Deprioritizing Content

Martin Szydlowski*

March 23, 2023

Abstract

I study platform design in a dynamic setting with inattentive followers. An influencer faces a stream of followers who randomly arrive on a platform and whose value from paying attention changes stochastically. The influencer chooses the quality of information to disclose each instant and followers choose whether to process that information. In the unique equilibrium, the influencer provides low quality information at all times. The optimal platform design which implements high quality deprioritizes the influencer's content periodically. The least-cost monetization scheme which implements high quality awards advertising revenue only to influencers with sufficiently many followers.

JEL Classification: D82
Keywords: Platform design, monetization, dynamic information disclosure, costly attention

1 Introduction

Deceptive content is a prevalent issue on social media platforms.[1] To rein in deceptive content, the European Union is set to pass Digital Services Act,[2] which requires platforms to remove such content and to demonetize its creators. In the US, various platforms have adopted voluntary measures. For example, Facebook, Twitter, and Youtube maintain strike systems, under which accounts are punished by prohibiting them from posting content in a certain window of time.[3] Facebook and Twitter also deprioritize and demonetize accounts which are deemed to post deceptive content. That is, they make it less likely that future content is shown to followers, or they reduce or eliminate the amount of advertising revenue that is shared with a content creator. However, effectively policing deceptive content is a

*University of Minnesota. Email: szydl002@umn.edu.

[1]See e.g. https://www.nytimes.com/2022/08/14/business/media/on-tiktok-election-misinformation.html.

[2]See https://digital-strategy.ec.europa.eu/en/library/2022-strengthened-code-practice-disinformation.

[3]See https://transparency.fb.com/enforcement/taking-action/restricting-accounts/, https://support.google.com/youtube/answer/2802032, and https://help.twitter.com/en/rules-and-policies/enforcement-options.

difficult task, because what is and is not deceptive is highly subjective. As Facebook notes "Misinformation is different from other types of speech [...] because there is no way to articulate a comprehensive list of what is prohibited. [...] The world is changing constantly, and what is true one minute may not be true the next minute."[4]

In this paper, I characterize a model of misinformation based on limited attention. I then characterize *content-agnostic* mechanisms which ensure that platforms are free from misinformation. That is, to curb misinformation, the platform does not need to take a stand on what constitutes deception. I identify two specific mechanisms - time-based prioritization rules, which penalize posting repeat content, and monetization rules in which only accounts with sufficiently many followers receive advertising revenue.

In the model, a sender (e.g. an influencer or a content creator) faces a stream of followers who randomly arrive on a platform. At each moment in time, the sender decides whether to post high quality information, low quality information, or both. Followers on the platform use the sender's information to make a binary decision. For example, they may look to the sender's advice on whether to buy a certain product, whether to try a medical treatment, or whether to vote for a political candidate. To access the sender's information, followers must pay an attention cost, which is higher for high quality information. Intuitively, high quality information may be inherently more complex and therefore more difficult to process. Low quality information is easier to process, but is less likely to lead followers to make the correct decision. In this sense, low quality information represents *misinformation* or *deceptive content* in the model. Followers' value from acquiring information is either high or low, and changes over time according to a time-homogeneous Markov chain. Intuitively, followers may be subject to random distractions, which reduce their value of acquiring information, or their inherent value from making a decision may change over time. When acquiring high quality information is sufficiently costly, low value followers do not acquire this type of information. Instead, they wait until their value of information increases. As a result, a "latent" mass of followers builds up over time, waiting for their value to change. However, this mass of followers creates a temptation for the sender. By revealing low quality information, the sender can induce all these followers to instantly pay attention, which earns her an instantaneous gain. For this reason, an equilibrium where the sender only provides high quality information is not sustainable. Instead, the sender provides low quality at all times, which low value followers acquire instantly. Thus, the model provides a rationale for why low quality content is prevalent.

To eliminate low quality content, the platform can use prioritization or monetization rules. A prioritization rule specifies the likelihood that the sender's information is shown to followers at each time. Deprioritizing content means lowering this likelihood, which reduces the low value follower's option value of waiting. That is, followers who wait may not receive information in the future, which encourages them to acquire information immediately instead. This eliminates the sender's incentive to provide low quality content. The sender-optimal prioritization rule involves showing the sender's content to followers periodically. Specifically, once the sender's content is shown to followers, future content is hidden for a fixed amount of time, which is designed to render the low types indifferent between acquiring information and waiting. Then, the sender's content is shown again and the cycle

[4]See https://transparency.fb.com/policies/community-standards/misinformation/.

continues. Thus, in the optimal scheme, followers can only sporadically acquire information. In applied terms, followers' ability to always see content leads to content quality being low. By restricting that ability, content quality improves.

Alternatively, the platform may use monetization to incentivize high quality content. The sender's content generates revenue from external advertisers, and the platform can adjust the share of revenue that the sender receives. I characterize the least-cost monetization scheme, i.e. the scheme which awards the sender the lowest possible share of revenue, subject to implementing high quality information. This scheme features two tiers: senders with few followers receive less advertising revenue, while senders with more followers receive more. The intuition is as follows. When the sender provides high quality, the mass of low types accumulates over time. By choosing low quality, the sender can capture these low types' attention, but in doing so, she resets the mass of active low types to zero. Under the optimal scheme, this reduces her future advertising revenue. To receive payments, the sender needs to wait again until sufficiently many followers accumulate. The optimal scheme then only provides ad revenue when the sender has sufficiently many followers, to discourage such deviations. Indeed, many platforms use "tiers," where content creators' ad revenue increases in their number of followers or the number of views.

Both the prioritization rules and the monetization schemes in this paper do not depend on the actual content quality. That is, to implement these schemes, a platform does not need to measure content quality, which is an inherently difficult task. Instead, the optimal prioritization scheme is simply a function of time, whereas the optimal monetization scheme depends on the number of current followers.

