding the influencer longer. The optimal contract economizes on that cost, while still incentivizing good advice when possible.

To understand a key intuition for the model, consider the impact of changing the return to advertising for the influencer. For a given duration of following, lowering the returns to the advertising technology by a constant fraction (like a tax on the influencer's profits) has no impact on the following or advising behavior. The reason is that the lower returns both lower the current reward to advertising and the future benefit of the follower's future attention. Scaling the value of advertising by a constant fraction simply lowers the influencer's payoff by that fraction. This result comes directly from the central feature of the model: i.e., the reward for good behavior is the future opportunity to operate the technology and not a direct monetary transfer between the parties.

This basic force is at the heart of the key results on disclosure rules such as the ones proposed by the FTC. Suppose that undisclosed and disclosed advice have different efficiency. If disclosure policy impacts both by the same amount, it is like a tax and improves nothing; in fact, it lowers influencer returns that may be passed on to followers. The effect of disclosure is only beneficial if it is sufficiently strong relative to the impact of disclosure on the profitability of the advertising technology; even then, the impact has to be strong enough to offset the costly taxation effect that has no beneficial effect on advice. Mandatory disclosure may be costly.

The FTC's proposed disclosure rules for social media are motivated by the notion that disclosure can improve transaction value. This intuition comes from models in which advice is provided in exchange for money; here, the influencer's reward for providing information is the future ads themselves, which might also be impacted by the disclosure rules. This distinction is why the dynamic model of exchange is essential to understanding the policy impact, including the possibility of lower welfare. The model suggests alternative policies that might improve welfare, such as allowing influencers to opt in to disclosure rules, which can give followers higher returns than blanket mandatory disclosure can. Influencers in the sow period would be expected to publicly opt in (or else not be followed), while influencers with good track records would opt out and get the full value of their advertising technology. Such a policy can improve followers' welfare even when mandatory disclosure cannot, as it simultaneously strengthens incentives for the influencers who are expected to maximize good advice and, at the same time, makes the technology by which good advice is rewarded (the advertisements in the reap period) as unconstrained as possible. Making the payoff high in the reap period is essential to making the relationship efficient throughout since it is precisely these rewards that encourage influencers to provide good advice.

The body of the paper is organized as follows. Section 2 introduces the model. The incentive-feasible Pareto frontier is constructed in Section 3. Section 4 then uses the model to study the policy issues of mandatory disclosure and market power. Section 5 discusses an extension of the model of disclosure used in the main analysis, which considers the possibility of incomplete attention by followers and advice that is both advertiser-supported as well as valuable to followers.

1.1 Literature

1.1.1 Dynamic Contracts without Money

The model is a dynamic contracting problem without money and is broadly similar to models in that literature. In particular, the model is most similar to that of Li et al. [2015] and Bird and Frug [2019], who study a dynamic version of a trust game. Following and good advice can be viewed as a form of favor exchange, as in Hauser and Hopenhayn [2008]. The model here differs from favor exchange in that, although favors occur in both directions, private information is one-sided. Such an arrangement is at the heart of papers such as Lipnowski and Ramos [2019]. Rather than payoffs being unknown, as in Lipnowski and Ramos [2019], the feasible set (that is, whether or not good advice can be generated) is private information of the influencer. That element makes the model most like that of Li et al. [2015] and Bird and Frug [2019]. The model here is somewhat simplified, in the sense that the feasible set is either one of two possibilities, and the follower (the principal in their language) has only two choices: follow or not follow. In the finance literature, DeMarzo and Fishman [2007] have a similar trust-model structure. The model here is cast in continuous time, which allows characterization and comparative statics, as well as policy analysis. The model proceeds by describing contracts in terms of a sufficient statistic that bears a resemblance to the experimentation model of Guo [2016] and papers in the patent literature such as Hopenhayn et al. [2006].

1.1.2 Disclosure and Internet Policy

Two recent papers discuss the influencer-follower relationship explicitly. Perhaps the most closely related paper, and one that is a natural complement to this work, is Fainmesser and Galeotti [2019]. That paper models the influencer-follower relationship as static and with complete information: followers know the exact amount of content of each type, sponsored and unsponsored. They take the private information problem studied here as solved by dynamic consideration, whereas this paper models that private information problem explicitly. The focus of Fainmesser and Galeotti [2019] is to generate equilibrium predictions for relative levels of sponsorship across influencers who differ by an exogenous characteristic that can be interpreted as "celebrity status." They also study transparency policies through their different channel. Pei and Mayzlin [2017] also study recommendations by influencers. In their paper, the influencer faces an explicit informational model in persuading a potential consumer. This generates an endogenous limit on the degree of endorsement that the influencer can give before recommendations are no longer followed. In that model, some form of credible commitment to what is and is not endorsed (like an FTC rule) is necessary for the market to function. In contrast, in this paper the market can function even in the absence of this form of commitment due to dynamic concerns.

Although focused on a different application, Inderst and Ottaviani [2012] study a static model of regulating advice, especially in financial markets. In their model, the reason for the adviser to want to give some good advice is exogenous, but the nature of the static relationship is modeled in much more detail. Disclosure can reduce welfare because it undoes the information value that advisers sometimes have. This model complements that one by focusing on the dynamic aspect, with the static impact of disclosure modeled in a more reduced-form way that is consistent with the static effects in Inderst and Ottaviani [2012].

Although many papers have studied ratings systems like the ones commonly employed on the internet, fewer have studied the repeated relationship between follower and influencer studied here. Much of the literature has been focused on search engines, which are an important and related example of free advice, but the focus here is on a different set of influencers with different policy questions. Burguet et al. [2015] model the bias in "organic" results for an optimizing search engine that also shows paid results. Their results focus on the interrelationship between disclosed and undisclosed ads in a static setting. For search engines, Yang and Ghose [2010] and Edelman and Lai [2014] study how the organic side interacts with disclosed, paid search results. Evidence suggests that the two are linked. Yang and Ghose [2010] show that paid advertisements are associated with higher click-through on organic results. Edelman and Lai [2014] directly study the role of Google's display of its own property (flight results) on users' behavior. They show that Google's flight results generate clicks on both Google properties and paid ads, suggesting that Google has at least two channels by which it is incentivized to bias listings. Rayo and Segal [2010] study a static model with commitment to disclosure rules. This model departs from the commitment assumption and instead penalizes undisclosed messages by a fixed amount.

Other papers have studied the integration of search engines and publishers, which brings up related issues of self-serving advice. In particular, de Cornière and Taylor [2014] study the incentives of a search engine and show that bias can result. Taylor [2011] studies the idea that inflated claims might attract additional visitors. Both of these papers study the static environment in greater detail; in that sense, the dynamic model here can be viewed as a natural complement to their work.

Related to online markets and advertising is advertising in two-sided markets more generally. Anderson and Jullien [2015] discuss the use of advertising to support a two-sided market, and Shi [2018] studies the policy issues related to taxation of advertising revenues in these markets. The taxation discussion is related to the reduced-form modeling of disclosure policy in this paper. This paper builds on that literature by expanding the focus to dynamic relationships and advertisements that are potentially hidden within content.

1.1.3 Repeated Advice

Advice is sometimes modeled as cheap talk, as in Crawford and Sobel [1982]. The model here differs from the cheap talk setup in that the bias of the sender determines the sender's payoff, but the sender's action can be only imperfectly monitored. The feedback effect of the paper relates broadly to the literature on reputation as trust in a repeated game, as described by Cabral [2005] and Mailath and Samuelson [2015]. These models of reputation in environments with monetary transactions go back at least to Klein and Leffler [1981]. This model in this paper includes dynamics and cycles of reputation such as those in Liu [2011] and Liu and Skrzypacz [2014]. In a signaling game context, Kaya [2009] discusses a reputation state variable that is similar, in the sense that it summarizes the state and evolves stochastically.

