Free Ad(vice)

Matthew Mitchell*

June 22, 2018

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 like 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 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 where 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 at the same time lowered the cost of providing advice through blogs and social media. Advice is often offered to potential consumers without any direct payment from the consumer to 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 include, seamlessly, paid endorsements. Twitter users provide sponsored recommendations to followers; in the U.S., FTC regulations suggest that the sponsorship should be disclosed, but it rarely is.[3] This differentiates this form of advertising from typical media advertisements or paid endorsements, which are largely transparent and explicitly separated from content. Websites like cnn.com include sponsored content alongside links news content in a way that makes the sponsored content seem like part of the news. 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 reward for providing good advice is to maintain followers for the influencer.

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 toward paying advertisers. On the other hand, it needs to maintain good

[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

advice on average, or following will not be valuable to followers. The goal of the model is to provide a simple implementation of the role of this tension in the dynamic relationship.

The model has positive implications about who make good influencers and how the relationship between influencers and followers evolves. The model is applied to policy questions, especially the proposed FTC disclosure regulations, to understand the impact on these dynamic relationships. Policy guidance in the model is different from a standard model where 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 providing current good advice. This second effect may make some regulations less effective, and suggest alternative regulations that are superior.

In the model, the relationship alternates between periods where the agent monetizes the opportunity to advise, and periods where unbiased advice is given, in a "reap and sow" cycle. A sufficiently long period without good advice causes the relationship to breakup permanently. The relationship can be summarized by the duration that the relationship is expected to last into the future. 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 no longer can 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 is extreme in the model, 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 through a longer time of following. The optimal contract economizes on that cost, while still incentivizing good advice when possible.

A key intuition for the model can be understood by considering 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. Therefore 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, that the reward for good behavior is future opportunity to operate the technology, and not direct monetary transfer between the parties.

This basic force is at the heart of the key results on disclosure rules like 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 which may be passed on to followers. The beneficial effect of disclosure is only 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 enough to offset the costly taxation effect that has no beneficial effect on advice. Therefore 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 where advice is provided in exchange for money; here the reward for the seller of providing information is the future ads themselves, which might also be impacted by the disclosure rules. This distinction 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 higher returns for followers than blanket mandatory disclosure. 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 the their advertising technology. Such a policy can improve welfare of consumers even when mandatory disclosure cannot, because 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 paper is organized as follows. The model, which is very sparse in its most basic form, is introduced in section 2. Much of the paper focuses on the case where the follower can commit to a contract. This is because commitment turns out to be irrelevant for many of the qualitative features of the contract, and therefore the model can be largely understood even with commitment, using simple tools of optimization. One goal of the paper is to provide a tractable model of this sort of relationship in the spirit of other papers in the literature described below, in order to be amenable to policy questions. The commitment contract is developed in section Section 3. Section 4 then uses the model to study the policy issues of mandatory disclosure and market power. Section 5 discusses the fully relational contract with no commitment on either side, in order to highlight the qualitative similarity to the benchmark discussion. Section 5 also considers extensions to allow for the influencer to make revenue from followers in other ways besides at the expense of good advice, and to allow for the possibility that ads lead to demonstrably bad advice, and provides a microfoundation of the model of disclosure used in the main analysis.

1.1 Literature

1.1.1 Dynamic Contracts without Money

The model of the contract is as a dynamic contract without money, and is therefore broadly similar to papers in that literature, and specifically most similar to Li et al. [2015] and Bird and Frug [2015] 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 like Lipnowski and Ramos [2015]. Rather than payoffs being unknown as in Lipnowski and Ramos [2015], the feasible set (that is, whether or not good advice can be generated) is private information of the influencer. That element is the one that puts the model most in common with Li et al. [2015] and Bird and Frug [2015]. 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. In the finance literature, the paper of DeMarzo and Fishman [2007] has 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 in terms of future time 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

Several papers have studied disclosure rules in markets similar to the ones studied here. 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 of Inderst and Ottaviani [2012].

The idea that advertisement and advice may be at odds on the internet dates back at least to the formative literature on search such as Brin and Page [1998], who stated: "[W]e expect that advertising funded search engines will be inherently biased towards the advertisers and away from the needs of the consumers." This paper contributes to thinking formally about the role of dynamic relationships in this bias, but in the context of advisers in places different from search engines per se. The most closely related relational contracting papers have been used to study employment relationships. A goal of this paper is to adopt that approach to understand industrial organization and regulation for situations where monetary transfers between the parties being modeled does not drive incentives.

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 to be sure 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] models 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, whereas this paper focuses on the dynamic incentives faced by the adviser. This paper provides further understanding of the problem faced by an adviser. 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. In Yang and Ghose [2010] it is shown that paid advertisements are associated with higher click-through on organic results. Edelman and Lai [2014] directly studies the role of Google's display of its own property (flight results) on users' behavior. They show that Google's flight results generate both clicks on the Google property and on paid ads, suggesting that indeed Google does have at least two channels by which it is incentivized to bias listings toward its own properties. 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, similar to the reduced form static model studied here, 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 [2017] studies the policy issues related to taxation of advertising revenues in these markets. The taxation discussion is related to the reduced form modelling of disclosure policy in this paper. This paper therefore builds on that literature by expanding the focus to dynamic relationships and advertisements that are potentially hidden within content.