2 Literature

The model features randomly arriving followers with stochastically changing values. The most related papers are Garrett (2016), Dilme and Li (2019), and Dilmé and Garrett (2022), which study dynamic pricing by a seller who commits to a price path. In these papers, any price which makes it attractive for existing low types to purchase leaves high types with rents, and the latter have an incentive to wait for prices to become low. The optimal price path is periodic, with times of high prices and (possibly random) discounts. Kapon (2022) studies a related model of amnesty programs, in which the optimal policy has similar features. In this paper, by contrast, the high type does not prefer the low type's information, and the sender is not concerned with separating the types. Instead, the main friction is the low types' ability to wait until their value becomes high. Then, the sender's lack of commitment leads her to provide low quality information, to prevent the low type from waiting. The accumulation of followers affects the sender's incentive to provide low quality, whereas in the above papers, the accumulation affects the cost of pooling low and high types.

A substantial literature studies internet platforms. Mitchell (2021) considers a dynamic model between a single influencer and a follower. The influencer decides how much to distort advice, and the equilibrium cycles between biased and unbiased advice. In this paper, the sender faces multiple followers with different values from acquiring information, and the quality of advice is not subject to cycles in equilibrium. Mitchell (2021) characterizes a tax scheme on advertising content, which is broadly similar to the monetization scheme in this paper. However, in his model, monetization fails to provide incentives, whereas it is effective in my model. Mitchell (2021) also considers conflict of interest disclosure rules, which play no role in my model, since the sender's preferences are commonly known.

Other related work includes Galeotti and Goyal (2009) who characterize optimal influence strategies in a social network, Hinnosaar and Hinnosaar (2021) who study collusion between cartels of inluencers to increase views, Cong and Li (2021) who study assortative matching between heterogeneous influencers and sellers of products, and Pei and Mayzlin (2022) who consider influencers' decisions to become affiliated with firms, which may lower the credibility of their recommendations. Fainmesser and Galeotti (2021) consider a static model where influencers decide on product endorsements, which affects follower participation and prices paid by advertisers. Compared to these papers, my model focuses on followers' dynamic choice of attention and the resulting distortions.

Finally, this paper is related to the literature on dynamic disclosure, most notably papers with limited commitment (Orlov et al. (2020), Bizzotto et al. (2021), Kaya (2022)), multiple receivers (Basak and Zhou (2020a), Basak and Zhou (2020b)), and limited attention (Che et al. (2020)). Relative to that literature, the key difference is the arrival of heterogeneous receivers, whose types change stochastically.

3 Model

Arrival. Time is continuous and infinite. Infinitesimal followers (he) arrive on a platform at rate λ. After arrival, each follower receives an exogenous shock at rate ρ which leads him to leave the platform.[5] Upon arrival, each follower has a high information value π_h with probability γ and a low value π_l with probability $1-\gamma$, where $0 < \pi_l < \pi_h$ and $\gamma \in (0,1)$. I refer to followers with value π_l as *low types*, and followers with value π_h as *high types*. Followers' value follows a time-invariant continuous-time Markov chain. Low types become high types at rate $\gamma_l > 0$ whereas high types remain high types. At $t = 0$, the initial mass of followers is zero. I denote with $i(t) \in \{l, h\}$ the type of a generic follower at time t and I index followers with i where no confusion can arise.

Strategies. Followers seek information about a state of the world $x \in \mathbb{R}$, where $x \sim F$. An sender (she) can reveal high quality information q_h, low quality information q_l, or both at each time. Formally, she chooses a disclosure strategy $\{q_t\}_{t\geq 0}$, where $q_t = \{q_{nt}\}_{n=1,2}$, $q_{nt} \in \{q_l, q_h\}$, and $q_h > q_l$.[6] With slight abuse of notation, I write $q_t = q$ if $q_{1t} = q_{2t} = q$ for $q \in \{q_l, q_h\}$. Information disclosure is costless for simplicity. Given the sender's choice of information, each follower chooses whether to pay attention. Given quality q, followers who pay attention learn x truthfully if $x \geq 1/q$ and receive no information if $x < 1/q$.

[5] I take the intuitive approach to aggregating random variables over infinitesimal followers, which is standard in this literature. See e.g. Garrett (2016).

[6] Thus, the sender posts a menu of high quality and low quality content at each time. Allowing the sender to post a menu is necessary for an equilibrium to exist. If the sender is restricted to strategies $q_t \in \{q_l, q_h\}$ instead, no equilibrium exists. Intuitively, in any equilibrium, the mass of active high types or low types accumulates over time, and whenever this mass of types is positive, the sender strictly profits from deviating by offering each type his preferred information quality.

Paying attention is costly, with cost c_q, where $c_h > c_l$. Intuitively, processing high quality information is more difficult than processing low quality information. Followers who pay attention then choose action $a_t \in \{0, 1\}$.[7] For example, followers may seek the sender's information on whether to buy a product, visit an external website, or vote for a particular candidate. Followers who have chosen an action exit the platform, as they have no further need for information. Thus, each follower's attention decisions can be summarized by a stopping time $\tau_i \geq 0$. A follower is *active* at time t if he has entered before time t and not exited yet. Figure 1 contains a heuristic timeline.

Payoffs. There are no transfers between sender and followers. Given the sender's information, followers choose a_t to maximize the expectation of $xa - \kappa$ for some $\kappa \in (0, 1)$. I assume that $E(x) < \kappa$, i.e. the followers choose $a = 0$ if the sender does not provide information, and that $1/q_h > \kappa$.[8] The expected value from paying attention for type $i \in \{l, h\}$ is given by

$$v_i(q) = \max_{a(x), a_{ND} \in \{0,1\}} \pi_i \int_{\frac{1}{q}}^{\infty} (x - \kappa) a(x) \, dF(x) + F\left(\frac{1}{q}\right) \left(E\left(x | x < \frac{1}{q}\right) - \kappa\right) a_{ND}.$$

The assumptions above yield $a_{ND} = 0$, since $E(x|x < 1/q) < E(x) < \kappa$, and $a(x) = 1$ for $x \geq 1/q$, so that

$$v_i(q) = \pi_i \int_{\frac{1}{q}}^{\infty} (x - \kappa) \, dF(x).$$

The sender prefers followers to choose $a = 1$, e.g. she prefers them to buy an advertised product or vote for a particular candidate. The sender's expected payoff from a follower paying attention is given by

$$u(q) = \int_{\frac{1}{q}}^{\infty} a(x) \, dF(x) + \int_{-\infty}^{\frac{1}{q}} a_{ND} dF(x).$$

The assumptions above yield

$$u(q) = 1 - F\left(\frac{1}{q}\right).$$

We have

$$u(q_h) > u(q_l), \ v_i(q_h) > v_i(q_l) \ \text{ for } i \in \{l, h\}, \text{ and } v_h(q) > v_l(q) \ \text{ for } q \in \{q_l, q_h\}. \tag{1}$$

Thus, the main friction which limits information quality is whether followers pay attention, and not the fact that the sender prefers $a = 1$.

