Online News and Editorial Standards

Atara Oliver*

March 19, 2022

Abstract

The internet enables a media firm to post articles at any time. I compare this situation to one in which news can only be released once. I determine the editorial standard, a cutoff for how confident a firm must be in order to post an article. If changing a story is costly, the firm's editorial standard is weakly higher when it can post at any time, and this standard decreases over time. This implies that a firm may be more cautious with releasing internet news. However, if the firm has a strong prior, it may post earlier with less information.

JEL Classification: D21, L82.

*ESMT Berlin Email: atara.oliver@esmt.org. This article is part of a project Digital Platforms: Pricing, Variety, and Quality Provision (DIPVAR) that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme (grant agreement No 853123).
An earlier version of the article was included in the author's doctoral dissertation at Rice University.
I thank Hulya Eraslan; Mallesh Pai; K. Ramesh; Özlem Bedrie-Defolie; Willy Lefez; David Ronayne; Saltuk Ozerturk, Mark Agerton; Ying Chen; Caroline Thomas; Tetsuya Shinkai; and seminar and conference participants at the EARIE Conference, ESMT, UCSD Rady School of Business, the US Food and Drug Administration Rice University, the Western Economic Association International Conference in Newcastle, the 29th Jerusalem Summer School in Economics, the 29th International Conference on Game Theory at Stony Brook, and the 2018 Texas Economic Theory Camp.

1 Introduction

The internet has changed the way that news is produced and accessed. Although some 24-hour cable news stations existed before the internet news era, the internet has enabled any media firm with a website to post news around-the-clock, and for consumers to read this news at any time. Additionally, the internet enables firms to update news stories as new information arrives. As consumers are unlikely to wait to read articles that are not available when they browse the internet, firms have an incentive to publish news quickly in order to obtain a larger share of consumers, possibly updating the story later if necessary.

However, the fact that a firm *can* post or update its news story any time, does not necessarily mean that it *should*. Not all leads are true, and in certain cases the desire to produce stories quickly has led firms to release information later found to be false. Such incorrect reporting is harmful both for the consumers and for the firm whose reputation is harmed.

The concern that firms may release news too quickly in order to gain readers is strengthened when a firm's editor, ownership, or management initially believes that a story is likely to be true. This is due to a firm believing that information consistent with its prior is more likely to be correct.[1] An example of this is the coverage of a shooting at a mosque in Quebec City, Canada, on January 29th, 2017. Fox News—a conservative-leaning American news agency—reported that the suspect was of Moroccan origin. In reality, the police initially identified two suspects—a French-Canadian and a Moroccan—but stated shortly afterwards that the Moroccan was just a witness. Although some other organizations initially mentioned both suspects until the situation became clearer, Fox News only mentioned the Moroccan suspect and did not retract it until the Canadian government requested they do so.[2] One

[1] Although the belief actually belongs to the individuals responsible for selecting the firm's content, for simplicity I refer to this as the firm's belief.

[2] For more information on this story, see http://www.cbc.ca/news/politics/kate-purchase-bill-shine-1.3960772

explanation for this is that, given Fox's political leanings, they viewed the lead regarding a Moroccan suspect as very plausible and therefore decided to go ahead and post that message rather than wait, whereas more left-leaning news agencies would view this story as implausible and therefore choose to wait for further confirmation before posting.

In this article, I examine the effect of the 24-hour news cycle on a firm's decision of when to post or update a news story. To do so, I model the process of internet news production and compare it to news production prior to the internet. In particular, to capture the dynamic nature of internet news, I depart from the typical approach in the media literature of viewing the news release decision as consisting of a single action. Instead, I allow for the internet firm's decision to be fully dynamic in the sense that the firm can post an article or change an existing article at any time. I find that dynamic considerations—specifically, the option value of waiting—cause the firm to use more caution when choosing whether or not to release an article. Nevertheless, this may not lead to news that is based on more information.