2 The Model

There is an infinite horizon of continuous time, and two players: a follower (who will operate as the principal) and an influencer (who will operate as the agent). The future is discounted by a common discount rate r that is normalized to 1. The follower's choice of whether or not to follow the influencer at any time t is denoted $f_t \in [0, 1]$, where $f_t = 1$ indicates following and $f_t = 0$ is not following.[4] Following is costly to the follower, as it requires forgoing an outside opportunity with flow payoff s, so the outside payoff in any instant is $s(1 - f_t)$. One can interpret this as a cost of paying attention to advice that does not generate a benefit.[5] Assume that f_t is publicly observed, either through a direct measure of following or through an indirect measure via behavior, such as clicks that generate revenue for the influencer, described next.

When being followed, the influencer faces a trade-off between generating advice and generating ad revenue. The more intensively the ad technology is run, the less likely is good advice. Let the influencer's use of ad technology be denoted $a_t \in [0, 1]$. When being followed, the influencer gets a flow payoff λa_t from choosing a_t. Good advice arrives to the follower at Poisson arrival rate $\lambda(a_t) = (1 - a_t)\lambda$.[6] The follower gets a verifiable benefit of 1 from every piece of good advice it receives, so there is no efficiency rationale for good

[4] Mixtures are formally allowed but turn out not to be used at the optimum, and they can be ignored in understanding the central results.

[5] Throughout the analysis, s is taken to be fixed. This is, in a sense, partial equilibrium; for instance, when regulation is studied, regulation might impact the opportunity cost of following, as well, if it changes the next best alternative use of attention. Here, the interpretation is that the opportunity cost of attention to an influencer is not primarily driven by other influencers but, rather, by other uses of time outside of paying attention to influencers.

[6] The linear specification has the feature that the influencer can be interpreted as a strategic exponential bandit arm, where the arm returns a payoff of 1 and the influencer decides whether to keep ($a = 1$) the payoff or share it with the follower. In that interpretation, the influencer is also receiving a Poisson arrival, similar to the follower. This is the case that most closely corresponds to Li et al. [2015], in the sense that the feasible set (does the unit exist to be transferred) is exactly the private information, as in their model. That interpretation is not necessary, however.

advice over monetization through ads: the total payoff to the two parties is λ per unit of time regardless of a_t. The nature of the relationship is driven entirely by sharing these payoffs. Section 4.1 shows that the actions are the same for an ad technology whose payoff is λxa. Assuming that good advice is verifiable simplifies the analysis and is consistent with the notion that the influencer knows the sense in which the advice might be biased. Let following be efficient: $\lambda > 1 > s$.

The choice of a_t is the influencer's private information. Although the follower cannot explicitly observe and punish the influencer taking money, there is implicit punishment associated with the fact that the follower will punish a lack of good advice. The decreasing $\lambda(a_t)$ is the tension between good advice and monetization that generates the potential for inefficiency. Finally, suppose that the influencer needs to receive at least $\bar{W}$ to invest in setting up the advice technology. This will play a role only in determining the initial conditions of the relationship between influencer and follower.

An interpretation is that attention generates traffic for external sites, and endorsement by the influencer generates traffic. However, there is a tension between the sites that most want traffic (and, therefore, are willing to pay the most) and the ones that will generate good experiences for consumers. Below, we consider several extensions to this basic structure, which have interesting implications but do not change the central economics of the benchmark model. Section 4.1 allows for total surplus to depend on the level of the ad technology, so that, in particular, the ads might reduce total surplus. In Section 5, the advice technology is modified so that there is not a pure tradeoff between ad revenue and good advice; rather, some good advice might also be monetizable, and so the follower reduces attention with the current quality of advice.

3 The Dynamic Relationship

Assume that, at the outset, the follower can choose an entire public-history dependent path for f_t, which will be subject to a commitment constraint (described formally below), that guarantees that both players are better off at each point in time than in permanent reversion to autarky.[7] In particular,

[7]This is analogous to computing the principal-optimal public perfect Nash equilibrium of the game without commitment. Since each player can unilaterally deviate and get the static Nash payoff regardless of the other player, there trivially can be no Nash equilibrium

$f_t(h_t)$ is a function of the public history h_t, where h_t includes the history of f for all dates up to t, and a list of all dates at which good advice was received. It turns out to be sufficient in such a case to consider contracts in which, for any history, a sufficient statistic is the future discounted units of time during which the influencer will be followed, denoted d_{h_t}. In other words,

$$d_{h_t} = E\left(\int_0^\infty e^{-j} f_{t+j}(h_{t+j})dj|h_t\right),$$

where the expectation operator is taken over future histories h_{t+j}. This description of the contract in terms of d is later shown to be without loss of generality; for now, one can consider this class of contracts (those summarized by d for any history) to be a constraint on the contracting environment, which later will be shown to not bind. When unambiguous, the duration after a history will just be written as d_t or simply as d, and f_t will be written without its history argument.

The variable d at any time period can be defined recursively in terms of the current period f and a (and so subscripts are suppressed) by

$$d = f(1 + (1 - a)\lambda(d^+ - d)) + \dot{d}, \tag{1}$$

where d^+ is the duration the contract calls for if good advice is given in the current period, and $\dot{d}$ is the time derivative of d . The term multiplied by f is the net gain in duration if the influencer is followed: one instant for the current period, plus – with probability $(1 - a)\lambda$ – the gain (or, in principle, the loss, if negative), of $d^+ - d$. The remaining term is the time derivative that occurs at (almost) all instants.

Using this recursive construction of d allows for writing an optimal contract recursively. Indexing the contract by d is also useful because of its close relationship to total surplus. For any d , the total surplus can be written as a function of the payoffs for the follower (V) and influencer (W) as

$$W(d) + V(d) = s + (\lambda - s)d \equiv TS(d) \tag{2}$$

since the total surplus is s at any time when advice is not sought, and λ per unit of time that it is sought. This facilitates simplification of the follower's Bellman equation. Although (2) does not continue to hold when the ad

worse than static Nash, which is autarky. Therefore, this corresponds to the strongest threat that could possibly sustain any equilibrium.

technology has a different rate of return from the advice, the construction of this contract turns out to be useful in that context as well. Commitment requires that both players are always receiving at least the autarky payoff. This is no constraint for the influencer since all paths generate non-negative payoffs, but is relevant for the follower and requires that $V(d) \geq s$.

For comparison, if a_t were observable, the influencer could choose a sequence of a_t so that $f_t = 1$ for all t – i.e., $d = 1$, so $V + W = \lambda$, subject to the commitment constraint that $V \geq s$. To see this, first, imagine that the follower follows only if the influencer never advertises. It is then (weakly) optimal for the influencer to always choose $a = 0$; the follower gets $V = \lambda$ and the influencer gets $W = 0$. On the other hand, suppose that the follower follows as long as $a = s/(1 + \lambda)$. In this case, it is again optimal for the influencer to choose the recommended a; the follower gets s and the influencer gets $\lambda - s$. Values of a in between can be used to generate the rest of $V + W = \lambda$. Departures from the full information Pareto frontier when $d < 1$ are, therefore, purely due to information asymmetry in the choice of a.

3.1 Recursive Formulation of the Optimal Contract

The next step is to characterize the set of possible payoffs in the contract. This is done by computing follower-optimal allocations for a given d and, therefore, tracing out the payoff frontier by making a a choice variable of the follower subject to incentive compatibility. The contract will also be characterized by an initial duration that determines the surplus division, as discussed below.