Finally, I assume that

$$v_l(q_l) - c_l > v_l(q_h) - c_h > 0 \text{ and } v_h(q_h) - c_h > v_h(q_l) - c_l. \tag{2}$$

The first inequality ensures that low types prefer low quality to high quality information, while the last inequality ensures that high types prefer high quality information. The second inequality ensures that both types' value from acquiring information is positive.[9]

[7] Given the assumptions on $v(q)$ and c_q below, making a choice without acquiring information is dominated by acquiring information for both types.

[8] This also ensures that the followers always prefer to wait if they receive no information for some time, as in Section 4.2.

[9] Note that $v_h(q_l) - c_l > v_l(q_l) - c_l$.

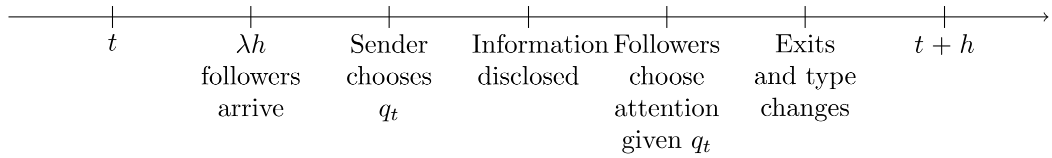

Figure 1: Heuristic Timeline.

Equilibrium. The equilibrium concept is a dynamic Stackelberg equilibrium, in which the sender does not have commitment across time. Intuitively, at any time t, the sender chooses a continuation strategy, given the best responses of all current and future followers. Followers take the sender's strategy as given and choose optimally when to acquire information. Formally, an equilibrium consists of a strategy for the sender $\{q_t\}_{t\geq 0}$, so that for any time $t \geq 0$, the continuation strategy $q^s \equiv \{q_s\}_{s\geq t}$ is optimal given the best responses $\{\tau_i(q^s, t')\}_{i\in\{l,h\},t'\geq t}$. Here, $\tau_i(q^s, t')$ is the optimal stopping strategy of a type $i \in \{l, h\}$ who is active at time $t' \geq t$ given the sender's continuation strategy q^s.

The sender's value is given by

$$U(t) = \sup_{\{q_s\}_{s\geq t}} E\left[\int_t^\infty e^{-r(s-t)} u(q_s)\, dn_s\right],$$

where n_s is the mass of followers who have paid attention up to time s, which implicitly depends on the stopping strategies $\tau_i(q^0, t)\}_{i\in\{l,h\},t\leq s}$. The followers' values are given by

$$V_l(t) = \sup_{\tau_l \geq 0} E\left[e^{-(r+\rho)(\tau_l - t)} \left(v_{i(\tau_l)}(q_{\tau_l}) - c_{q_{\tau_l}}\right)\middle| i(t) = l\right],$$

where $i(t) \in \{l, h\}$ is the follower's type at time t, and

$$V_h(t) = \sup_{\tau_h \geq 0} E\left[e^{-(r+\rho)(\tau_h - t)} \left(v_h(q_{\tau_h}) - c_{q_{\tau_h}}\right)\right].$$

4 Analysis

4.1 Equilibrium

Suppose that the sender chooses $q_t = q$ for all t. Then, the high type immediately acquires information, since there is no value of delaying. The low type either acquires information immediately or waits until he becomes the high type. His value from waiting is

$$V_l = (v_h(q) - c_q) \int_0^\infty \gamma_l e^{-(r+\rho+\gamma_l)t} dt = \frac{\gamma_l}{r+\rho+\gamma_l}(v_h(q) - c_q),$$

and the low type prefers to wait whenever

$$v_l(q) - c_q \leq \frac{\gamma_l}{r+\rho+\gamma_l}(v_h(q) - c_q).$$

Suppose that

$$\frac{v_l(q_l)-c_l}{v_h(q_l)-c_l} \geq \frac{\gamma_l}{r+\rho+\gamma_l} \geq \frac{v_l(q_h)-c_h}{v_h(q_h)-c_h}. \tag{3}$$

Then, the low type prefers to wait if $q_t = q_h$ for all t but prefers to acquire information immediately if $q_t = q_l$ for all t. This represents the key friction in the model, and I maintain this assumption throughout.[10]

The sender prefers to choose $q_t = q_h$ at all times. However, given Condition (3), this leads the low type to wait, so that the mass of active low types accumulates. Since the sender lacks commitment, whenever the mass of low types is positive, she prefers to deviate and choose q_l for an instant. Then, the accumulated low types pay attention and the sender receives a discrete gain. Because of this logic, there cannot be an equilibrium in which the sender exclusively provides high quality information. As I show in the proposition below, the unique equilibrium instead has the sender post high quality and low quality content.

Proposition 1. *Suppose that*

$$u(q_l) > \frac{\gamma_l}{r+\rho+\gamma_l} u(q_h). \tag{4}$$

Then, the unique equilibrium features $q_t = \{q_l, q_h\}$ for all $t \geq 0$.

The proof proceeds via a sequence of lemmas.