I consider a continuous-time finite-horizon model in which there are two states of the world. A monopolist firm receives leads according to a Poisson process that provide information about the state of the world. I compare two scenarios—a "Before-Internet Scenario" (BIS), in which the firm can only post once, and an "After-Internet Scenario" (AIS), in which the firm can post an article or update an existing article any time. In each scenario, the firm can post any message, and its revenue is proportional to the share of consumers who read it. However, news that was reported and turns out to be incorrect will eventually be discovered, resulting in a penalty for the firm, where this penalty is also proportionate to the readership.[3] In the AIS, the firm also incurs a retraction cost when it changes articles. Therefore, the firm faces a tradeoff between posting earlier and getting more readers, or posting later and reducing the risk of posting false information.

I define the firm's *editorial standard* as a cutoff that determines how high a firm's posterior

[3]This penalty could be in the form of a reduced reputation in the future, or possibly a defamation lawsuit.

belief at about a state of the world x at any given time must be in order to initially post an article claiming that the state is x. I find that there is a well-defined editorial standard—meaning that the firm uses a cutoff strategy—that only in certain circumstances depends on the posting time. In the BIS, there is a fixed cutoff such that, if the firm's posterior belief that a particular event occurred is greater than this cutoff, the firm will produce an article saying that the event occurred. In the AIS, the firm will use the same cutoff if there is zero retraction cost, because the actions that it takes at any given moment to not affect it's payoffs when taking a different action in the future, so the firm only focuses on its instantaneous payoffs. If the retraction cost is positive, however, it may use a stricter cutoff, meaning that it needs to be more certain about the state of the world in order to post. This cutoff is decreasing in time if there is a retraction cost due to the option value of waiting.

In the case that the same cutoff is used, the firm will post with (weakly) less information in the AIS than in the BIS, because the firm would no longer be forced to wait until a predetermined news-release time. If a stricter cutoff is used, however, this may not be the case, because there is now an option value of waiting. This implies that, when the firm can post at any time, if there is a positive retraction cost it would use a stricter editorial standard than it would have before the internet, with this standard decreasing over time. This means that the firm may be more cautious about posting on the internet, and thus be less likely to post incorrect information. However, if the firm has a particularly strong prior belief about the event, it would be more likely to have a posterior belief that exceeds the cutoff with less evidence. The firm may therefore end up posting incorrect news that it would not have posted if it were in the BIS, where it would be forced by the publication process to wait. This implies that the firm's prior is particularly important, and the result is that a firm with a strong prior is more likely to release a news a story, which may contribute to news becoming more polarized in the internet era.

This research relates to the timing and correctness of news production. To my knowledge, this article is the first to model the practice of news firms releasing a message and then updating it as desired, rather than just publishing once. Cancian et al. (1995) examine a model in which firms can choose when to broadcast a news report given that consumers can only watch news that is broadcast after they arrive home. However, their model does not consider the effect of timing on the content produced.

There are several other articles which model the news timing decision as a preemption game where there is a tradeoff between timing and correctness. Lin (2013) examines a game in which an event occurs, after which two news firms play a preemption game, in which both getting story posted first and getting it correct increase a firm's payoffs. The firms find out the true state of the world through a signal that arrives according to a state-dependent Poisson process. In article that was written concurrently with this one, Pant and Trombetta (2019) consider a two-period model in which there are two firms who each receive an initial signal about the state of the world. The firms may be high-type firms capable of doing research and learning the true state with certainty, or they may be low-type firms who cannot. The tradeoff in their model is between releasing news immediately and, depending on the action of the other firm, possibly getting a preemption benefit, or posting later and possibly demonstrating that they are a high-type firm. As in this article, they find that the internet timing incentives may not lead to lower news quality. In another concurrently written article, Shahanaghi (2021) considers a continuous-time model in which firms learn about the state over time, and then choose when to submit a report. Each firm receives a private signal confirming the state of the world according to a Poisson process, but can choose to post (just once) at any time, whether or not a signal has been received. In this model, a media firm can learn about the state of the world both from its own signal and the actions of its competitors.