1.1.3 Repeated Advice

Sometimes advice is 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 only be 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 includes dynamics and cycles of reputation such as 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 which is normalized to 1. The follower's choice of whether or not to follow the influencer at any time t is denoted $f_t \in [0, 1]$, were $f_t = 1$ indicates following and $f_t = 0$ is not following.[4] Following is costly to the follower, as it requires foregoing 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. Assume that f_t is publicly observed, either through a direct measure of following, or an indirect measure via behavior like clicks that generate the 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 $a_t \in [0, 1]$. The influencer gets flow payoff λa_t from choosing a_t. Good advice arrives to the follower at Poisson arrival rate $\lambda(a_t) = (1 - a_t)\lambda$.[5] In the basic model, the follower gets a verifiable benefit of 1 from every piece of good advice it receives, so there is no efficiency rationale for good advice over monetization through ads: the total payoff to the two parties is λ per

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

[5] 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. 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.

unit of time regardless of a_t. The nature of the relationship in the basic model is driven entirely by sharing these payoffs.[6] 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. So that good advice is ever given let $\lambda > 1$. So that following is efficient, assume $s < 1$.

The choice of a_t is private information of the influencer. Although the follower cannot explicitly observe and punish the influencer taking money, there is implicit punishment associated with the fact that the influencer 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 only play a role 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 3.5 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 5.2, the advice technology is modified so that there is not a pure tradeoff between ad revenue and good advice, but rather some good advice might also be monetizable. In section 5.3 the follower may have a choice of how much attention to pay. In section 5.4 there is not only good advice but also bad advice that is more likely to come the more intensively the ad technology is used, so that the model has three outcomes (good advice, bad advice, nothing) as opposed to two in the benchmark model. None of these changes alter the important conclusions from the benchmark model.

For comparison, if a_t were observable, the influencer could choose a sequence of a_t so that $f_t = 1$ for all t; the Pareto frontier would just be the set of all payoffs for the follower (V) and influencer (W) such that $V + W = \lambda$. To see this, if the influencer commits to follow as long as $a_t = 1$, and never after, it is (weakly) optimal for the influencer to always choose $a = 1$ and the follower gets exactly λ (while the follower gets nothing). If the follower

[6]It is shown in section 3.5 that the results are the same for an ad technology whose payoff is $\lambda x a$.

commits to follow regardless of a, the influencer gets λ and the follower gets nothing. Mixtures between these two strategies therefore trace out the entire total possible value of λ.[7] Departures from the full information Pareto frontier are purely due to information asymmetry in the choice of a.

3 The Dynamic Relationship

For the benchmark model, suppose the follower can choose at the outset an entire public-history dependent path for f_t.[8] In particular, $f_t(h_t)$ 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 where, 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 identical to one written in terms of promised utilities a la Abreu et al. [1990], and is therefore without loss of generality. For now, one can consider this class of contract (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 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 time derivatives are denoted with a dot over the variable. Using this recursive construction of d allows for writing an optimal contract

[7]Follower commitment merely imposes that $V \geq s$ but otherwise leaves the allocation unchanged.

[8]Section 5 shows that the qualitative characteristics of the allocation are unchanged whether or not the follower has commitment power.

recursively. Indexing the contract by d is also useful because of its close relationship to total surplus. For any d , the total surplus is

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

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

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 treating the follower as the principal, i.e. computing follower-optimal allocations for a given d, and therefore making a a choice variable of the follower subject to incentive compatibility, in order to trace out the frontier. The contract will also be characterized by an initial duration that determines the surplus division, discussed below.

The recursive problem, according to the principal of optimality, for any d is[9]

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

subject to incentive compatibility of a (to be described below) and the delivery of d according to the promise keeping constraint (1).[10] Denote the solution to this problem by $a(d)$ and $f(d)$. 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

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

[9]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 $\dot{d}$.

[10]There are also domain restrictions on d, a, and f (that they lie between zero and one). To keep the notation simple these are not explicitly included, but the discussion below always implicitly takes them into account, explicitly when they bind.

Incentive compatibility for a is thus

$$W(d^+) - W(d) \begin{cases} \geq 1 & \text{if } a(d) = 0 \\ \leq 1 & \text{if } a(d) = 1 \\ = 1 & \text{if } 0 < a(d) < 1 \end{cases} \tag{3}$$

3.2 The Pareto Frontier

The solution to the problem relies on concavity of V, which is essential to insuring that the IC constraint (3) binds when $a(d) < 1$. While the proof of the main characterization describe concavity and the binding IC constraint more formally, the model is simple enough that some intuition can be obtained without the complete argument.[11] For concavity, the following simple argument shows intuitively why one might expect that V is concave. Take 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[12]

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

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:

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

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

[11] The formal proof of this, and concavity itself, is contained in the appendix, as part of the proof to the characterization proposition 1.

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

Concavity implies that the IC constraint must bind, i.e. when $a(d) < 1$, $W(d^+) - W(d) = 1$. Intuitively, suppose d^+ is more than necessary for $a < 1$. To maintain the promise of d, that means $\dot{d}$ 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 for the follower when V is concave.