Lemma 1. *There exists no equilibrium in which $q_s = q_h$ for $s \in (t,t')$ and $q_{t'} \neq q_h$. Similarly, there exists no equilibrium in which $q_s = q_l$ for $s \in (t,t')$ and $q_{t'} \neq q_l$.*

Proof. Suppose by way of contradiction that $q_s = q_h$ for $s \in (t,t')$ and that $q_{t'} = q_l$.[11] Then, the low type prefers to wait whenever $s \in (t'',t') \subset (t,t')$ since

$$v_l(q_h) - c_h \leq e^{-(r+\rho)(t'-s)}\left((v_l(q_l)-c_l)(1-h(t'-t)) + (v_h(q_l)-c_l)h(t'-t)\right)$$

for any such s. Here, the RHS is the low type's value from waiting until t', and

$$h(t'-t) = \Pr(i(t') = h | i(s) = l) = e^{-\gamma_l(t'-s)} \tag{5}$$

is the probability that the low type has switched to high by time t'. The high type instantly acquires information for any such s, so that the mass of active low types m_s^l accumulates according to

$$dm_s^l = \left(\lambda(1-\gamma) - (\rho+\gamma_l) m_s^l\right) ds \tag{6}$$

and

$$m_s^l = \left(m_{t''}^l - \frac{\lambda(1-\gamma)}{\rho+\gamma_l}\right) e^{-(\rho+\gamma_l)(s-t'')} + \frac{\lambda(1-\gamma)}{\rho+\gamma_l}.$$

Also, we have

$$dn_s = \left(\lambda\gamma + \gamma_l m_s^l\right) ds,$$

[10] Without Condition (3), the equilibrium features $q_t = q_h$ for all t. If $(v_l(q_h)-c_h)/(v_h(q_h)-c_h) > \gamma_l/(r+\rho+\gamma_l)$, both types immediately acquire information if $q_t = q_h$, and if $\gamma_l/(r+\rho+\gamma_l) > (v_l(q_l)-c_h)/(v_h(q_l)-c_h)$, the low type waits given $q_t = q_l$. In both cases, the sender optimally chooses $q_t = q_h$.

[11] In the case $q_{t'} = \{q_l, q_h\}$, the analysis is identical.

since at any time s, newly arriving high types and low types who switch to become high types pay attention. The sender's payoff at $t' - \varepsilon$ is

$$U(t' - \varepsilon) = u(q_h) \int_0^\varepsilon e^{-rs} \left(\lambda\gamma + \gamma_l m^l_{t'-\varepsilon+s}\right) ds + e^{-r\varepsilon} \left(u(q_l) m^l_{t'} + U(t')\right).$$

Take ε sufficiently small and consider the continuation policy $\{\hat{q}_s\}_{s \geq t'-\varepsilon}$ so that $\hat{q}_{t'-\varepsilon} = q_l$ and $\hat{q}_s = q_s$ for $s > t' - \varepsilon$. This yields a payoff

$$\hat{U}(t' - \varepsilon) = m^l_{t'-\varepsilon} u(q_l) + u(q_h) \int_0^\varepsilon e^{-rs} \left(\lambda\gamma + \gamma_l \hat{m}^l_s\right) ds + e^{-r\varepsilon} \left(u(q_l) \hat{m}^l_\varepsilon + U(t')\right),$$

where

$$\hat{m}^l_s = \frac{\lambda(1-\gamma)}{\rho + \gamma_l} \left(1 - e^{-(\rho+\gamma_l)s}\right) \tag{7}$$

is the mass of active low types under $\{\hat{q}_s\}_{s \geq t'-\varepsilon}$. We have

$$\begin{aligned} \hat{U}(t' - \varepsilon) - U(t' - \varepsilon) &= m^l_{t'-\varepsilon} u(q_l) + \gamma_l u(q_h) \int_0^\varepsilon e^{-rs} \left(\hat{m}^l_s - m^l_{t'-\varepsilon+u}\right) ds + e^{-r\varepsilon} u(q_l) \left(\hat{m}^l_\varepsilon - m^l_{t'}\right) \\ &= \left(u(q_l) - \frac{\gamma_l}{r + \rho + \gamma_l} u(q_h)\right) \\ &\quad \cdot \left(\frac{\lambda(1-\gamma)}{\rho + \gamma_l} + \left(m^l_{t''} - \frac{\lambda(1-\gamma)}{\rho+\gamma_l}\right) e^{-(\rho+\gamma_l)(t'-\varepsilon-t'')}\right) \left(1 - e^{-(r+\rho+\gamma_l)\varepsilon}\right) \end{aligned}$$

which is positive given Condition (4). Thus, a strictly profitable deviation exists.

Suppose instead that $q_s = q_l$ for $s \in (t, t')$ and $q_{t'} = q_h$ or $q_{t'} \in \{q_l, q_h\}$. A similar argument as above implies that there is a profitable deviation, by choosing $\hat{q}_{t'-\varepsilon} = q_h$. See Internet Appendix A. □

Lemma 2. *There exists no equilibrium in which $q_s = q$ for all $s \geq t$ and any $t \geq 0$.*

Proof. If $q_s = q_h$ for all $s \geq t$, then low types wait and therefore $m^l_s > 0$ for all $s > t$. But then the same argument as above implies that the sender can deviate by momentarily switching to q_l. Thus, no such equilibrium exists. If $q_t = q_l$, all types pay attention and thus $m^h_s = m^l_s = 0$. Then,

$$dn_t = \lambda dt$$

and the sender's payoff at time t is given by

$$U(t) = u(q_l) \int_t^\infty e^{-r(s-t)} \lambda ds = \frac{\lambda}{r} u(q_l).$$

The strategy $\hat{q}(s) = q_h$ for $s \in (t, t+h)$ and $\hat{q}(s) = q_l$ for $s \geq t + h$ implies that high types instantly acquire information and low types wait, so that

$$\hat{U}(t) = u(q_h) \int_0^h e^{-rs} \left(\lambda\gamma + \gamma_l m^l_s\right) ds + m^l_h e^{-rh} u(q_l) + e^{-rh} U(t+h)$$

where m_s^l is given by Equation (7). Thus,

$$\hat{U}(t) - U(t) \geq u(q_h)\lambda\gamma \int_0^h e^{-rs} ds + m_h^l e^{-rh} u(q_l) - \lambda u(q_l) \int_0^h e^{-rs} ds$$