The recursive problem, according to the principle of optimality, for any d is[8]

$$V(d) = max_{a,f,d^+}(1 - f)s + f(1 - a)\lambda(1 + V(d^+) - V(d)) + V'(d)\dot{d} \quad (3)$$

subject to incentive compatibility of a (described below) and the delivery of d according to the promise-keeping constraint (1), as well as the commitment constraint that $V(d^+) \geq s$.[9] Denote the solution to this problem by $a(d)$ and

[8] For the functions W and V, derivatives are denoted with primes – i.e., $V'(d)$. The derivatives $V'(d)$ and $W'(d)$ can always be interpreted as the appropriate left- or right-hand derivative given the sign of $\dot{d}$.

[9] There are also domain restrictions on d, a, and f (that they lie between zero and one).

$f(d)$. The influencer's payoff given the solution is

$$W(d) = f(d)\lambda(a(d) + (1 - a(d))(W(d^+) - W(d)) + W'(d)\dot{d}.$$

The influencer's choice of a can, therefore, be written as

$$max_{a\in[0,1]} a\lambda + (1-a)\lambda(W(d^+) - W(d)).$$

Incentive compatibility for a is, thus,

$$W(d^+) - W(d) \begin{cases} \geq 1 & \text{if } a(d) = 0 \\ \leq 1 & \text{if } a(d) = 1 \\ = 1 & \text{if } 0 < a(d) < 1. \end{cases} \tag{4}$$

An important step in showing the nature of a solution is to determine which durations satisfy the commitment constraint $V(d) \geq s$. The following shows that this set is an interval. This is intuitive since the value function turns out to be concave, but this proof does not rely on concavity.

Lemma 1. *Suppose that there is a feasible plan that has $f = 1$ for duration $\bar{d}$. Then, for all $d < \bar{d}$, there exists a feasible plan where $f = 1$ for duration d. Moreover, for the largest feasible duration d, $V(d) = s$.*

The result implies that the range of d that is not feasible is an interval $(\bar{d}, 1]$. It is immediate that $V(\bar{d}) = s$ since if it were more, then there would be a feasible plan for some $d > \bar{d}$: let $f = 1$ and $a = 1$ until d falls to $\bar{d}$. For d close to $\bar{d}$, this makes almost as much as $V(\bar{d})$.

3.2 The Pareto Frontier

The solution to the problem relies on concavity of V, which is essential since it guarantees that V can be computed when the IC constraint (4) binds for $a(d) < 1$. While the proof of the main characterization describes concavity and the binding IC constraint more formally, the model is tractable enough that some intuition can be obtained without the complete argument.[10] Take

To keep the notation simple, these are not explicitly included, but the discussion below always implicitly takes them into account – explicitly when they bind.

[10]The formal proof of this, and concavity itself, is contained in the appendix, as part of the proof of Proposition 2.

some d with follower's value $V(d)$. For $\tilde{d} < d$, a feasible strategy for the follower, which delivers $\tilde{d}$ units of following time, is to wait (with $f = 0$) a fixed interval of time (in discounted terms, $\frac{d-\tilde{d}}{d}$ units of time) and then follow the plan that delivered $V(d)$. The discounted amount of following time is[11]

$$\frac{d-\tilde{d}}{d}0 + \frac{\tilde{d}}{d}d = \tilde{d}.$$

The payoff from such a strategy for the follower, who receives s while waiting and $V(d)$ from the moment that the waiting period ends, is

$$\frac{d-\tilde{d}}{d}s + \frac{\tilde{d}}{d}V(d).$$

But since $s = V(0)$ (if the follower will never follow again, $d = 0$, then the follower gets the outside option s forever), and the maximized value $V(\tilde{d})$ must be at least as high as this feasible strategy, we have:

$$V(\tilde{d}) \geq \frac{d-\tilde{d}}{d}V(0) + \frac{\tilde{d}}{d}V(d).$$

Although this is not a full proof of concavity, it shows that feasible "waiting" strategies can accomplish convex combinations of payoffs.[12]

Concavity implies that the value can be computed with the IC constraint binding – i.e., when $a(d) < 1$, $W(d^+) - W(d) = 1$. Intuitively, suppose that d^+ is more than necessary for $a < 1$. To maintain the promise of d, that means that $\dot{d}$ must be lower than if the IC constraint binds. This is effectively a randomization of future duration (based on whether or not good advice arrives given a); such a randomization is not beneficial to the follower when V is concave. Moreover, concave V immediately implies increasing W: longer duration is unambiguously good for the influencer. The full characterization shows that, whenever $a(d) < 1$, the value function is strictly concave and, in fact, the binding incentive constraint is optimal.

Imposing the IC constraint, in turn, helps explain the follower's incentives for choosing a. When the IC constraint binds, the difference between $V(d^+)$ and $V(d)$ can be rewritten using (2):

$$\begin{aligned} V(d^+) - V(d) &= (\lambda - s)(d^+ - d) - (W(d^+) - W(d)) \\ &= (\lambda - s)(d^+ - d) - 1. \end{aligned}$$

[11] This can also be verified from (1)

[12] Since $V(d) \geq s$, this policy always delivers at least s in value during the waiting period.

Replacing $V(d^+) - V(d)$ in the follower's problem:

$$V(d) = max_{a,f}(1-f)s + f(1-a)\lambda(\lambda - s)(d^+ - d) + V'(d)\dot{d}. \qquad (5)$$

subject to promise keeping, (1), and the incentive constraint that $d^+(d)$ is implicitly obtained from $W(d^+) - W(d) = 1$, as well as $V(d^+) \geq s$. When $a < 1$, the influencer's payoff simplifies to

$$W(d) = f(d)\lambda + W'(d)\dot{d}.$$

Whenever the current duration allows the follower to follow for one additional unit of discounted time, achieving the IC constraint is feasible. Define $\hat{d}$ by

$$W(\bar{d}) - W(\hat{d}) = 1.$$

When $d = \hat{d}$, there is one unit of reward to give if the duration rises to the maximum feasible; there will be exactly one unit of reward, and, therefore, the incentive compatibility can only be achieved if $d \leq \hat{d}$. The following characterizes the solution to the follower's problem.

Proposition 2. *The follower's problem in (3) is solved by*

$$\begin{aligned} f(d) &= 1 \text{ if } d > 0 \\ a(d) &= \begin{cases} 0 & \text{if } d \leq \hat{d} \\ 1 & \text{if } \bar{d} \geq d > \hat{d} \end{cases} \end{aligned}$$

for some $\hat{d}$ and $\bar{d}$ where $W(\bar{d}) - W(\hat{d}) = 1$. Moreover, this solution solves the problem among all history-dependent policies.

Linearity in a when $f = 1$ suggests that corners are optimal when following occurs. Consider the total benefit to the follower in motivating $a = 0$ instead of $a = 1$. When $a = 1$, the follower gets nothing when a piece of advice might otherwise have arrived. When $a = 0$, for every arrival, the follower gets 1, plus the change in total surplus $W + V$ that results from changing the duration promise to d^+, minus the change in influencer value. Write the sum of these three components as

$$1 + TS(d^+) - TS(d) - \left(W(d^+) - W(d)\right),$$

Since the IC constraint binds, the difference in W is exactly 1 and the benefit to the follower is the increase in future total surplus. Since total surplus in

(2) is increasing, this is positive; therefore, whenever feasible, the follower incentivizes $a = 0$. Since W is increasing, and d^+ can be no higher than 1 , $a < 1$ is not feasible for high enough d. In particular, if $d > \hat{d}$, where $W(1) - W(\hat{d}) = 1$, it is impossible to offer enough future duration to have $a = 0$. When d grows too high to feasibly get good advice, the influencer is rewarded with ads, setting $a = 1$. The follower incentivizes good advice fully whenever feasible and stops following only when $d = 0$; i.e., severance is permanent.