The binding IC constraint in turn helps to understand the incentives for the follower in choosing a. When the IC constraint binds, the difference between $V(d^+)$ and $V(d)$ can be rewritten using (2):

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

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

$$V(d) = max_{a,f}(1 - f)s + f(1 - a)\lambda(\lambda - s)(d^+ - d) + V'(d)\dot{d} \qquad (4)$$

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

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

The Bellman equation in (4) is linear in a, suggesting that corners are optimal. To understand the solution, 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:

$$1 + TS(d^+) - TS(d) - \big(W(d^+) - W(d)\big)$$

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 and 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 full solution, if any following ever occurs, is characterized in the following.

Proposition 1. *Suppose that* $f(d) > 0$ *for some* d*. Then*

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

where $W(1) - W(\hat{d}) = 1$*. Moreover,* W *is increasing and convex and* V *is concave (strictly for* $d < \hat{d}$*, linear for* $d > \hat{d}$*).*

The follower incentivizes good advice fully whenever feasible, and stops following only when $d = 0$; i.e. severance is permanent. The value function is depicted below.

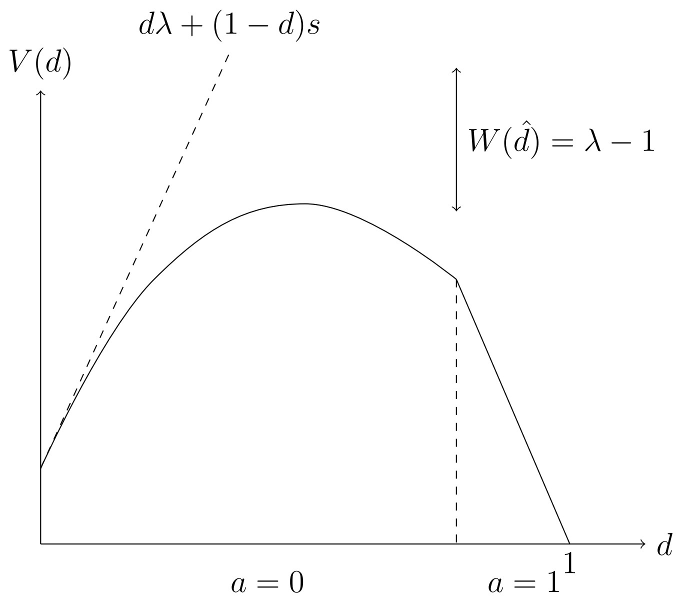

In addition to studying the case without commitment for the follower, Section 5 introduces attention on the follower side, as well as good advice sometimes arriving when $a = 1$, in order to further rationalize following in the $a = 1$ region. Although the bang-bang nature of the solution is the result of linearity, the basic idea is intuitive. Take any structure where breakup is costly, so total surplus is increasing in d, but 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 occur for low values of d.

Since the value function is linear for $d > \hat{d}$, the contract could randomize between $\hat{d}$ and 1 any time d falls in this range. However doing so is not essential for optimality, and therefore entrenchment in the sense of Li et al. [2015] does not necessarily occur. The only absorbing state that is required by optimality is $d = 0$. It is immediate that relationships end eventually with probability 1 if no randomization is used, since $d = 1$ can only be achieved if good advice comes with exactly duration $\hat{d}$:

Corollary 2. *For the solution described in 1, for all $\epsilon > 0$ there exists T such that the probability of following T periods from period zero, is less than ϵ, i.e. $E_0(f_T) < \epsilon$*

Finally: for large enough λ it must be the case that the follower does some following, and therefore the solution is not just $f(d) = 0$ for all d.

Lemma 3. *For λ large enough, $f(d) > 0$ for some d.*

3.3 The contract with promised utilities and interpretation as a "chip" mechanism

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 principal as a function of the promise W to the agent is

$$V_W(W) = max_{a,p}(1-f)s + f\lambda(1-a)(1 + V(W^+) - V(W)) + V'\dot{W} \quad (5)$$

subject to

$$W = f\lambda((1-a)(W^+ - W) + a) + \dot{W} \quad (6)$$

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

Proposition 4. *The problem described in (5) has the same solution as in Proposition 1, i.e. if $f(W) > 0$ for some W, 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 $\hat{W} = \lambda - 1$.

The utility possibility frontier can be depicted graphically:

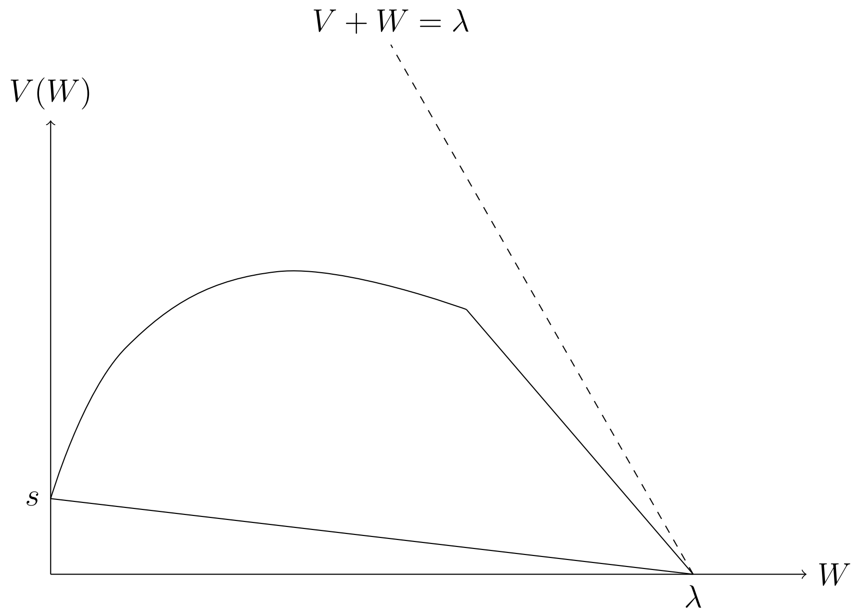

The change of variables facilitates interpreting the contract as being decentralized as a transfer of chips. Let the stock of (divisible) chips held by the influencer be given by C. Whenever good advice is delivered, the influencer is paid one chip. At every point in time where they are followed but good advice is not delivered, the chip stock changes at rate $C - \lambda$, since when $f = 1$, from (6), $\dot{W} = C - \lambda$. This can be interpreted as the chips being paid by the influencer to the follower (in exchange for following) at a constant rate λ, with the stock earning interest at the common interest rate of 1. For instance when $C = \lambda$ the influencer has enough chips to pay to be followed at every future period, using only the interest earned on the stock of chips, and can be followed forever. For any lower value, some good advice will be necessary to keep the following going.

Although the chip mechanism looks like a transfer, actual access to transfers could achieve much more. With transfers, there is no contradiction between achieving points on the full information Pareto frontier with $a = 1$ even with static contracts: if the follower is paid some about x to follow at t, and pays 1 if a piece of good advice received, the static contract generates payouts ranging from $V = \lambda$ and $W = 0$ to $V = s$ and $W = \lambda - s$. In other words, the departure of the outcome from the Pareto frontier is a function of the lack of rich enough transfers between the influencer and follower. The chip mechanism is limited by the fact that the chips can only be used to buy future following behavior, and not consumed directly.

3.4 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 payoff. This is partially for expositional simplicity, and partially to keep the focus in the policy analysis away from the trade-off between the follower and influencer at date zero, for which the model has less to say.[13] The initial condition that maximizes the follower's payoff is

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

It is often relevant to the comparative statics whether or not the constraint in that problem binds. The initial condition will be called unconstrained when the constraint is slack, i.e. $\bar{W} \leq max_d V(d)$. The unconstrained case corresponds to the case of no supply-side response by influencers. The constrained case corresponds to an extreme form of supply response for influencers (locally completely inelastic in W) and allows for analysis of how supply side forces might impact the results. Influencers could be either scarce (when the constraint binds and the initial conditions of the contract must respond in terms of delivered utility to the influencer) or not (when it does not). The market completely breaks down if $V(\bar{W}) < s$, since the arrangement would be worse for followers than just taking the outside option forever, so it is

[13] One could easily allow for other forms of bargaining over the Pareto frontier to determine d_0.

assumed throughout that $\bar{W}$ is small enough that $V(\bar{W}) > s$. This is always satisfied for some $\bar{W}$ close enough to zero.

3.5 Varying the Payoff to Ads

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

Lemma 5. *Suppose the influencer's payoff from the advertising technology is $x\lambda a$ for all d. Then $W_x(d) = xW(d)$ and $V_x(d) = V(d)$*

Section 5.2 introduces the idea of a cost or benefit from the following relationship unrelated to advice. It is shown that such value improves the follower's payoff, and therefore one pro-follower policy is to tax influencer income *and subsidize influencer-follower relationships*, for instance through making the internet faster or less expensive. Another interpretation is that, from the standpoint of generating good advice, a tax is not equivalent to a quota that had real effects on the level of the ad technology. The reason is that the tax impacts both 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 6. *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 and tax revenue plus the payoffs of the follower and influencer follow (2).

This result is relevant in the next section where disclosure rules are studied because such rules may have a taxation-like effect on the relationship, and therefore disclosure can have a negative consequence on welfare.

4 Disclosure Policy

An important policy consideration in these relationships is whether there would be benefit in mandating disclosure of monetary compensation by influencers. Such a policy is actively pursued, for instance, by the Federal Trade Commission in the US. In this section the model is augmented to include a disclosure decision in a reduced form way that is consistent with Inderst and Ottaviani [2012]. A more complete model of disclosure that justifies the reduced form is provided in section 5.3.

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 the technology.

4.1 A Reduced Form Model of Disclosure

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 might 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 be have direct costs: Twitter's character count means that characters used in disclosure are costly. The model of section 5.3 models in more detail the role that attention on the part of the follower might have on total surplus.

[15] 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.

Valuable (but paid) advertisements might be lost if consumers can reduce attention as a result of disclosure.

The idea that disclosure generates costs for both sides might come out of an economic model like the one in 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 underprovision of incentives. The positive role of commissions is frequently overlooked notably in policy debate.

Assume the ad technology has two modes of operation: disclosed and undisclosed. The amount of disclosed ads is $a_m \leq a$, and is observable. Since a_m is observable, it can be treated as a direct choice of the follower. An authority regulates disclosure by imposing a cost on any ads in excess of a_m, so that undisclosed ads return $\lambda u a_u$, where $a_u \geq 0$ is the difference between a and a_m. 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 these tweets so far.[16]

To model the lower return to disclosed ads, let the payoff from disclosed ads be $\lambda m a_m$ with $m \leq 1$. The details of how disclosure impacts follower's perceptions is described in more detail below, but to focus on its role in the dynamic relationship it is enough to take these returns as a parameter in this section.