so that

$$\begin{aligned} r\frac{\hat{U}(t) - U(t)}{1 - e^{-rh}} &\geq \lambda\gamma u(q_h) + r\frac{\lambda(1-\gamma)}{\rho + \gamma_l} \frac{\left(1 - e^{-(\rho+\gamma_l)h}\right)}{\left(1 - e^{-rh}\right)} e^{-rh} u(q_l) - \lambda u(q_l) \\ &\to \lambda\gamma\left(u(q_h) - u(q_l)\right) > 0 \end{aligned}$$

as $h \to 0$. Thus, the deviation is profitable for h sufficiently small. □

Finally, consider a candidate equilibrium with $q_t = \{q_l, q_h\}$ for all t. Then, high type acquires high information and low type acquires low information, so that the sender's value is simply given by

$$U(t) = \frac{\lambda}{r}\left(\gamma u(q_h) + (1-\gamma) u(q_l)\right),$$

and $m_t^l = m_t^h = 0$ for all t. By the dynamic programming principle, it is sufficient to consider deviations for some finite interval $(t, t+h)$. Deviating to $\hat{q}(s) = q_h$ for $s \in (t, t+h)$ yields

$$\hat{U}(t) = u(q_h) \int_0^h e^{-rs}\left(\lambda\gamma + \gamma_l m_s^l\right) ds + e^{-rh} m_h^l u(q_l) + e^{-rh} U(t)$$

and

$$\frac{d\hat{U}(t)}{dh} = \left(u(q_h)\gamma_l - (r + \rho + \gamma_l) u(q_l)\right) \frac{\lambda(1-\gamma)}{\rho + \gamma_l} e^{-rh}\left(1 - e^{-(\rho+\gamma_l)h}\right),$$

which is negative given Condition (2). Thus, no such deviation is profitable. Similarly, deviating to $q_t = q_l$ for some time yields a strictly lower value, since $u(q_l) < u(q_h)$. Thus, the unique equilibrium features $q_t = \{q_l, q_h\}$ for all t.

4.2 Deprioritizing Content

I now consider prioritization rules which implement high quality information, i.e. where the sender chooses $q_t = q_h$ for all t. A *prioritization rule* determines the likelihood that the sender's information is shown to followers at each time. Formally, a prioritization rule consists of a sequence of (possibly random) arrival times $\{\tau_n\}_{n\geq 1}$, so that the inter-arrival times are distributed according to the cdf G_n, i.e. $\tau_{n+1} - \tau_n \sim G_n(t)$ conditional on a realization τ_n. Thus, $G_n(t+h) - G_n(t)$ is the probability that the sender's information reaches followers on any interval of time $(t, t+h)$. Unlike in the baseline model, the sender's information may be *deprioritized*, i.e. it is hidden from followers with a certain probability. A prioritization rule is optimal if it maximizes the sender's value subject to implementing $q_t = q_h$ for all t.[12] Figure (2) provides a heuristic timeline.

[12] Alternatively, I could define an optimal prioritization rule to maximize the high type's or the low type's value, subject to implementing high quality. The optimum is identical in those cases.

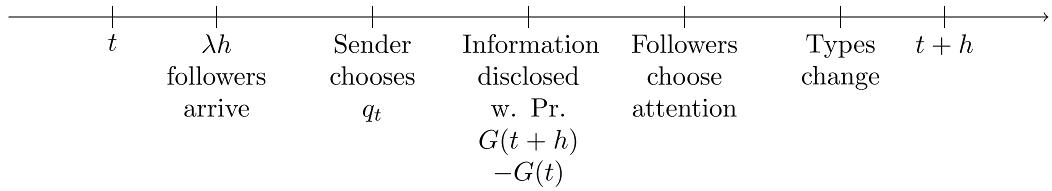

Figure 2: Timeline given G.

The optimal prioritization rule is remarkably simple. Once information arrives, the sender is deprioritized for a fixed amount of time t^*, during which none of her information reaches any followers. After that time, her information reaches followers with probability one, and after that she is again deprioritized. Thus, the sender gets to provide information at times $t \in \{0, t^*, 2t^*, ...\}$.

Proposition 2. *The optimal prioritization rule features $G_n = G$ for all n such that $dG(t) = 0$ for $t < t^*$ and $dG(t^*) = 1$.*

In the benchmark model, the main friction is that low types choose to wait, which then creates an incentive for the sender to provide low quality information. By deprioritizing future information, the platform discourages low types from waiting, and they instead acquire information whenever it is available. But then, the sender is no longer tempted to provide low quality information.

Proof. Consider an equilibrium in which $q_t = q_h$ and in which the low type acquires information instantly. Given disclosure time τ_n, the low type's value of waiting until the next disclosure time is given by

$$V_l(\tau_n) = E\left[e^{-(r+\rho)(\tau_{n+1}-\tau_n)}\left(h(\tau_{n+1}-\tau_n)(v_h(q_h)-c_h)+(1-h(\tau_{n+1}-\tau_n))(v_l(q_h)-c_h)\right)\right],$$

where the expectation is taken with respect to G, which is stationary without loss of generality. The low type prefers to acquire information at τ_n if

$$V_l(\tau_n) \leq v_l(q_h) - c_h, \tag{8}$$

and sender's value is given by

$$U(\tau_n) \;=\; E\left[e^{-r(\tau_{n+1}-\tau_n)}\left(m_{\tau_{n+1}-\tau_n} u(q_h) + U(\tau_{n+1})\right)\right],$$

where

$$m_t = \frac{\lambda}{\rho}\left(1 - e^{-\rho t}\right)$$

is the mass of all active followers if all followers wait between time 0 and t. Since G is stationary, we have $U(\tau_{n+1}) = U(\tau_n) = U$ so that the optimal G is characterized by the fixed point equation

$$U^* = \sup_G \int_0^\infty e^{-rt}(m_t u(q_h) + U^*)\, dG(t) \tag{9}$$

subject to Condition (8) and the feasibility condition $\int_0^\infty dG(t) = 1$.