The approach here differs from these articles above in several important ways. As the internet allows firms to post news at any time—and therefore firms can nearly immediately copy each others posting decisions—preemption concerns are less salient in the internet era, and I therefore deliberately construct a model in which preemption is not a concern.[4] In the world with the internet, I view the main tradeoff that a firm faces is that if it delays posting, it will lose customers, whereas if it posts sooner it will be more likely to post incorrect news. This article also differs from the others mentioned above in that the firm in my model receives multiple partially-informative signals, and can update the article whenever it chooses to rather than simply reporting once. In addition, I explicitly compare the firm's decision in the internet era to the decision it would have made before the internet, when the news release time was fixed. Moreover, as the firm in my model cannot learn the state with certainty until after its decisions have been made, my analysis focuses on the role of its posterior about the state of the world in its decision-making process.

Although this article is not directly about media bias, it indirectly relates to the media bias literature by considering the effect of the firm's posterior at a given time on its posting decision. There has been substantial recent literature on media bias; a summary of the theoretical literature can found in Gentzkow, Shapiro, and Stone (2016) and of empirical literature can be found in Puglisi and Snyder (2016). Strömberg (2004) considers a model of media bias in which newspapers choose to cover news which appeals to larger groups in order to increase their profit from readership. In Chan and Suen (2008) and Duggan and Martinelli (2011), firms coarsen the message space in order to give readers an impression of the state of the world which the media outlets view as more favorable, and in Anderson and McLaren (2012) in which biased firms strategically fail to report negative information in order to give readers a different impression of the state of the world. Additionally, the setup in my model is similar to that of Gentzkow and Shapiro (2006), with their model including consumers

[4]This is not to say that preemption is entirely irrelevant in the internet era. There are some stories, such as those that require an extensive amount of research to produce, in which preemption still plays an important role. However, when reporting current events, such as a terrorist attack, in which any firm can easily release information after the story is known, preemption is no longer particularly important because it is no longer clear which firm came first.

who judge the accuracy of reporting based on their priors and therefore unintentionally view biased news as more accurate, just as the firms in my model view leads that confirm their priors as more accurate. Other explanations of supply-driven media bias include a media firm wanting to maintain access to politicians (Ozerturk (2019)) and wanting to appeal to advertisers (Ellman and Germano (2009); Germano and Meier (2013)).

This article also relates to the emerging literature on fake news. The term "fake news" refers to the *intentional* production of incorrect news, whereas this research considers the *unintentional* production of incorrect news. Grossman and Helpman (2019) considers a model in which politicians and partisan media outlets can choose to release incorrect information about their positions, the political positions of their rivals, or alternative policies. Allcott and Gentzkow (2017) consider a model in which media firms may produce fake news in order to cater to consumers who are interested in reading news that confirms their priors. They find empirical evidence about the prevalence of fake news and its role in the 2016 United States presidential election. Although this research is not about "fake news," it can be difficult for consumers to distinguish between news that is intentionally fake and news that is unintentionally incorrect, and the negative societal effects caused by fake news can also be found with incorrect news in general. Additionally, the prevalence of fake news may increase the reputational harm to a firm that unintentionally produces incorrect news, as consumers may assume it was done on purpose.

Moreover, this article is an application of the literature on single-agent decision problems in continuous time. In particular, the After Internet Scenario is an example of a continuous-time decision problem in which the state changes according to a jump process. There are well-known difficulties with defining a policy or strategy in continuous time (see Simon and Stinchcombe, 1989). To avoid these issues and to take advantage of the fact that the firm's information only changes at discrete times, I adapt the method of Khan and Stinchcombe (2015), which considers a continuous-time and infinite-horizon model. In this model, the state changing according to a jump process, and at each of these state changes, the agent chooses a plan. The model in this article is finite horizon, so I apply Bäuerle and Rieder (2010) and set up a discrete-time like Bellman equation, which is similar to that in Khan and Stinchcombe (2015) except for its time-dependence.

