Free (Ad)vice

Matthew Mitchell*

April 19, 2017

Abstract

Consumers increasingly rely on intermediaries ("influencers") to provide information about products, often because product choice is vast. Examples include blogs, Twitter endorsements, and search engine results. Such advice is typically not paid for directly by the consumer, but instead the benefit to the influencer comes from mixing advice and endorsement, often in a way that is unobservable to the follower. Giving enough good advice is necessary to keep followers, but there is a tension between the best advice and most revenue. This paper models such a dynamic relationship between such an influencer and their follower. The relationship between influencer and follower evolves between periods of less and more ads. Influencers who inherently value attention provide better advice for followers. The model can provide insight into stricter enforcement of policies like the FTCs mandate of disclosure on paid Twitter endorsements. If disclosure makes adds less valuable, it may be that superior policies to tweet-by-tweet disclosure might exist. For instance a opt-in policy that effectively deregulates influencers with good reputations. The model can also be interpreted as a search engine that biases organic search results to maximize profits, potentially at the expense of providing advice that leads to competing services. Market power by the influencer may be good or bad for welfare, despite bias, suggesting that biased search results by a dominant engine is not necessarily a justification for antitrust-type action.

1

1   Introduction

In many markets where product differentiation is huge, consumers rely on intermediaries to provide information about options. An estimate of the number of products on Amazon alone is over 300 million.[1] Finding the right content, whether it is a song or physical product, or in the case of Google a website itself, is limited mostly by search and not the price of the content (often near zero, or at least constant across goods). The internet has both increased the scope of product differentiation,[2] 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 payment from the consumer. 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 recommendation to followers, often paid; in the U.S., FTC regulations suggest that this should be disclosed, but it rarely is.[3] Websites like cnn.com include sponsored content alongside links to provided content in a way that makes them seem like part of the news. Google provides, in addition to the disclosed advertisements, links to its own products within the "organic" results that can be thought of as ads for products it profits from. Facebook chooses trending topics in a way that can steer users to different products or sponsors. Of course Google results or Twitter feeds are not typically entirely ads: the "influencer" mixes advice with various messages from sponsors in order to earn income from the advice. The small size of each piece of advice on the side of the people receiving 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 manner in the tradition of the recent literature on dynamic contracting without money, and especially the model of delegation in organizations in Li et al. [2015]. The model is based on a tension between good

---

[1]This can be computed directly by a "negative search" for a string of gibberish.

[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 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 advice on average, or following will not be valuable to followers. 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, like the proposed FTC regulations and the impact of market power, to understand the impact on these dynamic relationships. The results show the sense in which guidance is different from a world in which advice would be motivated by money rather than future attention that might lead to more opportunities to give advice.

The contract alternates between periods where the agent is able to monetize the opportunity to advise, and periods where good advice is given. In the long run a sufficiently bad period without good advice causes the relationship to breakup permanently. This "reap and sow" cycle that is reminiscent of a model of reputation in repeated games. 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.

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 works because the influencer's payoff is increasing in this duration: the longer the pair will be together, the more the influencer can extract. The follower's relationship is not monotonic. On the one hand, the relationship generates value in aggregate, 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 faces a cost of rewarding the influencer through a longer time of following, and must economize on that cost, while still incentivizing good advice.

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

3

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.

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 Twitter are motivated by the usual notion that disclosure can improve transaction value. This intuition comes from transactions with 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 is the key sense in which a 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, including an opt-in policy for disclosure, that give higher welfare in the model than blanket mandatory disclosure. Influencers in the sow period would be expected to opt-in (or else not be followed) while influencers with good reputations 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 sow period) as unconstrained as possible. Making the payoff high in the "reap" period is essential to making the relationship efficient throughout.

The model can also be used to understand the role of market power in these relationships. In the case of search engines, paid advertisements by third parties are listed as such, but promotion of the search engine's own pages is difficult to distinguish from the optimal advice given the search term. Google has been accused of biasing search results in favor of its own

products.[4]The treatment of search engines relative to this conduct is a central policy question, in particular in the recent action against Google in Europe about the placement of Google's own pages relative to its best algorithmic advice for a given search. Google's acquisition strategy, including YouTube, can be seen as not independent of the goal of favoring some results over others.