Although the bang-bang nature of the solution is the result of linearity, the basic idea is intuitive. Take any structure where breakup is costly, so that total surplus is increasing in d, but is possibly a concave function rather than a linear one. Suppose that the incentive constraint requires that the influencer be rewarded in the form of higher d, by enough to offset the lost ad revenues. Take the potential for revenue to be some constant. Then, the follower's net benefit is the change in total surplus minus the amount of the ad revenue. Since the change in total surplus is declining in d, so does the net return to good advice for the follower, after accounting for the cost of incentives. Therefore, as here, good advice will naturally occur for low values of d.

Corollary 3. *$V(d)$ is strictly concave on $[0, \hat{d})$ and linear on $[\hat{d}, \bar{d}]$*

Strict concavity is important because, for histories in a strictly concave portion, mixtures (i.e., a between zero and one) make the follower strictly worse off. Whenever V is strictly concave, a and f are pinned down by the characterization in Proposition 2, and the evolution from (1) pins down the duration. However, since the value function is linear for d in $[\hat{d}, \bar{d}]$, the contract could equally well randomize, so long as all realizations keep the contract in the linear region. For instance, it would be equally good for the follower to pick, based on the d, a constant Poisson "escape" rule, where $a = 1$ prevails until some public random signal is observed, at which point the contract would move to $d = \hat{d}$ and ads would stop. The random "escape" signal arrives at a rate that is pinned down by d and $\hat{d}$. Until the "escape" signal occurs, $a = 1$, and then everything returns to the strictly concave region afterwards. However, since the randomization must remain in the feasible region, and $d = 1$ is not feasible, full entrenchment in the sense of Li et al. [2015] does not occur. The only absorbing state is $d = 0$, and the contract arrives there with a sufficiently long period of no good advice arriving, following the evolution in (1).[13]

3.3 Initial d

The final element of the contract is the initial duration d_0. Recall that the influencer is assumed to need at least $\bar{W}$ to invest in setting up the advice technology. While there are various ways to think of the initial conditions, we focus on the one that maximizes the follower's payoff. This is partially for expositional simplicity and partially to keep the focus of the policy analysis away from the surplus division between the follower and influencer at date zero, on which the model has less to say. The initial condition that maximizes the follower's payoff is

$$d_0 = argmax_{d:W(d)\geq\bar{W}} V(d).$$

Often relevant to the comparative statics is whether or not the constraint in that problem binds. The initial condition will be called unconstrained when the constraint is slack – i.e., $\bar{W} < W(d^*)$, where $d^* = argmax_d V(d)$. One can think of the cases as a comparative static on the parameter $\bar{W}$. In the unconstrained case, $\bar{W}$ is low, perhaps because they are in short supply; changes in the contracting environment (for instance, via the policies considered below) can improve the follower's payoff without regard to the influencer's payoff (although they may increase or decrease W, since for some d, W and V are both increasing). On the other hand, in the constrained case, $\bar{W}$ is high enough (perhaps because good advice is scarce) that any change in the contract has to fully compensate the influencer.

The market completely breaks down if $max_d V(d) < s$ since the arrangement would be worse for followers than just taking the outside option forever, so it is assumed throughout that $\bar{W}$ is small enough that $max_d V(d) > s$.

[13]Even in the weakly concave region, where future duration promise is not pinned down uniquely, it is not feasible to have anything but $a = 1$ (since a big enough duration promise to get lower a is impossible) and since duration evolves deterministically for $f = 0$, all choices of f for a given d have the same return, which implies that the payoff is highest when $f = 1$. In other words, a and f are still pinned down for a given d.

4 Policy Impact on the Dynamic Relationship

4.1 Taxing Ads

Thus far, the model has so far made the total surplus independent of a. It might seem more natural that ads are inefficient; this section verifies that the basic structure of the contract is as described above. Moreover, this allows the model to address what the impact would be of taxing ad revenue. This might be one way that a disclosure policy might impact the relationship.

Let ads generate $x\lambda a$. This allows for the possibility that ads produce less surplus than good advice ($x < 1$) and, therefore, have a cost in terms of total surplus. Define the influencer's payoff by

$$W_x(d) = f(d)\lambda(xa(d) + (1 - a(d))(W_x(d^+) - W_x(d)) + W_x'(d)\dot{d}, \qquad (6)$$

and let $V_x(d)$ be defined as in the problem in Section 3, but where $W(d)$ is replaced by $W_x(d)$.

One might imagine that taxes on monetization would discourage monetization and encourage good advice. The next lemma shows that this isn't true: nothing about the allocation changes. The reason is that this tax reduces both the current incentive to run ads and the future payoff from improving the relationship, as the payoff comes in the form of future ads.

Lemma 4. *Suppose that the influencer's payoff from the advertising technology is $x\lambda a$ for all d. Then, $W_x(d) = xW(d)$ and $V_x(d) = V(d)$.*

The tax impacts the current incentives and the future returns to ads symmetrically. When the initial d_0 is unconstrained, the contract is unchanged as a result of the tax. When d_0 is constrained, the tax must be passed on to followers. However, in either case, the sum of follower and influencer well-being is reduced when x is lower:[14]

Proposition 5. *Let $d_0(x)$ be the initial duration when the return is x. Then $W_x(d_0(x)) + V_x(d_0(x))$ is non-decreasing in x.*

[14] If one interprets x as a tax, total surplus *plus tax revenue* increases in the constrained case since it generates higher initial duration; tax revenue plus the payoffs of the follower and influencer must be given by (2).

4.2 A Reduced-Form Model of Disclosure

An important policy consideration in these relationships is whether there would be any benefit from mandating disclosure of monetary compensation by influencers. Agencies such as the U.S. Federal Trade Commission and several jurisdictions in the EU have actively pursued such a policy. In this section, the model is augmented to include a disclosure decision in a reduced-form way that is consistent with static disclosure papers, such as Inderst and Ottaviani [2012]. Section 5 provides a discussion of a richer model of disclosure that can deliver the reduced form used here.

Regulation by a body like the FTC is assumed to be an additional technology not feasible for influencer or follower. This is essential since, otherwise, the optimal contract would subsume optimal use of the regulation technology. This assumption is consistent with the notion that the scale required to use such a technology makes it difficult for an individual follower to operate it.

An alternative interpretation of the disclosure regulations is that they are imposed by the platform on which the influencers operate. It seems natural that the platform might have at least as good a technology to regulate content. A separate question is whether or not the platform would have the incentive to maximize total surplus, as a regulator might; it seems at least plausible that a platform would want to seek this goal.

We assume that disclosed and undisclosed ads might have different returns and, in particular, that disclosure might lower the return to the ad technology.[15] In the case of disclosed ads, the return might be impacted by the fact that the disclosure could make the ad less effective in terms of net value between the influencer and follower. This is consistent with the fact that, without disclosure rules, endorsements on Twitter and other social platforms are rarely disclosed. It is also consistent with the idea that disclosure might have direct costs: Twitter's character count means that characters used in disclosure are costly.[16]

[15] According to Jaclyn Johnson, president of creative services at Small Girls PR: "...bloggers we work with say, 'I want you to know, my engagement on posts that are tagged "#ad" or "#spon" get lower engagement than if that wasn't there."' http://www.nytimes.com/2016/08/30/business/media/instagram-ads-marketing-kardashian.html?emc=eta1&_r=0.

[16] Section 5 discusses the role that attention on the part of the follower might have on total surplus when some advice might be both paid and valuable to followers. There, valuable (but paid) advertisements might be lost if consumers can reduce attention as a

The idea that disclosure generates costs for both sides might come out of economic models such as Inderst and Ottaviani [2012]. In that model, disclosure can lower the informativeness of ads because it creates greater disincentive to advertise among more-efficient firms. In a related summary, Inderst [2015] states:

> Various policies can limit the use of commissions or dampen the impact that they can have on advisers' recommendations, such as a cap or an outright prohibition, mandatory disclosure, restrictions on the steepness of incentives, or their mandatory deferral. One of the key insights is that this may however not always increase welfare. In fact, when commissions serve a welfare enhancing role, such as to steer recommendations to more efficient products, such policies may generate or aggravate a problem of under provision of incentives. The positive role of commissions is frequently overlooked notably in policy debate.