[16] For instance, the famous Ken Bone tweet for Uber following the US Presidential Town Hall was taken down after the FTC said 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.

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

Lemma 7. *If $m > u$, then $a_m = a$. 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, since the ad technology always makes λua as in section 3.5 with $x = u$. According to Lemma 5, it therefore decreases W and has no impact on V for given d, and cannot improve total welfare once the initial condition is taken into account according to Proposition 6. Such a weak disclosure policy is effectively a burden on monetization that does not benefit followers.

Corollary 8. *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 9. *Suppose $u < m$. Then, for all d, $V(d)$ is decreasing in u.*

Disclosure is good for followers but bad for influencers; the net impact on welfare is ambiguous. Unambiguously, however, a policy near $u = m$ is not welfare improving. A disclosure policy must be sufficiently harsh to offset any "taxation" effect it has on the $a = 1$ part of the policy. There is no reason why 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. This is because, to the extent that disclosure acts like a tax, it can potentially harm influencers without improving the quality of advice provided.

4.2 Alternative policies

The model suggests alternative policies that could be an improvement. Suppose that disclosure rules only applied to influencers below $\hat{d}$. High d influencers are free to make the full ad technology return. Then the follower gets the benefit of the tighter IC constraint without the cost of making the reward to good advice lower. This policy could be implemented on an opt-in basis. Suppose the influencer could announce whether or not disclosure rules would apply to them before the follower chooses f.[17] For $d < \hat{d}$, the follower only follows if the announcement is that disclosure rules apply; for $d > \hat{d}$ no such requirement is imposed. Influencers with low d announce that disclosure rules apply, and they are regulated. Influencers with high d do not. The policy can 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 10. *Suppose that $u > m$. Then the follower is better off under the opt-in rule than with regulation of all undisclosed ads.*

Even in a more general model, where a did not take corners, there would be at least some scope for deregulation of the influencers at the top, i.e. for duration close enough to 1. The ability to be unfettered when an influencer has been very successful increases the incentive to give good advice everywhere else.

5 Extensions

5.1 Limited commitment for the follower

Now suppose that the contract must be supported by the threat of reversion to the static Nash outcome for the game, i.e. the principal getting s and the agent getting 0. In order to impose this, assume that must be the case that V is least s at every point in time. [18]

Definition. A plan is commitment feasible if, for every h_t, $V_{h_t} \geq s$.

The solution with this constraint imposed is qualitatively similar to before. The main difference is that some durations of interaction d are not commitment-feasible. For instance it can never be possible to have $d = 1$,

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

[18] 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 worse than static Nash. Therefore this corresponds to the strongest threat that could possibly sustain any equilibrium.

since then the principal gets payoff zero and would be better off reverting to static Nash. Whether or not a given d is commitment feasible is a cutoff rule:

Lemma 11. *Suppose there is a commitment-feasible plan that has $f = 1$ for duration $\bar{d}$. Then for all $d < \bar{d}$, there exists a commitment feasible plan where $f = 1$ for duration d.*

The result implies that the range of d that is not commitment 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 commitment 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 $V(\bar{d})$.

Define $\hat{d}$ by

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

Proposition 12. *The solution restricted to the set of commitment feasible plans is*

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

for some $\hat{d}$ and $\bar{d}$ where $W(\bar{d}) - W(\hat{d}) = 1$.

The impact of the lack of commitment shifts the utility possibility frontier:

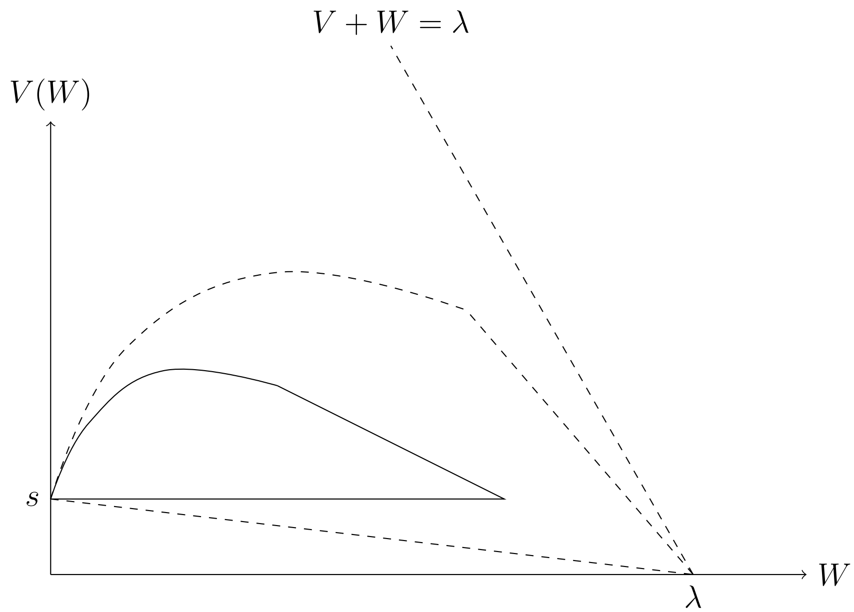

The results are qualitatively the same as with commitment; in particular, for instance, the argument for the taxation result in Proposition 5 is unchanged.