2 Model

Consider an environment with a single profit-maximizing firm and a unit measure of consumers. A news *event* occurs at time $t = 0$ that is a realization of a random variable X, which takes values in $\{0, 1\}$. The firm has a prior π_0^1 that the event is 1.[5] The firm receives leads, which arrive according to a Poisson process on $[0, 1]$ with constant arrival rate λ. A lead that arrives at time t is a message m_t, where $m_t \in \{0, 1\}$ is a statement that the event is x. If the state is 1, the message m_t is drawn from a $Bernoulli(q)$ distribution, where $q \in (0.5, 1)$, whereas if the state is 0 the message is drawn from a $Bernoulli(1 - q)$ distribution. At each time a lead arrives, the firm updates its prior using Bayes' rule. As the lead arrival rate is independent of the state, the firm learns nothing from the absence of leads.

I consider two scenarios—a "Before-Internet Scenario" (BIS) and an "After-Internet Scenario" (AIS). In the BIS, the firm can only choose to post an article at a set time, $\bar{t}$, which is the set daily time for news release. Because the news is always released at the same time, all consumers arrive at that time.[6] In the AIS, consumers arrive according to a Uniform distribution on $[0, 1]$. In this scenario, the firm can initially post or update its article at any time. After time $t = 1$, the consumers lose interest in the news event and therefore no new consumers arrive.[7] I assume that, at some point after the $t = 1$, the true state of the world is

[5] Note that it is not necessarily the case that this prior is equal to the true prior probability. The firm makes its decisions based on its subjective individual beliefs about the state of the world.

[6] In the case of a TV broadcast, the consumers were literally all arriving at once. In the case of a newspaper, this should be viewed more figuratively. The newspaper is printed once, cannot be changed once printed, and all of the consumers see the same content.

[7] This assumption reflect the fact that consumers lose interest in news after a relatively short time period,

revealed and becomes common knowledge. Each consumer can only read an article which is posted—meaning either initially posted or kept posted—at the instant at which she arrives.

2.1 Before Internet

In the BIS, the firm can post either message or no message at $\bar{t}$, so $a \in A = \{\varnothing, 0, 1\}$, with $a = \varnothing$ if the firm does not post and $a = x$ if the firm posts a message that the event is x.[8] If the firm does not post, its profit is normalized to 0, whereas if it posts, it receives advertising revenue which is normalized to 1, which is the mass of consumers who read it. A news story is considered to be *correct* if the message posted is equal to x when the event is x, and it is *incorrect* if the other message is posted. Let $\kappa > 1$ be the cost of posting an incorrect message.[9] When the event is x, the firm's profit from posting an article a is:

$$U(a|X = x) = \begin{cases} 0 & \text{if } a = \varnothing, \\ 1 & \text{if } a = X, \\ 1 - \kappa & \text{otherwise.} \end{cases} \tag{2.1}$$

If the event is 1, the Poisson process can be split into two independent Poisson processes—one with arrival rate $\lambda_1 = q\lambda$ and another with arrival rate $\lambda_0 = (1-q)\lambda$. If the event is 0, then $\lambda_1 = (1-q)\lambda$ and $\lambda_0 = q\lambda$. I refer to messages saying that the event is 1 as *type*-1 *messages* and messages saying that the event is 0 as *type*-0 *messages*. As this is a model with symmetric binary signals, the difference between the number of type-1 and type-0 messages, denoted by d, is sufficient to determine the posterior (as in Brocas and Carrillo (2007)), so

because newer news grabs their attention.

[8]This implies that the firm cannot post the leads themselves or its posterior. As media firms often need to put concise messages in their headlines and can at best put disclaimers in the article itself, this restriction is reasonable.