Assume that the influencer's chosen level of the ad technology now has two elements: disclosed and undisclosed. The amount of disclosed ads is a_m and is observable. Since a_m is observable, it can be treated as a direct choice of the *follower*. Incentive compatibility will be needed for the separate choice of undisclosed ads a_u. Let $a = a_m + a_u$ be the total ads; $a \leq 1$ as before.

To model the lower return to disclosed ads, let the payoff from disclosed ads be $\lambda m a_m$ with $m \leq 1$. An authority regulates disclosure by imposing a cost on any ads in excess of a_m, so that undisclosed ads return $\lambda u a_u$. The variable u is the policy variable considered by the FTC. One interpretation is that the FTC can intercept a fraction u of all advertisements and force them to be taken down; in fact, this channel has been a common one for the FTC to use in regulating tweets so far.[17]

To gain intuition for the role of disclosed (observable) ads, imagine that only this mode were feasible. Then, the model would have no tension: the follower could allow the influencer to run ads a fraction of the time (for

result of disclosure.

[17]For instance, the famous Ken Bone tweet for Uber following the US Presidential Town Hall was taken down after the FTC said that it was likely in violation of disclosure rules. (https://www.engadget.com/2016/10/13/ken-bone-may-have-violated-ftc-rules-with-uber-tweet/). Fines were threatened, but not implemented in the CS:GO Lotto case (https://www.engadget.com/2017/09/08/youtube-csgo-lotto-fcc-no-fine/) and Warner Bros./Shadow of Mordor, which involved both Twitter and YouTube.

instance, oscillating at high frequency, although that is not necessary), so that the influencer got $\bar{W}$ and the follower got the rest. Any deviation from the preset ad schedule would be met with permanent unfollowing. The specific dynamics would not be pinned down.

The follower always wants to allow recommended ads to be run in their most efficient form:

Lemma 6. *If $m > u$, then $a_u = 0$. If $m < u$, then $a_m = 0$.*

As a result, when regulation is weak (so that undisclosed ads are more profitable than disclosed ones, $m < u$), the impact of u is identical to a tax on ads, as the ad technology always makes λua, as in Section 4.1 with $x = u$. Therefore, according to Lemma 4, it decreases W and has no impact on V for given d, and it cannot improve total welfare once the initial condition is taken into account according to Proposition 5. Such a weak disclosure policy is effectively a burden on monetization that does not benefit followers.

Corollary 7. *If $u > m$, then $V(d)$ is independent of u and $W(d)$ is increasing in u, so $V + W$ is increasing in u.*

On the other hand, if $u < m$, the policy changes $V(d)$ for fixed d since it impacts the incentive constraint differently from the current payoff.

Lemma 8. *Suppose that $u < m$. Then, for all d, $V(d)$ is decreasing in u.*

Strict disclosure is good *at the margin* for followers but bad for influencers; the net impact on welfare is ambiguous. When $u = m$, however, the disclosure has already had a negative welfare effect that has to be overcome in order to make a stricter disclosure standard, with $u < m$, net welfare improving.[18] In other words, a strict disclosure policy must be sufficiently harsh to offset any "taxation" effect it has. Further, there is no reason that welfare needs to be higher when $u = 0$ (an FTC policy that completely eliminated any incentive to run undisclosed ads) compared to $u = 1$, unless m is close enough to one. Moreover, very low values for u might be very costly, or even impossible, to implement. One way to interpret the results here is that incomplete regulation may reduce total welfare.

[18] Usual arguments imply that V is continuous in u.

4.3 Alternative policies

The model suggests alternative policies that could be an improvement. Suppose that disclosure rules applied only to influencers below $\hat{d}$. High d influencers would be free to make the full ad technology return. Then, the follower would get the benefit of the looser IC constraint without the cost of making the reward to good advice lower. This policy could be implemented on an opt-in basis. Suppose that influencers could announce whether or not disclosure rules would apply to them before the follower chose f.[19] For $d < \hat{d}$, the follower would only follow if the announcement were that disclosure rules apply; for $d > \hat{d}$, no such requirement would be imposed. Influencers with low d would announce that disclosure rules apply, and they would be regulated. Influencers with high d would not. The policy could always be implemented as an opt-in arrangement.

This can be incorporated in the model by adding an observable variable y that indicates opt-out if $y = 1$ and opt-in if $y = 0$. When they opt-out they can have tweets be undisclosed but make 1 instead of u. This change is always unambiguously better than a weak disclosure rule for consumers:

Proposition 9. *Suppose that $u > m$. Then, the follower is better off under the opt-in rule than with regulation of all undisclosed ads.*

The ability to be unfettered when an influencer has been very successful increases the incentive to give good advice everywhere else.

5 Discussion: Disclosure when advice and income are not always in conflict

This section adds three key differences to the benchmark model. First, some paid advice might also be beneficial to followers. Influencers often argue that they endorse products and services that they would recommend even in the absence of sponsorship.[20] Second, even when an influencer wants to monetize,

[19] In the Twitter example, this could be part of the influencer's profile information.

[20] Fashion blogger Kim France explained: "I make money on the blog through affiliate linking. This means that when I link to, say, a dress from Nordstrom or Shopbop or another major retailer and you buy it, I get a small commission. There are many, many items included on this blog that are from smaller retailers that aren't part of any affiliate program, however. And I never, ever link to anything I wouldn't want to buy for myself, commission

it has limited ability to find sponsors; not all content when $a = 1$ is an ad, so the current public history does not immediately reveal whether or not a particular post is sponsored. Finally, in the benchmark model, followers paid attention when $a = 1$, even though the current return was negative (net of s); they did so because $V(d)$ was still greater than s. In this section, we will assume that whenever the current-period return to following is negative, the follower pays less than full attention.

To incorporate good advice that is also sponsored, let good advice arrive at a positive rate $p\lambda$ when it is sponsored, instead of zero, The interpretation is that a fraction p of sponsored advice is also good advice.[21] Content that is sponsored, therefore, generates good advice at rate λp, and content that is unsponsored generates good advice at rate λ, as before. Since it may be impossible to have sponsored content at all times, let ρa be the sponsored content at any given time. This means that, when paying attention to all content, a follower gets good advice at rate

$$\lambda(1 - \rho a(1 - p))$$

instead of $\lambda(1 - a)$ in the benchmark model, where $\rho = 1$ and $p = 0$. The current net return on all content, then, for the follower is $\lambda(1 - \rho a(1 - p)) - s$. For sponsored content, it is λp.

Suppose that the follower does not give full attention when the current return to attention is non-positive. Simply put, followers pay less attention when their instantaneous return to paying attention is low; one can think of this as following but ignoring. Ignoring causes the follower to lose any benefit from the advice, and, instead generates a benefit $0 < b < s$ – i.e., it is less than the outside option of unfollowing. If the influencer is being followed but ignored, it can get only m rather than 1 from each unit of sponsored content. One interpretation of the lost revenue is that if the advice is a

or no commission." http://www.girlofacertainage.com/2016/07/25/your-every-question-answered/. Similarly, Google contends that it links to its own products on searches not because of revenue, but to enhance the user's experience.

[21] There also may be revenue streams that depend on f but not on any unobserved choice by the influencer: influencers might get revenue from a source outside of the advertising channel that generates additional value of followers such as separate, disclosed and verifiable ads that run alongside the advice. Celebrities may inherently value followers for professional and personal reasons. Mitchell [2019] shows that such a value unambiguously benefits followers (since it increases influencers' incentives to keep duration high) but has an ambiguous impact on influencers.

product endorsement, then the consumer gets value 1 with probability p, but the "click rate" on disclosed paid ads falls, leading to lower surplus to be shared between the influencer and seller. Since $b < s$, the follower always wants to pay attention to unpaid advice, which gives a unit benefit at rate λ. However if $p\lambda < b$, the follower would rather ignore sponsored advice since it produces surplus with only probability p.