The model is used to ask to what extent a dominant search provider makes advice better or worse. Strengthening an influencer's value can both increase a notion of "market power" for the influencer and also increase the incentives for the influencer to keep the relationship alive by providing enough good advice. In this sense the relationship contains a form of capital that leads to the possibility of natural monopoly, relevant when assessing the appropriate competition policy for a large influencer like Google. In section 4.2, market power, defined as more value from the influencer relative to the outside option, can increase social welfare. The reason is that greater "inside value" makes the relationship easier to maintain, and therefore can benefit both sides. There is a sense that this sort of favor-trading model has a sort of natural monopoly, in the sense that inside value unambiguously makes the relationship more efficient. On the other hand market power skews value to the influencer, so it is unclear whether or not the follower benefits. In section 4.2, it is shown that the efficiency effect never benefits the follower on net, and therefore market power causes consumer harm, even when it increases welfare.

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 is a bit of a distraction. 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. The section 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.

---

[4]Blogs have no restrictions; there is no way to tell what things are being recommended because of compensation.

1.1 Literature

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. 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.

1.1.1 Reputation

Sometimes advice is modeled as cheap talk as in Crawford and Sobel [1982]. The model here differs in that the bias of the sender determines the sender's payoff, but the sender's action can only be imperfectly monitored. 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 has dynamic reputation, and includes dynamics and cycles of reputation such as in Liu [2011], 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.

1.1.2 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 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. The model here is cast in continuous time which allows characterization and comparative statics, as well as policy analysis, that are the motivation here.

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]. Halac and Prat [2014]study an employment model where the employer gets a periodic signal and has private information about whether the monitoring technology is operating. The employee responds by shirking depending on whether or not they believe the monitoring technology to be operating. In this paper the private information is entirely on one side (the influencer), but the model shares the Poisson structure with periodic improvements in the state and continuous degradation of the state when no jump occurs.

1.1.3 Disclosure and Internet Policy

Several papers have studied disclosure rules in markets similar to the ones studied here. The closest is Inderst and Ottaviani [2012] who 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 be bad 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.

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. For search engines, papers like Yang and Ghose [2010] and Edelman and Lai [2014] studies 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 gen-

erate 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. 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. In that sense this paper provides further understanding of the problem faced by an adviser. 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.

2 The Model

A follower can choose to follow an influencer in continuous time. The future is discounted by a common discount rate r which is normalized to 1.[5] The choice of following at any time is denoted $f \in [0,1]$, were $f = 1$ indicates following and $f = 0$ is not following.[6] Following is costly to the follower, as requires foregoing an outside opportunity with flow payoff s. Therefore this outside payoff in any instant is $s(1-f)$. 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 use of ad technology be denoted $a \in [0,1]$. The influencer gets flow payoff λa from choosing a. Good advice arrives to the follower at Poisson arrival rate $\lambda(a) = (1-a)\lambda$. 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 the basic model, the follower gets a 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 unit of time regardless of a. The nature of the relationship in the basic model is driven entirely by sharing these payoffs. For good advice to ever be given it is necessary that $\lambda > 1$. So that following is efficient, assume $s < 1$.

---

[5]This is a normalization since varying the technological parameter λ by a constant factor has the same impact as adjusting r in the opposite direction by the same factor.

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

The choice of a 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 $\lambda(a)$ is the tension between good advice and monetization that generates the potential for inefficiency. For simplicity, and so that the role of the private information on a is the focus of the model, it is assumed that the arrival of good advice is known by both sides of the transaction.

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 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 were observable, the influencer could choose a constant a so that $f = 1$ forever; the Pareto frontier would just be the set of all payoffs for the follower (V) and influencer (W) such that $V + W = \lambda$. Follower commitment merely imposes that $V \geq s$ but otherwise leaves the allocation unchanged. Departures from the full information Pareto frontier are purely due to information asymmetry in the choice of a. In the next section the follower can commit to fully history dependent time path for f, but a subsequent section shows that the qualitative characteristics of the allocation are unchanged whether or not the follower has commitment power.

3   The Dynamic Relationship

The follower can choose at the outset an entire public-history dependent path for f_t. In particular, f_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

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

where the expectation operator is taken over future histories that could occur following h_t. 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]. The construction of the contract in d leads naturally to the construction of the contract without commitment on either side, which is qualitatively similar and described below. For now, one can consider this class of contract (those summarized by d for any history) to be an assumed restriction on the contracting environment, which later will be shown to be without loss. When unambiguous, the duration after a history will just be written as d_t or simply d. Suppose that the influencer needs to receive at least $\bar{W}$ to invest in setting up the advice technology.