[9]If $\kappa < 1$, then it is always in the firm's best interest to post an article, so the decision is trivial. Moreover, having $\kappa > 1$ reflects the concern that media firms have about reputational damage from posting false news. Gentzkow and Shapiro (2006) provide examples that indicate that firms do care about their reputations as providers of correct news.

it is not necessary to consider the entire history.

A (pure) policy is a function σ such that, for any prior $\pi_0^1 \in [0,1]$ and any lead difference $d \in \mathbb{Z}$, $\sigma(\pi_0^1, d) \in A$. An optimal policy is a policy σ^* such that $\sigma^*(\pi_0^1, d) \in \arg\max_{a \in A} E[U(a|X)|\pi_0^1, d]$ for all $(\pi_0^1, d) \in [0,1] \times \mathbb{Z}$.

2.2 After Internet

In the AIS, at each time $t \in [0,1]$, the firm chooses an action $a_t \in A$, with $a_t = \varnothing$ if the firm does not post and $a_t = x$ if the firm posts a message that the event is $x \in \{0,1\}$. The per-consumer advertising revenue a firm makes from having any news story posted at time t is equal to 1, which is the density of the consumers who read it, and the per-consumer cost of having incorrect news posted is κ, with $\kappa > 1$.[10] This implies that at any time t, the firm's flow profit is $u(a_t|X = x) = (\mathbb{1}(a_t \neq \varnothing) - \mathbb{1}(a_t \notin \{\varnothing, x\})\kappa)$. As the time interval over which consumers are interested in a story is short, firms do not discount the future. Additionally, if the firm posts a story which it then removes, it faces a retraction cost. Formally, let a_t be the article that is posted at time t and a_{t+} be the article that is posted immediately afterwards. Then the retraction cost can be described by a function $r : A \times A \to \{0, \rho\}$, where $r(a_t, a_{t+})$ is equal to a constant $\rho \geq 0$ if $a_t \neq \varnothing$ and $a_t \neq a_{t+}$ and equal to 0 otherwise. In other words, the retraction cost is positive if and only if the firm had an article up at that instant, which it then removed.

The only information from the history that a firm needs to make a decision at time t is the time itself;[11] its posterior belief that the event is 1, π_t^1; and the action taken at t, a_t. In particular, it is not necessary for the firm to keep track of the leads received as those leads are summarized by the posterior. This implies that the firm's decision can be expressed as

[10] This could be viewed as a cost proportional to the probability of getting caught with false news, which may be more likely with a larger measure of readers. Alternatively, the damage to the firm's reputation may be proportional to the readership.

[11] More precisely, the firm is making the decision immediately after time t.

a Markov decision problem and analyzed using a recursive framework.[12]

I define the *expected instantaneous payoff* from an action as $\hat{u} : [0,1] \times A \to \mathbb{R}$, with:

$$\hat{u}(\pi_t^1, a_t) = \pi_t^1 u(a_t|X = 1) + (1 - \pi_t^1)u(a_t|X = 0). \tag{2.2}$$

This is equal to 0 if $a_t = \varnothing$, $(1 - (1 - \pi_t^1)\kappa)$ if $a_t = 1$, and $(1 - \pi_t^1\kappa)$ if $a_t = 0$.

As this is a continuous-time problem in which the firm can move any time, one cannot simply define the strategies as defining an action taken at each time t given π_t and a_t. Because the real-line is not well-ordered, there is no clearly defined last action prior to time-t, so induction is not well-defined. In other words, one cannot properly define a policy dependent on the history, which means that a given policy does not necessarily map to a unique outcome (see Simon and Stinchcombe (1989)). In particular, given the definition of the retraction cost in this model, if there is no restriction on the the firm could get away with switching an action without paying the retraction cost. For example, suppose that the firm's policy was to post 1 initially, but then post 0 if her posterior π_t^1 falls below $\hat{\pi}$, and to keep 0 posted for the remainder of the time if this occurs. If at time $t = \frac{1}{3}$ the firm's posterior falls below $\hat{\pi}$, the firm could simply have message 1 posted from $[0, \frac{1}{3})$ and message 0 posted from $[\frac{1}{3}, 1]$, and never actually pay the retraction cost as there is no time t at which $a_t \neq a_{t+}$.