Treating the optimal value $U^* > 0$ as a constant in the objective function (9), this problem admits the Lagrangian[13]

$$\mathcal{L} = \sup_G \inf_{\eta,\mu} \int_0^\infty u_t dG(t) + \mu \left(v_l(q_h) - c_h - \int_0^\infty v_t dG(t) \right) + \eta \left(1 - \int_0^\infty dG(t) \right),$$

where

$$u_t = e^{-rt}\left(m_t u(q_h) + U^*\right) \text{ and } v_t = e^{-(r+\rho)t}\left(v_h(q_h) - c_h - e^{-\gamma_l t}\Delta v(q_h)\right).$$

Rewriting the objective function yields

$$\mathcal{L} = \sup_G \inf_{\eta,\mu} \int_0^\infty (u_t - \mu v_t - \eta)\, dG(t) + \eta + \mu\left(v_l(q_h) - c_h\right).$$

A necessary condition for the objective to be finite is

$$u_t - \mu v_t - \eta \leq 0 \; \forall t \geq 0,$$

and the dual problem is given by

$$\begin{aligned} &\inf_{\mu \geq 0, \eta} \mu\left(v_l(q_h) - c_h\right) + \eta \\ s.t. \quad & u_t - \mu v_t - \eta \leq 0 \; \forall t \geq 0. \end{aligned}$$

For any given μ, the optimal η in the dual problem satisfies $\eta = \max_t u_t - \mu v_t$. Note that we have $\max_t u_t - \mu v_t \in (0, \infty)$. Otherwise, if $u_t - \mu v_t < 0$ for all t and then the primal problem would feature $dG(t) = 0$ for all t, which cannot be optimal. Moreover, we have $\lim_{t\to\infty} u_t = \lim_{t\to\infty} v_t = 0$, and thus $\lim_{t\to\infty} u_t - \mu v_t = 0$, so that the maximum is finite and attained at a finite time.

We have

$$u_t'' - \mu v_t'' = -r\left(u_t' - \mu v_t'\right) - (r+\rho)^2 e^{-(r+\rho)t}\frac{\lambda}{\rho}u(q_h) - \mu(r+\rho+\gamma_l)e^{-(r+\rho+\gamma_l)t}\Delta v(q_h) + \mu\rho v_t',$$

where

$$v_t' = -(r+\rho)v_t + (r+\rho+\gamma_l)e^{-(r+\rho+\gamma_l)t}\Delta v(q_h).$$

Thus, there is a unique global maximum, because $u_t' - \mu v_t' = 0$ implies that $u_t'' - \mu v_t'' \neq 0$.

Letting $t^* = \arg\max_t u_t - \mu v_t$, we have $dG(t) = 0$ for all $t \neq t^*$ and $dG(t^*) = 1$. That is, under the optimal policy, the platform delays disclosure until t^*. Since the IC constraint (8) must bind, we have $t^* = t_l$, where

$$v_l(q_h) - c_h = v_{t_l}. \tag{10}$$

The function v_t is either monotonically decreasing or hump-shaped, since

$$v_t'' = -(r+\rho)v_t' - (r+\rho+\gamma_l)e^{-(r+\rho+\gamma_l)t}\Delta v(q_h),$$

[13]See e.g. Anderson and Nash (1987), Th. 2.1.

so that $v''_t < v'_t$. In particular, $v''_t < 0$ whenever $v'_t \leq 0$. We have $v_0 = v_h(q_h) - c_h > v_l(q_h) - c_h$. Thus, there exists a unique t_l such that the indifference condition (10) holds. Finally, the multiplier μ is determined implicitly, so we indeed have $t^* = t_l$.

It remains to characterize the fixed point in Equation (9). Since $t^* = t_l$, and since t_l is independent of U^*, we simply have

$$U^* = u_{t_l} + e^{-rt_l} U^*$$

or equivalently

$$U^* = \frac{u_{t_l}}{1 - e^{-rt_l}}.$$

Finally, since given t^* both types pay attention at any disclosure time, the sender optimally chooses $q(\tau_n) = q_h$, since any other choice leads to a strictly lower value at the disclosure times. Since no information reaches followers for $t \neq \tau_n$, the sender is indifferent over any choice at those time, and we can set $q_t = q_h$ without loss of generality. □

The following comparative statics immediately follow from Equation (10).

Corollary 1. *Time t^* is strictly increasing in c_h, γ_l, and v_h, and strictly decreasing in r, ρ, and v_l.*

Intuitively, the information acquisition cost increases the low type's value of waiting relative to acquiring information now, so that t^* has to increase to keep the low type indifferent. Similarly, if the low type is more likely to become high, i.e. γ_l is higher, the value of waiting increases. Conversely, the value of waiting decreases in the discount factor r and the rate of exit ρ, and the value of acquiring information immediately increases in v_l.

4.3 Monetizing Content

The analysis so far abstracted from monetary payments. In reality, content on platforms generates advertising revenue, which is shared between content creators and the platform. I now characterize the least-cost advertising split which implements $q_t = q_h$ for all t. That is, the platform aims to ensure that content is high quality, while minimizing the share of advertising revenue the sender receives.

Formally, the platform generates advertising revenue A per view. Recall that n_t is the cumulative mass of followers who have paid attention up to time t. Then, the sender generates advertising revenue $A dn_t$ per unit of time. The advertising split $\{\alpha_t\}_{t\geq 0}$, where $\alpha_t \in [0, A]$, determines the share of revenue the sender receives. I consider advertising splits that are functions of current views only, i.e. $\alpha_t = \alpha(dn_t)$ with some abuse of notation. Intuitively, on a small window of time $(t, t+h)$, $n_t - n_{t-h}$ followers view the sender's content. Then, the sender receives advertising revenue $\alpha(n_t - n_{t-h})$ per view, so that her total revenue equals $\alpha(n_t - n_{t-h}) \cdot (n_t - n_{t-h})$.[14] I denote the set of all such policies with $\mathcal{A}$. Figure 3 provides a heuristic timeline.

[14]In equilibrium, dn_t is of order dt, so total revenue is simply $\alpha_t \cdot \left(\lambda\gamma + \gamma_l m_t^l\right) dt$. See Equation (11) below.