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

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

since the total surplus is s at any time when advice is not sought, and λ per unit of time that it is. This relationship facilitates simplification of the follower's Bellman equation whenever W is determined by a binding incentive compatibility constraint below. The relationship in (2) does not continue to hold when the ad technology has a different rate of return from the advice, however 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 an optimal 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. The recursive problem, according to the principal of optimality, is

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

subject to incentive compatibility of a (to be described below) and the delivery of d according to the promise keeping constraint (1).[7] 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

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

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 ensures that the IC constraint (3) binds when $a(d) < 1$. The following simple argument shows intuitively why one might expect that V is indeed concave. Take some d with follower's value $V(d)$. For $x < d$, a feasible strategy for the follower, which delivers x units of following time, is to wait (with $f = 0$) a fixed interval of time (in discounted terms, $\frac{d-x}{d}$ units of time) and then follow the plan that delivered $V(d)$. The discounted amount of following time is[8]

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

---

[7]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. The derivative $V'(d)$ can always be interpreted as the appropriate left or right hand derivative given the sign of $\dot{d}$.

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

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

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

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 $V(x)$ must be at least as high as this feasible strategy:

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

Although this is not a full proof of concavity, it hints at the sense in which "waiting" strategies can accomplish convex combinations of payoffs.

Concavity implies that the IC constraint must bind, i.e. when $a(d) < 1$, $W(d') - W(d) = 1$.[9] 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); such a randomization is not beneficial for the follower when V is concave.

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

---

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

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. Denote the total surplus by $TS(d)$. The follower gets, from motivating $a = 0$,

$$1 + TS(d') - TS(d) - (W(d') - W(d))$$

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. [10]

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

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

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.

---

[10]Formal proofs are in the appendix.

Since the value function is linear for $d > \hat{d}$, the contract could randomize between $\hat{d}$ and 1 any time d fell in this range. However doing so is not essential for optimality, and therefore entrenchment in the sense of Li et al. [2015] does not 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 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]. The value function for the principal as a function of the promise W to the agent is

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

14

subject to

$$
\begin{aligned}
W &= f\lambda((1-a)+a)+\dot{W} \\
&= f\lambda+\dot{W}
\end{aligned}
$$

Since $W(d)$ is monotone, applying the change of variables $W = W(d)$ recovers the same three equations for V, W, and d as were determined above, and therefore the solution to the problem is as described in the problem in terms of d. The contract as described above is essentially a monotonic transformation of the contract in promised utilities. 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 be given by C. Whenever good advice is delivered, the number of chips grows by one. At every point in time where good advice is not delivered, the chip stock changes at rate $C - \lambda$. 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. Therefore when $C = \lambda$ the influencer has enough chips to ask for advice at every future period, just from the interest earned on the stock of chips.

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.

3.4 Initial d

The final element of the contract to determine is the initial duration d_0. Recall that the influencer needs to receive at least $\bar{W}$ to invest in setting up the advice technology. Then 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, which is the focus, 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 is useful to understand how supply side forces might impact the results.

3.5 Lower Payoff to Ads

The model up to now kept the total surplus independent of a. It might seem more natural that ads are inefficient (or possibly generate net surplus); 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 4. *Suppose the influencer's payoff from the advertising technology is xa for all d. Then $W_x(d) = xW(d)$ and $V_x(d) = V(d)$*

Proof. Suppose $W_x(d) = xW(d)$. We then verify that the principal's problem is identical, and generates $W_x(d)$, 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 $W_x(d) = xW(d)$ □

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 implication is that the contract is unchanged as a result of the tax. When d_0 is constrained, the tax must be passed on to followers. In that case, since $W(d)$ is increasing in x, d_0 increases in response to the tax. $V(d_0)$ falls (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). In either case, total surplus $W + V$ decreases when the tax is imposed. [11]

---

[11]Total surplus plus tax revenue, however, increases; in this case a lump sum tax on influencers increases welfare.

4 Disclosure and Competition Policy

In this section we apply the model to consider two relevant policy issues: disclosure of ads and the role of market power by influencers.