To address this issue, I define a policy such that the firm only makes a decision at the times at which leads arrive. Specifically, I adapt the framework developed Khan and Stinchcombe (2015) (henceforth KS). KS develop a method for solving single-agent decision problems in infinite-horizon continuous-time models where the state changes according to a possibly nonstationary jump process. In their framework, at each time at which the state changes, the agent chooses a *plan*, which is a function that governs her behavior from that

[12]As neither the prior actions of the firm before this instant nor the exact history of leads received affect its expected payoffs, there is no reason for the firm to consider the entire history.

time until the next (random) time at which the state changes again. In KS, no aspects of the model depend on calendar time, so time resets at each jump time and the value function is not time dependent.

In the AIS, however, time itself is relevant to the firm's decision. This is due to the model being finite horizon and the option value of waiting. In order to allow for this type of non-stationarity, I use the method of Bäuerle and Rieder (2010). As the state does not change between the Poisson arrivals, the AIS is a *Piecewise Deterministic Markov Process*. This problem can thus be formulated recursively using a (discrete-time) Bellman equation. In the remainder of this section, I give a formal description of the firm's problem.

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\in[0,1]}, P)$ be a filtered probability space, meaning that Ω is the space containing all sample paths; $\{\mathcal{F}_t\}_{t\in[0,1]}$ is a filtration[13] defined on Ω with $\mathcal{F} \supseteq \cup_{t\in[0,1]}\mathcal{F}_t$, where the filtration is right-continuous and complete; and P is a probability measure on Ω. Let $\{T_k\}_{k\in\mathbb{Z}_+}$ be a sequence of random lead arrival times, where the first lead arrival time is always at the time the news event occurs, i.e. $T_0(\omega) = 0 \; \forall \omega \in \Omega$.[14] To simplify notation, I allow leads to arrive after $t = 1$, but these leads play no role in the firm's decision as no more customers arrive after $t = 1$. The probability that the time between lead k and lead $k + 1$ is less than $y \in \mathbb{R}_+$ is given by the exponential distribution with parameter λ, whose density I denote by $f(y)$.

At each random time T_k,[15] the firm chooses a *plan*,[16] which is a a left-continuous function $p_k : [T_k, \infty) \to A$ such that $p_k(\tau)$ assigns an action $a_\tau \in A$ to each time after T_k, *provided that*

[13] A *filtration* is a collection of sub-σ-algebras $\{\mathcal{F}_t\}_{t\in[0,1]}$ such that $\mathcal{F}_s < \mathcal{F}_t$ for all $s < t$.

[14] One can view this as the time at which the firm is initially informed about both possibilities and gets its prior.

[15] T_k is a random variable, but rather than specifying that something occurs at $T_k(\omega)$ for each ω, I follow the common convention of stating that something occurs at T_k.

[16] I define the plans on $[T_k, \infty)$ for notational simplicity. As the game ends at $t = 1$, the later part of the plan is irrelevant.