If we assume that $p < b/\lambda$ but $(1 - \rho(1 - p)) > b/\lambda$, then if $a = 1$, the influencer is still worth paying attention to given the fraction ρ of paid advice and the fraction p of paid advice that is good. If the follower knows that a particular piece of advice is sponsored, though, it leads to less attention. From the influencer's point of view, disclosure is a tax: each ad makes m times as much revenue under a disclosure regime, as in the model presented above. On the other hand, disclosure benefits consumers: the follower gets $\rho b + \lambda(1 - \rho)$ when $a = 1$, instead of $\lambda(1 - \rho(1 - p))$, so followers benefit by $\rho b - \lambda\rho p$. When the benefit to the consumer is small (i.e., λp is close to b), the model (nearly) matches the "disclosure as taxation" assumption used above. More generally, even if followers benefit more substantially from b, disclosure regulation could still make total surplus at instants when $a = 1$ go up or down. Since these instants are, in a dynamic sense, benefiting both parties, reductions to total surplus in these periods may not help consumers, even if they are enjoying some gains from lower attention under disclosure regulation. As stressed in the model, the ad technology serves as both the temptation and the reward for the influencer, so lowering its return can be harmful.

6 Conclusion

This paper has introduced a model of the dynamic interaction between an influencer and a follower. In a market for advice without prices, dynamic incentives come through future attention and advice. The model builds on the approach used in the dynamic contracting literature without monetary transfers to consider industrial organization questions such as regulatory policies for such a market.

A policy that taxes monetization in the advice process does not necessarily change the amount of good advice. Disclosure policy in such an environment can, therefore, be ineffective to the extent that it acts as a tax. A superior policy might involve regulating only selectively, so that the incentive to provide good advice would include the incentive to escape from costly disclosure rules. Such a policy can potentially improve overall welfare, and in particular can benefit followers who receive information.

Many interesting directions could be developed from this starting point. Further research could consider other policies, such as a tax on monetization that funds infrastructure for making relationships less costly (for instance, by making the internet faster). The model could be adapted to include having the follower learn about the rate of arrival of good advice from the influencer, so that the problem could have the experimentation aspect of a traditional bandit problem. Another interesting dimension would be to include equilibrium between many influencers and many followers. Understanding equilibrium arrangements in this sort of dynamic relational contracting environment is, more generally, an interesting avenue for future research.

References

Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica*, 58(5):1041–1063, September 1990. URL `https://ideas.repec.org/a/ecm/emetrp/v58y1990i5p1041-63.html`.

Simon P. Anderson and Bruno Jullien. Chapter 2 - the advertising-financed business model in two-sided media markets. In Simon P. Anderson, Joel Waldfogel, and David Stromberg, editors, *Handbook of Media Economics*, volume 1 of *Handbook of Media Economics*, pages 41 – 90. North-Holland, 2015. doi: https://doi.org/10.1016/B978-0-444-62721-6.00002-0. URL `http://www.sciencedirect.com/science/article/pii/B9780444627216000020`.

Daniel Bird and Alexander Frug. Dynamic non-monetary incentives. *American Economic Journal: Microeconomics*, 11(4): 111–50, November 2019. doi: 10.1257/mic.20170025. URL `http://www.aeaweb.org/articles?id=10.1257/mic.20170025`.

Roberto Burguet, Ramon Caminal, and Matthew Ellman. In Google we trust? *International Journal of Industrial Organization*, 39:44–55, 2015. ISSN 01677187. doi: 10.1016/j.ijindorg.2015.02.003.

Luis Cabral. The economics of trust and reputation: A primer. *New York University and CEPR*, 2005. URL http://pages.stern.nyu.edu/ lcabral/reputation/Reputation_June05.pdf.

Michael G. Crandall and Pierre-Louis Lions. Viscosity solutions of Hamilton-Jacobi equations. *Transactions of the American Mathematical Society*, 277(1):1–1, jan 1983. ISSN 0002-9947. doi: 10.1090/S0002-9947-1983-0690039-8. URL http://www.ams.org/jourcgi/jour-getitem?pii=S0002-9947-1983-0690039-8.

Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1913390.

Alexandre de Cornière and Greg Taylor. Integration and search engine bias. *The RAND Journal of Economics*, 45(3):576–597, sep 2014. ISSN 07416261. doi: 10.1111/1756-2171.12063. URL http://doi.wiley.com/10.1111/1756-2171.12063.

Peter M DeMarzo and Michael J Fishman. Optimal Long-Term Financial Contracting. *The Review of Financial Studies*, 20(6):2079, 2007. doi: 10.1093/rfs/hhm031. URL + http://dx.doi.org/10.1093/rfs/hhm031.

B Edelman and Z Lai. Design of Search Engine Services: Channel Interdependence in Search Engine Results. *Journal of Marketing Research*, 2014. URL http://journals.ama.org/doi/abs/10.1509/jmr.14.0528.

Itay Perah Fainmesser and Andrea Galeotti. The Market for Online Influence. *SSRN Electronic Journal*, aug 2019. ISSN 1556-5068. doi: 10.2139/ssrn.3207810. URL https://www.ssrn.com/abstract=3207810.

Yingni Guo. Dynamic delegation of experimentation. *American Economic Review*, 106(8):1969–2008, 2016. ISSN 00028282. doi: 10.1257/aer.20141215.

Christine Hauser and Hugo Hopenhayn. Trading Favors: Optimal Exchange and Forgiveness. *Carlo Alberto Notebooks*, 2008. URL http://ideas.repec.org/p/cca/wpaper/88.html.

Hugo Hopenhayn, Gerard Llobet, and Matthew F Mitchell. Rewarding Sequential Innovators: Prizes, Patents, and Buyouts. *Journal of Political Economy*, 114(6):1041–1068, dec 2006.

Roman Inderst. Regulating commissions in markets with advice. *International Journal of Industrial Organization*, 43:137–141, 2015. ISSN 01677187. doi: 10.1016/j.ijindorg.2015.05.005.

Roman Inderst and Marco Ottaviani. Competition through commissions and kickbacks. *American Economic Review*, 102 (2):780–809, April 2012. doi: 10.1257/aer.102.2.780. URL `http://www.aeaweb.org/articles?id=10.1257/aer.102.2.780`.

Ayça Kaya. Repeated signaling games. *Games and Economic Behavior*, 66 (2):841–854, 2009. ISSN 08998256. doi: 10.1016/j.geb.2008.09.030.

Benjamin Klein and Keith B Leffler. The Role of Market Forces in Assuring Contractual Performance. *The Journal of Political Economy*, 89(4):615–641, 1981. ISSN 0022-3808. doi: 10.1086/260996. URL `http://www.jstor.org/stable/1833028`.

J Li, N Matouschek, and M Powell. Power Dynamics in Organizations. *American Economic Journal: Microeconomics*, 2015. URL `http://pareto-optimal.com/s/LMP-Submission.pdf`.

E Lipnowski and J Ramos. Repeated Delegation. 2019. URL `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2552926`.

Q Liu. Information acquisition and reputation dynamics. *The Review of Economic Studies*, 2011. URL `http://restud.oxfordjournals.org/content/78/4/1400.short`.

Qingmin Liu and Andrzej Skrzypacz. Limited records and reputation bubbles. *Journal of Economic Theory*, 151(1):2–29, 2014. ISSN 10957235. doi: 10.1016/j.jet.2013.12.014.

George J. Mailath and Larry Samuelson. Reputations in Repeated Games. *Handbook of Game Theory with Economic Applications*, 4(1):165–238, 2015. ISSN 15740005. doi: 10.1016/B978-0-444-53766-9.00004-5.