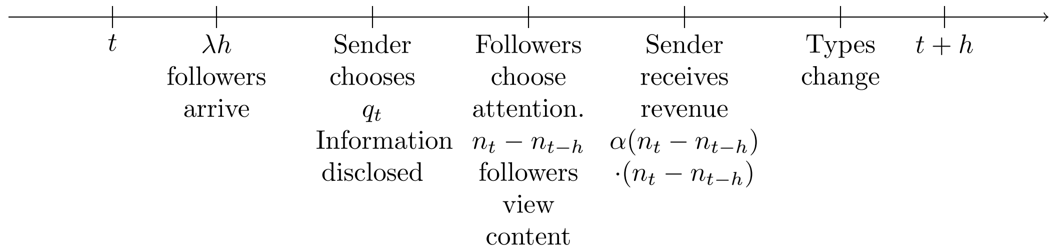

Figure 3: Timeline with advertising revenue.

Given $\{\alpha_t\}_{t\geq 0}$, the sender's payoff is

$$U(t) = \sup_{\{q_s\}_{s\geq t}} E\left[\int_t^\infty e^{-r(s-t)}\left(u(q_s) + \alpha_s\right) dn_s\right].$$

In an equilibrium where $q_t = q_h$ for all t, high types pay attention immediately and low types wait until they become high types. Therefore,

$$dn_t = \left(\lambda\gamma + \gamma_l m_t^l\right) dt, \tag{11}$$

where m_t^l is given in Equation (6). Then, the sender's payoff becomes

$$U(t) = \int_t^\infty e^{-r(s-t)}\left(\lambda\gamma + \gamma_l m_s^l\right)\left(u(q_h) + \alpha_s\right) ds.$$

Since low types wait to become high types, we have $m_t^l > 0$ for any $t > 0$. Just as in the baseline model, the sender now has an incentive to deviate and to set $q_t = q_l$ to get all low types to instantly pay attention. Then, $n_t - n_{t-} = m_{t-}^l > 0$ and the sender receives a discrete gain of $(u(q_l) + \alpha_t)m_{t-}^l$. To discourage such a deviation, the optimal ad split sets $\alpha_t = 0$ whenever $n_t - n_{t-} > 0$. Thus, by deviating, the sender gains $u(q_l)m_{t-}^l$, but then since all low types pay attention, we have $m_t^l = 0$. Effectively, the sender "resets" the mass of active low types to zero, which affects her future advertising revenue. Her continuation value is then given by $U(0)$. Overall, deviating is suboptimal whenever

$$u(q_l)\, m_t^l + U(0) \leq U(t) \;\; \forall t \geq 0. \tag{12}$$

The least-cost advertising split which implements $q_t = q_h$ is hence given by

$$\min_{\{\alpha_t\}_{t\geq 0}\in\mathcal{A}} \int_0^\infty e^{-rt}\left(\lambda\gamma + \gamma_l m_t^l\right)\alpha_t dt \tag{13}$$

$$s.t.\; u(q_l)\, m_t^l + U(0) \leq U(t) \;\; \forall t \geq 0$$

$$0 \leq \alpha_t \leq A \;\forall t \geq 0$$

Proposition 3. *If*

$$A \geq \frac{\lambda(1-\gamma)}{\lambda\gamma + \gamma_l \frac{\lambda(1-\gamma)}{r+\rho+\gamma_l}}\left(u(q_l) - u(q_h)\frac{\gamma_l}{r+\rho+\gamma_l}\right), \tag{14}$$

then the optimal policy features $\alpha_t = 0$ for $t < t^{**}$ and $\alpha_t = A$ for $t \geq t^{**}$, where $t^{**} \in [0, \infty)$. Otherwise, implementing $q_t = q_h$ for all t is not feasible.

Intuitively, since

$$dn_t = \left(\lambda\gamma + \gamma_l m_t^l\right) dt,$$

the number of "views" whenever the sender provides information is a deterministic function of m_t^l. When the sender deviates, the mass of active low types resets to zero, and she must wait until the mass becomes sufficiently large (i.e. it reaches $m_{t^{**}}^l$ at time t^{**}, so that dn_t is sufficiently large) before she receives ad revenue again. If she instead continues, she either receives ad revenue sooner (if $t < t^{**}$) or continues receiving ad revenue (if $t \geq t^{**}$). Thus, by providing payments only when the number of views is sufficient large, the platform can discourage deviations. The time t^{**} is the maximum delay which ensures that the sender's IC condition holds at $t = 0$.

To see this formally, decompose Constraint (12) as

$$g_t + f_t \leq 0,$$

where

$$\begin{aligned} g_t &= u(q_l)\, m_t^l + \int_0^\infty e^{-rt} u(q_h) \left(\lambda\gamma + \gamma_l m_t^l\right) dt \\ &\quad - \int_t^\infty e^{-r(s-t)} u(q_h) \left(\lambda\gamma + \gamma_l m_s^l\right) ds \end{aligned} \tag{15}$$

and

$$f_t = \int_0^\infty e^{-rt} \alpha_t \left(\lambda\gamma + \gamma_l m_t^l\right) dt - \int_t^\infty e^{-r(s-t)} \alpha_s \left(\lambda\gamma + \gamma_l m_s^l\right) ds.$$

Intuitively, g_t captures the sender's willingness to deviate if she never receives any ad revenue, and f_t captures how ad revenue affects her incentive to deviate.

Lemma 3. *If $\alpha_t > 0$, then $\alpha_{t'} = A$ for all $t' > t$.*

Proof. Suppose not. Then, for some $t' > t$, we have $\alpha_{t'} < A$. Consider the alternative policy $\{\hat{\alpha}_t\}_{t \geq 0}$ so that $\hat{\alpha}_s = \alpha_s$ for $s \notin \{t, t'\}$, $\hat{\alpha}_{t'} = \alpha_{t'} + \varepsilon$, and $\hat{\alpha}_t = \alpha_t + \hat{\varepsilon}$ where

$$\left(\lambda\gamma + \gamma_l m_t^l\right) \hat{\varepsilon} = -e^{-r(t'-t)} \left(\lambda\gamma + \gamma_l m_{t'}^l\right) \varepsilon. \tag{16}$$

This does not change the objective in Equation (13), weakly decreases f_s for $s \in (t, t')$, and otherwise leaves f_s unchanged. If the set of times at which $\alpha_t \in (0, A)$ has positive Lebesgue measure, then we can strictly relax Condition (12) by setting $\hat{\alpha}_t$ according to Equation (16). □