It is important to note that a *plan* is not a strategy. The strategy is the choice of plans from the set of all feasible plans.

another lead has not arrived.[17] In other words, a firm who chooses $p_k(\tau) = a_\tau$ is committing at time T_k to do that action at a future point—unless it receives additional information. There are many examples of plans that a firm could choose. For example, one plan may be to "post message 1 immediately, and then leave it posted until the next lead arrives," while another plan may be "do not post anything now, but post at time x if no new leads were received prior to then."[18] At the time the next lead arrives, the firm uses that new information to choose a different plan. For any action a_{T_k} and time T_k, the set of *feasible plans* $\mathcal{P}_k(a_{T_k})$ consists of the set of all plans such that $p_k(a_{T_k}) = a_{T_k}$ and there are at most finitely many times at which $p_k(a_t) \neq p_k(a_{t+})$.[19] Let $\mathcal{P}_k = \mathcal{P}_k(\varnothing) \cup \mathcal{P}_k(0) \cup \mathcal{P}_k(1)$.

For each $T_k < 1$, the firm's expected profit, excluding the retraction costs, from following a particular plan p_k is:

$$\hat{U}(T_k, p_k, \pi^1_{T_k}, a_{T_k}) = \int_{T_k}^{1} \left[\int_{T_k}^{s} \hat{u}(\pi^1_{T_k}, p_k(\tau)) d\tau \right] f(s) ds. \tag{2.3}$$

.

The inner integral is the expected profit (excluding retraction costs) that the firm will make from following plan p_k from T_k to s, and f is the density over timing of T_{k+1}, which determines how long the plan will be in place for.[20]

Let $\mathcal{T}_k(p_k)$ denote the set of times τ at which a firm would change its action when following plan p_k if it never receives another lead. The firm's expected retraction costs when following plan p_k can be represented by $\hat{R}(T_k, p_k, \pi^1_{T_k}, a_{T_k}) = \int_{T_k}^{1} \left[\sum_{t \in \mathcal{T}_k(p_k), t < s} r(p(t), p(t+)) \right] f(s) ds$. Given this, one can denote the firm's total expected profit from following plan p_k as:

[17] Note that the left-continuity condition ensures that there is precise last time at which an action a_t is taken, meaning that if the firm changes from action a to action $a' \neq a_t$, then $a_{t+} \neq a_t$.

[18] As I show later, given the option value of waiting, delay may be optimal.

[19] Note that, when $\rho > 0$, it would never be in the firm's interest to change plans infinitely many times regardless of this restriction, as its costs would be infinite!

[20] As f is Poisson, the lead arrivals are independent of each other, so this can without loss be written as the density rather than the conditional density.

$$\Pi(T_k, p_k, \pi^1_{T_k}, a_{T_k}) = \hat{U}(T_k, p_k, \pi^1_{T_k}, a_{T_k}) - \hat{R}(T_k, p_k, \pi^1_{T_k}, a_{T_k}). \tag{2.4}$$

A policy for the firm is a sequence $\{\sigma_k\}_{k\in\mathbb{Z}_+}$, with $\sigma_k : [0,1]\times[0,1]\times A \to \mathcal{P}$ for all $k > 0$, where $\sigma_k(T_k, \pi^1_{T_k}, a_{T_k}) \in \mathcal{F}_k(a_{T_k})$ assigns a plan at time T_k given the firm's posterior that the state is 1 and its action at time T_k. An optimal policy for the firm is a sequence $\{\sigma^*_k\}_{k\in\mathbb{Z}_+}$, such that, for each $k \in \mathbb{Z}_+$ and each $(\pi^1_{T_k}, a_{T_k}) \in [0,1] \times A$, if $T_k < 1$:

$$\sigma^*_k(T_k, \pi^1_{T_k}, a_{T_k}) \in \operatorname*{arg\,max}_{p_k\in\mathcal{P}(a_{T_k})} \Pi(T_k, p_k, \pi^1_{T_k}, a_{T_k}) + \mathbb{E}\sum_{\ell=k+1}^{\infty} \Pi(T_\ell, p_\ell, \pi^1_\ell, p_{\ell-1}(T_\ell - T_{\ell-1})). \tag{2.5}$$