4.1 Disclosure

An important policy consideration in these relationships is whether there would be benefit in mandating disclosure of monetary compensation by influencers. In order to model a meaningful trade-off, we assume that disclosed and undisclosed ads might have different returns, and in particular that disclosure might lower the return to the ad technology. On the topic specifically of internet endorsement, [12]

> Audiences "have a very visceral reaction to '#ad' or '#spon' or whatever it is, where they don't want to know people are getting paid for stuff even if they are," said Jaclyn Johnson, president of creative services at Small Girls PR, where she connects brands like L'Oréal Paris and Urban Decay cosmetics to influencers who have large social media followings. "A few 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."'

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.

4.1.1 Modeling the FTC policy

Denote the influencer's choice of disclosed ads be a_m and undisclosed ads be a_u. 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

---

[12]http://www.nytimes.com/2016/08/30/business/media/instagram-ads-marketing-kardashian.html?emc=eta1&_r=0

perceptions is interesting but left unmodeled, and the value to the follower of the advice is fixed as one. It is straightforward to also allow the value to the follower to differ as disclosure varies; the important feature is that the net value of the advertising technology between influencer and follower is lower when disclosure is forced upon it. 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.

Meanwhile an authority might impose a cost on undisclosed ads, so that undisclosed ads return $\lambda u a_u$ with $a_u \leq 1 - a_m$. 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.[13]

Without further assumption, the ad technology would effectively be whichever gave a higher return. To make the policy meaningful, assume that disclosed ads cannot be greater than the recommended ad level a, since otherwise the ad level would be known to be different from the right one. Implicitly this assumes that "fake" ads, i.e. disclosure of ads that were not actually paid, is possible. The FTC technology checks to see if undisclosed ads were run, but not whether disclosed ads were genuinely paid for. As a result the payoff to

---

[13]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/

running ad level $\hat{a}$ when the policy calls for a is

$$max\{m, u\}a + max\{\hat{a} - a, 0\}u.$$

If $u \geq m$, then the model is identical to one where the ad technology has rate of return u, as in section 3.5. According to Lemma 4, it therefore decreases W and has no impact on V for given d. Such a weak disclosure policy is effectively a burden on monetization that does not benefit followers for a given d.

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

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 $d \in (0, 1)$. □

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.

4.1.2 Alternative policies

The model suggests several policies that could be an improvement. For instance, suppose that disclosure rules only applied to influencers below $\hat{d}$. High d influencers were 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.[14] 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. 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. To the extent that exposure to regulation is observable, the policy can always be implemented as an opt-in arrangement.

4.2 Market power

One interpretation of the model is a relationship between a powerful influencer like Google. Google can choose to distort advice in order to increase

---

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

value to Google, possibly at a cost to consumers. Although the main antitrust concern is whether or not this conduct is anti-competitive relative to other firms (like linking users on the "organic" side of the search to Google travel rather than TripAdvisor), an alternative concern is that Google's monopoly power impacts its incentives to give good advice. Supposed anti-competitive behavior by Google relies on the ability to distort away from the best product; if Google's product is the best link, the conduct would be hard to describe as anti-competitive.

One way to understand market power is as the relative quality of the influencer relative to other options. Suppose that the value of output to both parties is scaled by γ, so that good advice is worth γ to the follower and the advertising technology generates γa for the influencer. [15] Since s is not scaled, higher γ signifies more market power: the follower gets more value inside than outside the relationship. It is easy to characterize the value function in terms of γ and s, since neither impacts the optimal allocation. Let $\underline{V}(d)$ denote V in the case where $\gamma = 1$ and $s = 0$.

Lemma 6. $V(d) = \gamma\underline{V}(d) + s(1 - d)$

Proof. Suppose the policy remains unchanged for any d. Then the IC constraints and feasibility constraints bind (since they are unchanged) for any other γ and s; all that remains is to show that the V indeed is as described. For s, note that we can compute dV/ds directly. The analog to (4) is

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

In the region where $a = 0$ and $f = 1$ it is direct to compute

$$\begin{aligned} dV/ds &= -\lambda(d' - d) - \dot{d} \\ &= -\lambda(d' - d) - (d - 1 - \lambda(d' - d)) \\ &= 1 - d \end{aligned}$$

and

$$dV/d\gamma = \lambda^2(d' - d) + \underline{V}'(d)\dot{d}$$

[15]The possibility that "value" is measured as rate of arrival λ is considered below; frequent arrival additional incentive benefits for the relationship that are held fixed by scaling value and not frequency.