5.2 When advice and income are not in conflict

Influencers often argue that they endorse things that they would recommend even in the absence of sponsorship.[19] 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 and 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.

[19]Fashion blogger Kim France specifically says: “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 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.

Those ideas can be incorporated in the benchmark model in several ways. This section introduces two. An outside source of revenue that does not interfere with good advice can be introduced by letting the revenue from the ad technology when $f = 1$ be $\lambda a + v$. This corresponds to disclosed ads, but could also be a celebrity's valuation of followers for outside reasons. On the other sponsored advice is a mix of valuable and less valuable advice, it may be that good advice arrives at a positive rate $p\lambda$ when $a = 1$, instead of zero, and revenue when $a = 0$ is $p\lambda$ since the influencer can send along the genuine advice that is also paid even when other sponsorship is minimized. The interpretation is that a fraction p of sponsored advice is also good advice. Each extension is discussed formally next.

5.2.1 $v > 0$

In this section let the technology for revenue be $\lambda a + v$, with $v > 0$ (but $p = 0$). The influencer's payoff becomes

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

and the total surplus is

$$W + V = d(\lambda + v - s) + s$$

The value v, although it accrues directly to the influencer, unambiguously benefits the follower:

Proposition 13. *$V(d)$ is increasing in v for $d \in (0, 1)$*

The impact on W is ambiguous: it lowers the information rent on λd but generates vd in extra returns. This implies that a celebrity influencer may be a better adviser to the extent that they have inherent desire for followers, but this does not necessarily improve their rents from giving advice. Good advice and the outside value v are complements, in the sense that it becomes easier to get good advice the higher is v, since, for a given d, higher $V(d)$ can only come about because there is more good advice during the time spent following.[20]

[20]It is natural to ask, given that increasing v improves incentives, whether or not welfare would be raised or lowered if the assumption that cash could not be transferred between

Whether or not taxation remains neutral depends on whether or not the taxation applies to the value v or not. If the taxation applies to all benefits for the influencer, including v, then the argument remains unchanged. If v is exempt from the tax, however, then the tax can have real effects, since, in the incentive constraint, the incentive to run ads when $a = 0$ is the difference between the payoff from $a = 0$ and $a = 1$ (which is reduced in proportion to the tax), while the cost of running ads is the difference in future value at d^+compared to at d (of which only a fraction of the future value is subject to the tax). This points to another way in which policy might be best positioned. In order to inventivize good advice, it is necessary to not tax forms of revenue that might serve as an indirect reward for providing good advice. This might, in many cases, be nearly impossible; it would require delineating benefits for the influencer in terms of which ones *could* be biased, compared to the ones that could not be biased, so that it would tax the temptation to run ads differently from the reward from making the relationship last longer. When such a distinction is impossible, and a tax simply impacts revenues equally, the results from $v = 0$ continue to hold.

5.2.2 $p > 0$

Alternative, suppose there is a technology that allows *paid good advice*. Let a fraction $p > 0$ of possible ads be both paid and good advice. The influencer's value function is

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

and the follower's value function is

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

follower and influencer were loosened to allow a subscription fee, i.e. a payment v from the follower to the influencer when following, so that the incentive effect is the one studied in this section, but at a cost of v to the follower, rather than being a gain in total surplus. Since such a subscription fee improves incentives, it might be the case that it could make the follower better off. However this logic that increasing subscription fees improving outcomes is not universal: the full information social welfare maximum $W + V = \lambda$ can be achieved by a constant *negative* subscription fee per instant when $f = 1$, between s and λ, together with $a = 1$. This suggests that the possibility of subscription fees from follower to influencer are not necessarily an improvement to the relationship.

In this case the follower cannot distinguish if good advice is coming from the fact that ads aren not being run, or because the influencer has been lucky to provide paid good advice. It is now no longer certain that $a = 0$ is chosen whenever feasible; encouraging $a = 0$ entails rewarding the influencer for good advice that they would be willing to give even if there was no dynamic reward. In the limit where p is near 1 there is no reason for the follower to give surplus for all arrivals when very few require a reward. This differentiates this example from the known value of following $v > 0$; when $p > 0$ it might be that paid good advice creates a disincentive for followers to encourage unbiased advice.[21]

The next section provides some microfoundation for the earlier assumption that disclosure lowers the value of ads by considering the possibility that paid advice and good advice may not always be at odds, and therefore if consumers pay less attention to paid advice (which is sometimes helpful and sometimes not), the avoidance of paid advice leads to a reduction in good advice as well.

5.3 Microfounding Disclosure as a Tax on Influencers: Following but Ignoring

In section 4 disclosure was modelled as a tax on influencers. This section provides a microfoundation for that assumption through a very simple model of *two sided* private information, where disclosure might lead to a change in

[21]If $a = 0$ for every arrival the follower receives

$$1 + V(d^+) - V(d)$$

Since social surplus is now $\lambda(1+p)d + s(1-d)$,

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

and so the benefit for follower can be written as $(\lambda(1+p) - s)(d^+ - d)$. This is positive, as in the benchmark with $p = 0$. Since the follower gets $p\lambda$ arrivals when $a = 1$, choosing $a = 0$ is only optimal if

$$(\lambda(1+p) - s)(d^+ - d) > \lambda p$$

The left hand side of the expression decreases in d, and therefore, for high enough d, it may be the case that the optimal contract now involves periods with $f = 1$ and $a = 1$ even though $a = 0$ is feasible, and would be chosen if $p = 0$.

private behavior of followers, in particular leading them to not pay attention to paid ads.