Matthew Mitchell. Free Ad(vice). 2019.

Amy Pei and Dina Mayzlin. Paid vs Independent Product Recommendation By Bloggers. 2017.

Luis Rayo and Ilya Segal. Optimal Information Disclosure. *The Journal of Political Economy*, 118(5):949–987, 2010. ISSN 00223808. doi: 10.1086/657922.

C Matthew Shi. Effects of an Advertising Tax in Two-sided Markets under Imperfect Competition. 2018. URL https://editorialexpress.com/cgi-bin/conference/download.cgi?db_name=IIOC2018&pa

Greg Taylor. The informativeness of on-line advertising. *International Journal of Industrial Organization*, 29(6):668–677, nov 2011. ISSN 0167-7187. doi: 10.1016/J.IJINDORG.2011.03.001. URL https://www.sciencedirect.com/science/article/pii/S0167718711000269.

S Yang and A Ghose. Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence? *Marketing Science*, 29(4):602–623, 2010. ISSN 0732-2399. doi: 10.1287/mksc.1090.0552.

Proofs

Proof of Lemma 1

Proof. For the plan starting from $\bar{d}$, $V(\bar{d}) \geq s$. For $d < \bar{d}$ let $f = 0$ (and so $\dot{d} = d$) until duration rises to $\bar{d}$. The return to such a plan is

$$\frac{\bar{d} - d}{\bar{d}} s + \frac{d}{\bar{d}} V(\bar{d}) \geq s$$

and therefore constitutes a feasible plan. Let $\bar{d}$ be the largest feasible $\bar{d}$. If $V(\bar{d}) > s$, then either $\bar{d} = 1$, in which case $V(\bar{d}) = s$ (a contradiction that $V(\bar{d}) > s$) or there exists $d > \bar{d}$ with

$$V(d) \geq \frac{\bar{d} - d}{\bar{d}} s + \frac{d}{\bar{d}} V(\bar{d}) > s$$

which contradicts that $\bar{d}$ is the largest duration with $V(d) \geq s$. Therefore $V(\bar{d}) = s$. □

Proof of Proposition 2

The proof of the structure of the optimal contract follows the following steps. First construct a solution to the dynamic program, then show that it is the unique viscosity solution, and therefore describes the optimum, following Crandall and Lions [1983]. To construct a solution to the dynamic program, suppose that V is concave; this implies W is increasing and convex. Then construct V, and verify that under that solution that V is indeed concave. This is accomplished through a series of claims.

For claims A1-A5, suppose that V is concave. Claim A6 then verifies concavity.

Claim (A1). W is increasing and convex

Proof. Since $W(d) = TS(d) - V(d)$, and total surplus is linear, convexity is implied by concavity of V. To show W is increasing it is therefore sufficient to show that it is increasing at zero. Since $W(0) = 0$ (the minimum possible amount), increasing at zero is guaranteed. □

Claim (A2). If $a < 1$, V can be computed for the case where the IC constraint binds.

Proof. Suppose $a < 1$. Then:

$$V(d) = max_{d^+}(1-a)s+ap\lambda(1+V(d^+)-V(d))+V'(d)(d-a(1+p\lambda(d^+-d)))$$

so the derivative with respect to $d^+$is

$$V'(d^+) - V'(d)$$

which is less than zero by concavity of V. Therefore if the IC constraint doesn't bind, d^+ can be reduced until it does without lowering the payoff. □

Claim (A3). Suppose $f = 1$ and $W(d) < \lambda - 1$. Then $a(d) = 0$.

Proof. If $f(d) = 1$ and $a(d) < 1$:

$$V(d) = (1 - a(d))\lambda(\lambda - s)(d^+ - d) + V'(d)(d - 1 - (1 - a(d))\lambda(d^+ - d)))$$

so

$$\begin{aligned} dV/da &= \lambda(d^+ - d)\left((\lambda - s) - V'\right) \\ &= \lambda(d^+ - d)W' < 0 \end{aligned}$$

Therefore either it is optimal to have $a = 0$ or $a = 1$. When $a = 1$, $V(d) = V'(d)(d-1)$, so

$$\begin{aligned} V_{a=0}(d) - V_{a=1}(d) &= \lambda((\lambda - s) - V'(d))(d^+ - d) \\ &= \lambda W'(d)(d^+ - d) > 0 \end{aligned}$$

Therefore $f(d) = 1$ implies $a = 0$ if feasible. □

Combing the fact that W is increasing and the fact that $a = 1$ whenever feasible implies that, for some $\hat{d}$, $a(d) = 1$ for $d > \hat{d}$, and $a(d) = 0$ for $0 < d < \hat{d}$. The next claim establishes that $f(d) = 1$ for all $0 < d < \hat{d}$.

Claim (A4). Suppose $f(d) > 0$ for some $0 < d < \hat{d}$. Then $f(d) = 1$ for $0 < d < \hat{d}$.

Since $a = 0$ in this range, the derivative of the follower's objective for f, letting $x = d^+ - d$, is

$$\begin{aligned} -s + \lambda(\lambda - s)x - V'(d)(1 + \lambda x) &= -s + \lambda x(\lambda - s - V') - V' \\ &= -s + \lambda x W' - V' \end{aligned}$$

If $f = 0$, so the derivative is negative, then W and V are linear and therefore the derivative is decreasing in d since x is decreasing. Therefore if $f = 0$ is optimal for some $\tilde{d}$ then it is also optimal for all d in the range $\hat{d} > d > \tilde{d}$. Therefore there will never be any good advice starting from $\tilde{d}$: duration will always be such that either $f = 0$ or $a = 1$. But then $V(\tilde{d}) = (1 - \tilde{d})s$, and since $V(d) \geq (1 - d)s$ for all d, it cannot also be that $V(d)$ is concave and $V(d) > (1 - d)s$ for some d.

Finally, $f(d) = 1$ for $d > \hat{d}$:

Claim (A5). Suppose $d > \hat{d}$. Then $f(d) = 1$.

In this range V and W are linear with $W'(d) > \lambda$ so that it intersects $W(\bar{d}) = TS(\bar{d}) - s$ from below. So since

$$V(d) = (1 - f(d))s + V'(d)(d - f(d))$$

then

$$\begin{aligned} dV/df &= -s - V'(d) \\ &= W'(d) - \lambda > 0 \end{aligned}$$

Therefore $f = 1$.

Finally, for this solution, verify that V is concave:

Claim (A6). The V described by Proposition 2 is concave.

Proof. Suppose $d > \hat{d}$, i.e. $a(d) = 1$ and $f(d) = 1$ Then $W(d) = \lambda + W'(d)(d-1)$ and so W (and therefore V) must be linear.

Suppose $d < \hat{d}$, i.e. $f(d) = 1$ and $a(d) = 0$. Then

$$V(d) = \lambda(\lambda - s)(d^+ - d) + V'(d)(d - 1 - \lambda(d^+ - d))$$

Let $d^+ - d = x$. Note that $x' \leq 0$ if W is convex. So

$$\begin{aligned} V' &= \lambda(\lambda - s)x' + V''(d - 1 - \lambda x) + V'(1 - \lambda x') \\ x'(\lambda V' - \lambda(\lambda - s))/\dot{d} &= V'' \end{aligned}$$

so

$$\begin{aligned} V'' &= -x'\lambda(\lambda - s - V')/\dot{d} \\ &= -x'\lambda W'/\dot{d} \end{aligned} \tag{7}$$

but both x' and $\dot{d}$ are negative, while W' is positive, so $V'' \leq 0$.