which verifies the conjecture by replacing $dV/d\gamma$ with $\underline{V}(d)$, the recursion for $s = 0$. Since the other two regions are merely s (when $d = 0$) and the discounted waiting time until $V(\hat{d})$ (for $d > \hat{d}$) the other regions follow directly. $\qquad\square$

As a comparative static, consider an increase in market power x by raising γ and lowering s so that $\gamma\lambda d_0 + (1 - d_0)s$ is constant, i.e. $ds/dx = -\lambda d_0/(1 - d_0)$. By doing the comparison in this way there is no sense in which market power, for the starting duration, increases welfare on its own. It may, however, change both the division of surplus and the initial d_0. The following result shows how changing market power impacts the surplus of both parties

Proposition 7. *(1) For fixed d_0, $dV(d_0)/dx < 0$. (2) Suppose d_0 adjusts to the change in x. If d_0 is unconstrained, $d(W(d_0) + V(d_0))/dx > 0$. If d_0 is constrained, $d(W(d_0) + V(d_0))/dx < 0$.*

Proof. Fix d_0. Since $dV(d_0)/dx = \underline{V}(d_0) + (1 - d_0)(ds/dx) = \underline{V}(d_0) - \lambda d_0$, but $\underline{V}(d_0) < \lambda d_0$, so $dV(d_0)/dx < 0$

For total surplus, start with the unconstrained case. In that case d_0 is determined by

$$\underline{V}'(d_0) = s/\gamma$$

Since s is decreasing in x and γ is increasing in x, d_0 increases, and therefore total surplus rises.

On the other hand, if d_0 is constrained, then higher x reduces $V(d_0)$ and therefore raises $W(d_0)$ for fixed d_0. This loosens the constraint and allow lower d_0, lowering total surplus. Since W does not change, it must be that V is decreasing in x. $\qquad\square$

Consider how a fall in market share impacts the parties in the unconstrained case. A fall in market share increases $V(d_0)$ for fixed d_0, and therefore must increase V after $d_0$is optimized. But such a decrease in market share decreases total surplus, so this increase must be more than compensated by a fall in W. In the unconstrained case, more market power strengthens the relationship, which increases surplus but makes the position of the influencer stronger. In the constrained case, the follower offsets any increase in W by (inefficiently) reducing d_0.

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 forever. In order to impose this, assume that must be the case that V is least s at every point in time. [16] We will call such optimal plans "committment-feasible."

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 $d = 1$, 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 8. *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.*

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 committment-feasible plan. □

The result immediately 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 $V(\bar{d})$.

Following the same arguments as in Proposition 1, with the domain restriction that $d \leq \bar{d}$, generates an analogous solution. Define $\hat{d}$ by

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

---

[16]This is analogous to computing the principal-optimal public perfect Nash equilibrium of the game without commitment.

For $d > \hat{d}$ it has to be that $p(d) = 0$. For $d < \hat{d}$, $p(d) = 1$ except when $d = 0$. $a(d) = 1$ except at $d = 0$.

Note that V concave with $V(0) = V(\bar{d}) = s$ implies that $V'(\bar{d}) < 0$. The impact of the lack of commitment shifts the utility possibility frontier:

Without commitment there is a further reason why market power, as defined in the previous section, might improve welfare: making the outside option higher makes the constraint $V(d) \geq s$ tighter. Suppose that the follower could find, without any search friction, a new influencer, so that $s = V(d_0)$. In the unconstrained case, this implies that the utility possibility set is the single point $(0, s)$: no good advice is ever given. Either commitment or a friction in finding a new influencer is essential in generating good advice. This is another sense in which market power may be beneficial in a relationship like this one.

5.2 When advice and income are not in conflict

Influencers like Kim France argue that they are often paid for things that they would recommend anyway. [17] Google contends that it links to its

---

[17]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

own products on searches not because of revenue, but to enhance the user's experience. 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 (for instance separate, disclosed and verifiable ads like Google's right hand side of the search bar, or an inherent value of having followers). Celebrities may inherently value followers.