There are three key elements. As in the prior section, suppose that a fraction p of paid advice also generates good advice for the follower. Additionally, assume that when $a = 1$, there is heterogeneity in advice in terms of being paid and unpaid. One can interpret this as paid opportunities being limited, even when the influencer would like sponsorship. This is essential to a microfoundation of disclosure since otherwise, when $a = 1$ all ads are paid and there is nothing to disclose. Let the probability that advice is paid at any instant when $a = 1$ be ρ. With probability $1 - \rho$ advice is unpaid, and generates good advice at rate λ, even though $a = 1$. This is therefore another channel by which, as in section 5.2, consumers might get good advice when $a = 1$. This setup keeps disclosure of *paid* from being identical to disclosure of *valuable*. The consumer therefore receives good advice at rate λ when $a = 0$ and at rate $\lambda(\rho p + (1 - \rho))$ when $a = 1$.

Finally, suppose that the follower can choose to follow but not pay attention to advice from the influencer. Following remains observable but attention is not part of the public history. Not paying attention loses any benefit from the advice for the follower but instead generates a benefit $0 < b < 1$, i.e. it is less than the unit benefit from good advice. For the influencer, if attention is not being paid at a particular history, the influencer can only get m rather than 1 from setting $a = 1$.[22]

One interpretation is that, for instance if the advice is a product endorsement, that 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 $b < 1$, 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 not follow paid advice, since it only produces surplus with probability p.

Assume that $p < b/\lambda$ but $\rho p + (1 - \rho) > b/\lambda$, so that 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 paid, it leads to no attention being paid. Therefore in that case the influencer receives ρm when $a = 1$. The follower receives

[22]This assumption requires, as in Lipnowski and Ramos [2015], that the influencer cannot immediately observe their own payoff, since then they could back out the attention variable directly.

$\rho b + \lambda(1-\rho)$, reflecting the fraction of time spent not paying attention, and the rate of good advice received when paying attention. However if the paid advice is not disclosed, the follower receives $\rho p\lambda + \lambda(1-\rho)$. If p is very close to b/λ, disclosure changes the followers payoff by a vanishingly small amount, and the influencer makes ρm instead of ρ, so disclosure is effectively a tax at rate m on the influencer. In this limiting case the regulation in that case acts like a tax at rate m when $a = 1$, as in the reduced form model.

5.4 Bad advice

Until now, advice was either good or neutral; the cost of neutral advice was implicit in s. To make the concept of bad advice more explicit, suppose that in addition to good advice, bad advice can arrive, at an increasing rate in a. Specifically let bad advice come at rate $\alpha + a\lambda_b$. Bad advice generates a payoff $-l$ for the follower.

Upon bad advice the follower updates duration to $\underline{d}$. The incentive constraint becomes, for $a = 1$,

$$\lambda - \lambda(W(d^+) - W(d)) + \lambda_b(W(\underline{d}) - W) \leq 0$$

or

$$1 - (W(d^+) - W(d)) + \frac{\lambda_b}{\lambda}(W(\underline{d}) - W) \leq 0$$

The IC constraint is loosened by choosing $W(\underline{d}) < W$. On the other hand there is a cost to tightening the IC constraint. The objective is

$$V(d) = max_{a,f,d^+}(1-f)s + f(1-a)\lambda(1+V(d^+)-V(d)) + f(\alpha + a\lambda_b)(-l+V(\underline{d})-V(d)) + V'(d)\dot{d}$$

Since V is concave, the follower faces a tradeoff by punishing bad advice by setting $V(\underline{d}) < V(d)$ when $a = 0$. On the one hand, incentives are improved; on the other hand there are losses due to concavity since $V(\underline{d}) < V(d)$ acts like a random fluctuation in duration for a given value of a.

One way to see the beneficial impact of potential bad advice is in the special case of $\alpha = 0$. In that case there is no bad advice when $a = 0$ so $W(\underline{d}) = 0$ whenever bad advice is received when the policy recommends $a = 0$, since on path there should be no such bad advice and therefore improving incentives through punishment is costless. The IC constraint reduces to

$$1 - (W^+ - W) - \frac{\lambda_b}{\lambda}W \leq 0$$

The IC constraint loosens more, relative to the benchmark case with no bad advice, the higher is W, since the threat of bad advice becomes more meaningful. [23]

6 Conclusion

This paper has introduced a simple 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 only regulating 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 not only overall welfare, but also can be beneficial for consumers who receive information.

Many interesting directions could be developed from this starting point. Other policies, like a tax on monetization that funded infrastructure for making relationships less costly (for instance by making the internet faster) could be considered. 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 had the experimentation aspect of an exponential bandit problem. Another interesting dimension would be to include equilibrium between possibly 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.

[23] When $\alpha = 0$ and $\lambda_b = \lambda$, the IC constraint is just $W^+ = 1$. The contract "resets" to $W = 1$ every time good advice arrives; on path, this is achieved by a fixed period of ads followed by a return to the regime where ads are not run, but good advice resets W to 1, leading to a fixed period of ads before $a = 0$ resumes. I thank Aaron Kolb for suggesting this particular example.