The final step to show concavity is to show that V is concave at the point $\hat{d}$. Since $W(d^+(\hat{d})) = \lambda$,

$$\lambda - W(\hat{d}) = 1$$

So the slope of W on for $d > \hat{d}$, since W linear, is $(\lambda - W(\hat{d}))/(1 - \hat{d}) = 1/(1 - \hat{d})$. Taking the limit from the left of $\hat{d}$:

$$\begin{aligned} W(\hat{d}) &= \lambda + W'(\hat{d})(\hat{d} - 1 - \lambda(1 - \hat{d})) \\ &= \lambda - W'(\hat{d})(1 - \hat{d})(\lambda + 1) \end{aligned}$$

so

$$W'(\hat{d}) = \frac{\lambda - W(\hat{d})}{(1 - \hat{d})(\lambda + 1)} = \frac{1}{(1 - \hat{d})(\lambda + 1)} < 1/(1 - \hat{d})$$

so W is convex, and therefore V is concave. □

Claim (A7). V is *strictly* concave on $0 < d < \hat{d}$

Proof. The previous claim showed strict concavity at $\hat{d}$. Consider $x = d^+ - d$. If $d^+ > \hat{d}$ and $d < \hat{d}$, the slope of W is strictly higher at d^+ than $\hat{d}$ so $x' < 0$. This implies, by (7), that $V'' < 0$ for all such d. Let this range of d be denoted $(\hat{d}_2, \hat{d})$. Now if, for some d, $d^+(d) > \hat{d}_2$, by the same argument V is strictly concave at d. Let all such d be denoted by the range $(\hat{d}_3, \hat{d}_2)$. Since $W(d)$ is bounded by $\lambda - s$ and $W(\hat{d}_i) - W(\hat{d}_{i+1}) = 1$, repeated iteration must cover all of $d < \hat{d}$ in finitely many intervals. Therefore $V(d)$ is strictly concave on $0 < d < \hat{d}$. □

Claim (A8). The solution described is a viscosity solution to the dynamic program, and therefore solves the optimization problem.

Proof. The final step of the proof is to show that the solution constitutes a viscosity solution to the dynamic program at $\hat{d}$. It is vacuously a supersolution. Take some smooth $\phi(d)$ where $\phi(d) - V(d)$ is at a local minimum at $\hat{d}$ and compute

$$(1 - a(\hat{d}))\lambda(\lambda - s)(d^+ - \hat{d}) + \phi'(d)(\hat{d} - 1 - (1 - a(\hat{d}))\lambda(d^+ - \hat{d}))) > V(\hat{d})$$

and therefore V is a viscosity subsolution and therefore a viscosity solution. The constructed V is therefore the unique viscosity solution by Crandall and Lions [1983]. □

Claim (A9). The solution to the dynamic program solves the problem for arbitrary history dependent policies for f and a.

Proof. Suppose instead contracts are indexed by promised utility to the influencer, W. This transforms the problem into the usual utility possibility set as in Abreu et al. [1990]. Since $V(d)$ is concave and $V(0)$ is exactly the total surplus, W is an increasing function of d, facilitating the transformation. Let W^+ be the promise after good advice is received, and $\dot{W}$ be the rate of change after no good advice is received. The value function for the follower as a function of the promise W to the influencer is

$$V_W(W) = max_{a,p}(1 - f)s + f\lambda(1 - a)(1 + V_W(W^+) - V_W(W)) + V_W'\dot{W}$$

subject to

$$W = f\lambda((1 - a)(W^+ - W) + a) + \dot{W}$$

Since $W(d)$ is monotone, applying the change of variables $W = W(d)$ recovers an identical solution. That is, then

$$\begin{aligned} f(W) &= 1 \text{ if } W > 0 \\ a(W) &= \begin{cases} 0 & \text{if } W \leq \hat{W} \\ 1 & \text{if } W > \hat{W} \end{cases} \end{aligned}$$

where $\hat{W} = \lambda - 1$. □

Proof of Lemma 4

Proof. Suppose $W_x(d) = xW(d)$. We show that the follower solves the same problem and therefore chooses the same policies, and therefore generates $V_x(d) = V(d)$. Replacing $W_x(d) = xW(d)$ in the follower's problem, x only enters the follower's problem through the IC constraint and the definition of total surplus. For surplus, for general x it must be that

$$W_x(d)/x + V_x(d) = d(\lambda - s) + s$$

and therefore the constraint is identical if $W_x(d) = xW(d)$. For incentive compatibility, $W_x(d^+) - W_x(d) \geq x$ is the same as $W(d^+) - W(d) \geq 1$. So if $W_x(d)$ is as stated, the follower's problem is identical and therefore $V(d)$ is the same and the optimal choices for the follower are identical. Substituting the same decision rules into the recursion of (6), x cancels which verifies that the decision rules generate $W_x(d) = xW(d)$ □

Proof of Proposition 5

Proof. If d_0 is unconstrained, then this follows directly since $W_x(d_0)$ falls and $V_x(d_0)$ remains unchanged. Ifd_0 is constrained, then it must increase in response to the tax. $V_x(d_0)$ cannot increase in d_0 (since, if it rose, it would have been better to choose higher d_0 in the absence of the tax as the function $V(d)$ is unchanged) and $W_x(d_0(x))$ remains at $\bar{W}$. □

Proof of Lemma 6

Proof. When $a = 0$, the claim is automatically satisfied since $a = a_m = 0$ is the only feasible amount of disclosed ads. Suppose that $a > 1$. The follower must reward good advice with increased payoff of at least u, since the return to excess ads is u. Therefore the IC constraint requires that, when $a > 0$, $W' \geq W + u$ The modified problem for the follower is

$$V_W(W) = max_{a,a_m,p}(1-f)s + f\lambda(1-a)(1 + V(W') - V(W)) + V'\dot{W}$$

subject to

$$W = f\lambda((1-a)(W' - W) + a_m m + (a - a_m)u) + \dot{W}$$

For any a_m, a combination where the statement does not hold, there is a combination with lower $\hat{a} < a$ and $\hat{a}_m$ with either $\hat{a}_m = \hat{a}$ (if $m > u$) or

$\hat{a}_m = 0$ (if $m > u$) and $a_m m + (a - a_m)u = \hat{a}_m m + (\hat{a} - \hat{a}_m)u$. This remains IC for the same W' and satisfies the promise keeping constraint for the same $\dot{W}$ as in the original contract, but increases $V_W(W)$. Therefore the statement must hold at an optimal contract. □

Proof of Lemma 8

Proof. Take two u, u' with $u' < u$. Suppose the policy for u is followed when the return to a_u is u'. Then by construction promise keeping holds and gives $W_u(d) = W_{u'}(d)$. Therefore the policy is incentive compatible (since choosing $\hat{a} > 0$ when $a = 0$ has a lower return and the same foregone value) at u'. Therefore following the u policy gives the same $V(d)$ in either case. But since the IC constraint is now slack for every $d < \hat{d}$, and strict concavity implies that the IC constraint binds at an optimum, there is a strict gain by moving to the optimal policy for all $d \leq \hat{d}$, and relaxes the constraint $V(d) \geq s$. □

Proof of Proposition 9

Proof. Since $m > u$, when $a = 1$ the ads would be undisclosed if the policy were always applied, and the solution is described by Proposition 2. Fix this policy for each d. Then by the same argument as in Lemma 4, the payoff to the influencer is scaled up by $1/m$ and the payoff to the follower is unchanged. Therefore the policy remains incentive compatible, since $W(d^+) - W(d) = 1/m > 1$. Therefore the same policy is feasible and therefore clearly the opt in policy leaves the consumer at least as well off as the regulation of all undisclosed tweets.

For that solution, it therefore must be that $W(1) - W(\hat{d}) = 1/m > 1$. Since the return to raising $\hat{d}$ is strictly positive given V in the original contract according to Claim A3, this slack constraint improves the follower's payoff near $\hat{d}$, and therefore $V(d)$ can be made strictly larger for all $d \leq \hat{d}$. Again this relaxes the $V(d) \geq s$ constraint. □