Those ideas can be incorporated in the benchmark model in a straight-forward way. This section introduces two. To model an outside source of revenue that does not interfere with good advice, let the revenue from the ad technology be $\lambda a + v$, so that there is a known value of followers even when $a = 0$, and that isn't influenced by a. This corresponds Google's disclosed ads, but could also be a celebrity's valuation of followers for career purposes.

When ads are unobserved, they may be a mix of valuable and less valuable advice. It is natural to model this trade-off with an additional form of revenue: good advice that is also paid. Let the choice of this sort of "paid good advice" advertising opportunities be $\tilde{a} \in [0, 1]$ with revenue$\tilde{\lambda}\tilde{a}$, but good advice comes at a rate *increasing* in $\tilde{a}$. In that case it is easy to verify that $\tilde{a} = 1$ and therefore, from the influencers side, it is just as if $v = \tilde{\lambda}$. The difference is that, in the case of advertising opportunities that are also good advice, the follower cannot distinguish if good advice is coming from the non-conflicted set, or because the influencer has avoided including the conflicted advertisements. The distinction turns out to be relevant.

The effect of $v > 0$ is that influencers' incentives are strengthened by other forms of income. In a sense, the model predicts that influencers that have other reasons to value relationships with followers make good advisers; a celebrities desire for attention helps make them a good adviser on Twitter just as Google's paid search results on the right hand side enhance its incentives to give good advice on the left hand side. The outside value increases good advice for any d. However when that value comes from paid good advice, it can be the case that the optimal policy switches away from incentivizing good advice, and therefore the paid good advice can crowd out good advice that is conflicted by advertising.

---

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/

5.2.1 Known Transaction Value

In this section let the technology for revenue be $\lambda a + v$, with $v > 0$, but no technology $\tilde{a}$ for "paid good advice. The influener's payoff becomes

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

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 9. $dV(d)/dv > 0$ for $d \in (0,1)$

Proof. Following the same line of argument as in Lemma 5, 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. □

The impact on W is ambiguous: it lowers the information rent on λd but generates vd in extra returns. This implies that an influencer like Kim Kardashian is a better adviser to the extent that she has inherent desire for followers, but this does not necessarily improve her 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.

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 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

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.

5.2.2 Paid good advice

Now suppose there is a technology that allows "paid good advice" through $\tilde{a}$, rather than for any follower via v. The influencer's value function is as above (replacing v with $\tilde{\lambda}$, since $\tilde{a} = 1$) but now the follower's value function changes to include the possibility of both types of good advice. The follower's problem becomes

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

Intuitively, it is now no longer certain that $a = 0$ is chosen whenever feasible. The reason is that encouraging $a = 0$ entails rewarding the influencer for good advice that they would be willing to give even if there was no dynamic reward. This is easy to see in the limit where $\tilde{\lambda}$ is large relative to λ, and therefore there is no reason for the follower to give surplus for all arrivals when very few require a reward.

To see this more formally, let $\gamma = \tilde{\lambda}/(\lambda + \tilde{\lambda})$ be the fraction of arrivals that will only be given if there is a reward. Then if $a = 0$ for every arrival the follower receives

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

Since social surplus is now $(\lambda + 2\tilde{\lambda})d + s(1 - d)$,

$$V(d') - V(d) = (\lambda + 2\tilde{\lambda} - s)(d' - d) - (W(d') - W(d))$$
$$= (\lambda + 2\tilde{\lambda} - s)(d' - d) - 1$$

and so the benefit for follower can be written as $(\lambda + 2\tilde{\lambda} - s)(d' - d)$. This is positive, as in the benchmark without $\tilde{\lambda}$, but if $a = 1$ the influencer gets γ for every arrival. Therefore $a = 0$ is only optimal if

$$(\lambda + 2\tilde{\lambda} - s)(d' - d) > \gamma$$

Since nothing changes about the convexity of W, 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 $\tilde{\lambda} = 0$. In other words, the fact that the two types of good advice can not be distinguished can serve to make them substitutes, in the sense that more of the paid good advice can crowd out unpaid good advice. This differentiates "paid good advice" from the known value of following v.

5.3 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 $-b$ for the follower. This section shows that the basic logic of the model without bad advice remains valid.