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. The Advertising-Financed Business Model in Two-Sided Media Markets. pages 41–90. 2015. doi: 10.1016/B978-0-444-62721-6.00002-0. URL http://linkinghub.elsevier.com/retrieve/pii/B9780444627216000020.

D Bird and A Frug. Dynamic Nonmonetary Incentives. 2015. URL http://sites.northwestern.edu/dbi475/files/2015/10/Download-Job-Market-Paper-PDF

S Brin and L Page. The anatomy of a large scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1/7):107–17, 1998. ISSN 01697552. doi: 10.1.1.109.4049.

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(May 2002), 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`.

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, 2012. ISSN 00028282.

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. 2015. 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.

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. 2017. 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 Proposition 1

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 (A1). W is increasing and convex

Proof. Since $W(d) = TS(d) - V(d)$, and total surplus is linear, convexity is immediate. Since $W(0) = 0$ and $V(d) \geq 0$, it must also be the case that W is increasing □

Claim (A2). If $a < 1$, 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. □

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 $0 < d < \hat{d}$.

$a(d) < 1$ for some d, i.e. $V(d) > (1 - d)s$ for some d. Then

For $a = 0$, 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 $W(1) = \lambda$ 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 1 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 $x' < 0$ 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

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

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

The final step is to show that V is concave at the boundary $\hat{d}$. 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. □

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].

Proof of Lemma 3

Proof. Suppose the principal asks for advice until stopping at rate γ. If good advice is received before stopping the agent gets asked advice for d units of time (starting from that point, discounted to that point), and then no advice is asked for, so $p = 0$ for the d units of time. Then if the agent gives good advice they get

$$W_\gamma = \lambda(d\lambda - W_\gamma) - \gamma W_\gamma$$

Set $d\lambda - W_\gamma = 1$, so $W_\gamma = \frac{\lambda}{1+\gamma}$. Now as λ grows, choose $\gamma(\lambda)$ so that W_γ is constant. That implies that $d\lambda$ is constant in λ. The derivative of $\gamma(\lambda)$ is

$$\frac{d\gamma}{d\lambda} = \frac{1/(1+\gamma)}{\lambda/(1+\gamma)^2} = \frac{(1+\gamma)}{\lambda} = \frac{1}{W_\gamma}$$

The principal's payoff is

$$\begin{aligned}
V_\gamma &= \lambda(1 + (1-d)s - V_\gamma) + \gamma(s - V_\gamma) \\
&= \frac{\lambda + (1-d)\lambda s + \gamma s}{1 + \lambda + \gamma} \\
&= \frac{\lambda}{1+\lambda+\gamma} + \frac{(\lambda+\gamma)s}{1+\lambda+\gamma} - \frac{d\lambda s}{1+\lambda+\gamma} \\
lim_{\lambda\to\infty} V_\gamma &= \frac{1}{1+1/W_\gamma} + \frac{(1+1/W_\gamma)s}{1+1/W_\gamma} \\
&= \frac{W_\gamma}{W_\gamma+1} + s
\end{aligned}$$

So $V_\gamma > s$. □

Proof of Proposition 4

Proof. By construction the influencer always prefers to choose $a = 0$ when the policy calls for it. It remains to be verified that the influencer would always (weakly) rather pay the follower at rate λ to be followed than not pay, and allow the chips to simply earn interest. Suppose the chip stock is C, and the influencer chooses to not be followed for t units of time, at which point the chip stock is $e^t C$. If the influencer then followed the recommended action from that point on, their discounted payoff would be

$$\hat{W}(C) = e^{-t} e^t C = C$$

Therefore the influencer always chooses to pay at rate λ when the chip stock is positive, as not paying does not improve their payoff. □

Proof of Lemma 5

Proof. Suppose $W_x(d) = xW(d)$. We then verify that the principal's problem is identical, and generates $V_x(d)$, as x only enters the principal'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 principal's problem is identical and therefore $V(d)$ is the same. Substituting the same decision rule into the recursion of $W_x(d)$ verifies that the decision rules generate $W_x(d) = xW(d)$ □

Proof of Proposition 6

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 7

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 9

Proof. Take two u, u' with $u' < u$. Suppose the policy for u is followed when the return to undisclosed tweets 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 concavity implies that the IC constraint binds at an optimum, there is a strict gain by moving to the optimal policy for all $d \in (0, 1)$. □

Proof of Proposition 10

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 1. Fix this policy for each d. Then by the same argument as in Lemma 5, 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 $0 < d < 1$. □

Proof of Lemma 11

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 such a commitment-feasible plan. □

Proof of Proposition 12

Proof. Following the steps for the proof in (A1)-(A6) is unchanged on the domain $[0, \bar{d}]$ for d. □

Proof of Proposition 13

Proof. Following the same line of argument as in Lemma 9, following the same policy at $v' > v$ as is optimal at v remains IC and therefore delivers the same V for$v$. But when $v' > v$, for the policy at v the influencer's payoff at v' is $W_v(d) + (v' - v)d$. Therefore W is steeper in d and the policy remains IC, and therefore is feasible at v'. But since the IC constraints don't bind for that policy at v', the optimal policy at v' generates an even higher payoff. □