Thus, without loss of generality, the optimal policy features $\alpha_t = 0$ for $t \leq t^{**}$ and $\alpha_t = A$ for $t > t^{**}$, where $t^{**} \in [0, \infty)$. It only remains to determine the optimal t^{**}. We have $g_0 = 0$, $g_t' > 0$ and $g_t'' < 0$ for all $t > 0$, which follows from differentiating Equation (15). Similarly, we have $f_0 = 0$, $f_t' < 0$ and $f_t'' < 0$ for all $t > 0$. See Internet Appendix B for derivations. Thus, $g_t + f_t$ is strictly concave. Since $g_0 + f_0 = 0$, we have $g_t + f_t \leq 0$ for all t if and only if

$$g_0' + f_0' \leq 0$$

or equivalently

$$(1-\gamma)\left(u\left(q_l\right)-u\left(q_h\right)\frac{\gamma_l}{r+\rho+\gamma_l}\right)-rAe^{-rt^{**}}\int_{t^{**}}^{\infty}e^{-r\left(s-t^{**}\right)}\left(\lambda\gamma+\gamma_l m_s^l\right)ds\leq 0. \tag{17}$$

There exists a t^{**} such that this inequality holds whenever

$$\max_T A\int_T^{\infty}e^{-rs}\left(\lambda\gamma+\gamma_l m_s^l\right)ds\geq\frac{\lambda\left(1-\gamma\right)}{r}\left(u\left(q_l\right)-u\left(q_h\right)\frac{\gamma_l}{r+\rho+\gamma_l}\right).$$

In this problem, the LHS is maximized at $T=0$. Thus, a necessary and sufficient condition is that

$$A\int_0^{\infty}e^{-rs}\left(\lambda\gamma+\gamma_l m_s^l\right)ds\geq\frac{\lambda\left(1-\gamma\right)}{r}\left(u\left(q_l\right)-u\left(q_h\right)\frac{\gamma_l}{r+\rho+\gamma_l}\right),$$

which is equivalent to Equation (14). Then, the optimal t^{**} is pinned down by letting Inequality (17) bind.

5 Concluding Discussion

Modeling of Information. This setting resembles Dye (1985) and is chosen for simplicity. It captures the idea that worse quality information leads to worse decisions in a parsimonious way and allows the followers to realize rents, which generates an option value of waiting for low type followers. Beyond these features, the particular modeling of information disclosure is irrelevant. Any setup that leads to values $u(q)$, $v_l(q)$, and $v_h(q)$ with the same properties as those in Equations (1) and (2) generates the same qualitative results.

Changing Types. Shocks to the followers' value of attention are necessary to generate a value of waiting for the low type. In the baseline model, I have assumed that low type followers may become high types, but that high types never become low types. This is mainly to simplify the algebra. Otherwise, when low and high types are forced to wait, as in Proposition 2, one has to solve a system of ODEs to determine both types' values. However, the results extend to that setting as long as the high type prefers to instantly acquire both high and low quality information. Thus, an alternative setup, where the high type becomes low at rate γ_h and where the mass of followers is initialized at the steady state (as in Garrett (2016)) yields qualitatively similar predictions.

Content-Specific Mechanisms. In the analysis, I have focused on mechanisms which do not depend on the quality chosen by the sender. Allowing for content-specific mechanisms renders the setting trivial. The platform can simply deprioritize the sender forever (i.e. ban her from the platform) after the sender posts low quality content. Such bans are indeed used in reality, in cases where content is clearly illegal. However, they are infeasible whenever content quality is difficult to measure.

References

Anderson, E. J. and P. Nash (1987). *Linear programming in infinite-dimensional spaces: theory and applications*. Wiley.

Basak, D. and Z. Zhou (2020a). Diffusing coordination risk. *American Economic Review 110*(1), 271–97.

Basak, D. and Z. Zhou (2020b). Timely persuasion. *Available at SSRN 3540708*.

Bizzotto, J., J. Rüdiger, and A. Vigier (2021). Dynamic persuasion with outside information. *American Economic Journal: Microeconomics 13*(1), 179–94.

Che, Y.-K., K. Kim, and K. Mierendorff (2020). Keeping the listener engaged: a dynamic model of bayesian persuasion. *arXiv preprint arXiv:2003.07338*.

Cong, L. W. and S. Li (2021). A model of influencer economy. *Available at SSRN 3974661*.

Dilmé, F. and D. F. Garrett (2022). A dynamic theory of random price discounts. Technical report, ECONtribute Discussion Paper.

Dilme, F. and F. Li (2019). Revenue management without commitment: Dynamic pricing and periodic flash sales. *Review of Economic Studies 86*(5), 1999–2034.

Dye, R. A. (1985). Disclosure of Nonproprietary Information. *Journal of Accounting Research 23*(1), pp. 123–145.

Fainmesser, I. P. and A. Galeotti (2021). The market for online influence. *American Economic Journal: Microeconomics 13*(4), 332–72.

Galeotti, A. and S. Goyal (2009). Influencing the influencers: a theory of strategic diffusion. *RAND Journal of Economics 40*(3), 509–532.

Garrett, D. F. (2016). Intertemporal price discrimination: Dynamic arrivals and changing values. *American Economic Review 106*(11), 3275–3299.

Hinnosaar, M. and T. Hinnosaar (2021). Influencer cartels. *Available at SSRN 3786617*.

Kapon, S. (2022). Dynamic amnesty programs. *American Economic Review 112*(12), 4041–75.

Kaya, A. (2022). Paying with information. *Theoretical Economics*. Publisher: Econometric Society.

Mitchell, M. (2021). Free ad (vice): Internet influencers and disclosure regulation. *RAND Journal of Economics 52*(1), 3–21.

Orlov, D., A. Skrzypacz, and P. Zryumov (2020). Persuading the Principal to Wait. *Journal of Political Economy 128*(7), 2542–2578.

Pei, A. and D. Mayzlin (2022). Influencing social media influencers through affiliation. *Marketing Science 41*(3), 593–615.