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

$$\lambda - \lambda(W(d') - W(d)) + \lambda_b(W(\underline{d}) - W) \le 0$$

or

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

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

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

Since V is concave, the follower faces a tradeoff between punishing bad advice when $a = 0$ (but note there is no punishment for $a = 1$) and losses due to concavity of when $V(\underline{d}) < V(d)$, which 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 $a = 0$ and it is immediate that the policy is $W(\underline{d}) = 0$ whenever bad advice is received when the policy recommends $a = 0$, since on path there should be no such bad advice. The IC constraint reduces to

$$1 - (W' - W) - \frac{\lambda_b}{\lambda}W \le 0$$

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. Nothing changes about the structure of the optimal contract compared to before.[18]

6   Conclusions

In a market for advice without prices, dynamic incentives come through future attention and advice. In such a market, excess competition can make good advice scarce. A policy that taxes monetization in the advice process does not change the amount of good advice, but the revenues from such a policy that are used to subsidize relationships independent of advice can help solve inefficiencies arising from limited commitment frictions.

Many interesting directions could be developed from this starting point. 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 character 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.

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.

---

[18]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 news resets W to 1. In this region, d is falling even though since $\lambda_b = \lambda$, the total probability of news is $\alpha + a\lambda + (1-a)\lambda = \alpha + \lambda$. In other words, duration falls not because a lack of news one way or another is evidence of ads, but rather because any division of surplus where $W < \lambda$ requires $d < 1$, and it is efficient to backload any separations. I thank Aaron Kolb for suggesting this particular example.

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.

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.

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.

M Halac and A Prat. Managerial attention and worker engagement. 2014. URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2501507.

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.

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 procedure. First suppose that the solution is such that V is continuous and concave. This implies W is increasing and convex. Then verify that the solution has the form described, then verify that under that solution that V is indeed concave. This implies that the solution solves the Bellman operator and therefore is an optimum.

Binding IC

To verify agent's IC constraint binds: suppose you raise d' beyond where the agent's IC constraint binds. The impact of that can be seen through choosing d'

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

so the impact of d' is

$$V'(d') - V'(d)$$

which is less than zero by concavity of V.

Proof of Lemma 1

For sufficiently high d, $a(d) = 0$ is not feasible, because W is increasing and continuous. Further, if $d = 0$, $f = 1$ is not feasible. Focus on the domain where $a(d) = 0$ and $f(d) = 1$ is feasible.

Step 1: Suppose $f = 1$. Then $a = 0$ is always optimal if feasible.

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

$$V_{a=0}(d) - V_{a=1}(d) = \lambda((\lambda - s) - V'(d))(d' - d)$$
$$= \lambda W'(d)(d' - d) > 0$$

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

Since W is increasing, this implies that there are two regions: one below $\hat{d}$ where $a < 1$ is feasible, and above where it is not.

Step 2: Suppose $a(d) < 1$ for some d, i.e. $V(d) > (1-d)s$ for some d. Then $f(d) = 1$ for $d < \hat{d}$.

For $a = 0$, the derivative of the follower's objective for f, letting $x = d' - d$, is

$$-s + \lambda(\lambda - s)x - V'(d)(1 + \lambda x) \quad = \quad -s + \lambda x(\lambda - s - V') - V'$$
$$= \quad -s + \lambda x W' - V'$$

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.

step 3: Suppose $a(d) = 1$, i.e. $a(d) = 0$ is not feasible, for $d > 0$. 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

$$dV/df \quad = \quad -s - V'(d)$$
$$= \quad W'(d) - \lambda > 0$$

Therefore $f = 1$.

34

Proof of Lemma

Step 1: shape of W (and therefore V) for any f and a.

- Suppose $a(d) = 1$ and $f(d) = 0$ Then $W(d) = \lambda + W'(d)(d-1)$ and so the W (and therefore V) must be linear.

- Suppose $f(d) = 0$. Then $W(d) = W'(d)d$. i.e. once again W and V are linear.

- Suppose $f(d) = 1$ 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

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

so

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

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

Step 2: kink point where $d' = 1$

The critical point is $\hat{d}$ such that $d' = 1$, where a goes from 0 to 1. Since $W(d') = \lambda$,
$$\lambda - W(\hat{d}) = 1$$
So slope on the RHS, since W linear, is $(\lambda - W(\hat{d}))/(1 - \hat{d}) = 1/(1 - \hat{d})$
But taking the limit from the left of $\hat{d}$

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

so

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

so W gets steeper at $\hat{d}$, and therefore W is convex, while V is concave.

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 $V_\gamma > s$. □