Recall that, although the firm can move at any time, the firm only needs to plan its decision at the times at which leads arrive. The firm's value function at time 0 can be defined as follows:

$$V^*(0, \pi^1_0, a_0) = \sup_{\{p_k\}_{k\in\mathbb{Z}_+}} \left\{ \Pi(0, p_0, \pi^1_0, a_0) + \mathbb{E}\sum_{\ell=1}^{\infty} \Pi(T_\ell, p_\ell, \pi^1_\ell, a_\ell) \right\}, \tag{2.6}$$

subject to the constraint that $p_k \in \mathcal{P}(a_{Tk})$ given any time-t action a_{Tk}. This can be defined analogously starting from any lead arrival time. This is equivalent to a discrete-time infinite-horizon stochastic problem in which the firm chooses a plan in each period given the time, its posterior, and its current action, and where there is almost surely an $N \in \mathbb{N}$, such that, for all $k > N$, $\Pi(T_k, p_k, \pi^1_{T_k}, a_{T_k}) = 0$.[21] This value must be less than 1, which is the payoff from posting the correct article the entire time, so this infinite sum is well-defined. One can formulate this problem recursively using a Bellman equation. Let $C([0,1]^2 \times A)$ denote the vector space of continuous functions on $[0,1]^2 \times A$ endowed with the sup-norm.

[21]This occurs when $T_k \geq 1$

Lemma 1 (**Recursive form**). *The functional* $M : C([0,1]^2 \times A) \to C([0,1]^2 \times A)$*, where, for each* $k \in \mathbb{Z}_+$*:*

$$M(V(T_k, \pi^1_{T_k}, a_{T_k})) = \sup_{p_k \in \mathcal{P}(a_{T_k})} \left\{ \Pi(T_k, p_k, \pi^1_{T_k}, a_{T_k}) + \mathbb{E}V(T_{k+1}, \pi^1_{T_{k+1}}, a_{T_{k+1}}) \right\} \tag{2.7}$$

is a contraction mapping, and therefore M *has a unique fixed point* V^* *which is equal to* V^**.*

(All proofs are in the appendix.)

3 Analysis

3.1 Before Internet

As the firm only makes one decision in the BIS, this is a static decision problem. Given its posterior, the firm decides whether to post a message, and if so, which message. I assume that if the firm is indifferent between posting 1 and 0, it will post 1, and if it is indifferent between posting something and posting nothing, it will post something. Let $\pi^* = \frac{\kappa-1}{\kappa}$, and $\pi^{**} = \max\{\pi^*, \frac{1}{2}\}$.

Proposition 1 (**BIS editorial standard**). *If the firm's posterior that the event is* 1 *is not equal to* $\frac{1}{2}$*, the firm will post message* 1 *iff* $\pi^1(d) \geq \pi^{**}$ *. If* $\pi^1(d) < \pi^{**}$*, it will post message* 0 *if* $\pi^0(d) \geq \pi^{**}$ *and nothing otherwise.*

The cutoff π^{**} defines the firm's editorial standard in the BIS. Proposition 1 implies that if the expected cost of posting a false article is sufficiently low relative to its revenue, the firm will choose to post the article. In fact, as I state in the following corollary, if the cost of posting an incorrect article is less than or equal to twice the benefit of posting, the firm will always post.

Corollary 1. *The probability with which the firm posts any message is:*

1. *strictly decreasing in the cost* κ *for all* $\kappa > 2$,

2. *constant in the cost for all* $\kappa \leq 2$.

This follows from the fact that, if $\kappa > 2$, then $\pi^{**} = \frac{\kappa-1}{\kappa}$, so the probability that either message is posted is strictly decreasing in κ, whereas if $\kappa \leq 2$, $\pi^{**} = \frac{1}{2}$, so the firm will always post a message and simply post the one it believes is most likely. The intuition behind this is displayed in figure 1.

[Figure 1 here]

The following proposition provides insight as to the effect of having a more correct prior on the probability of posting an incorrect message in the BIS.

Proposition 2 (Effect of firm's prior). *In the BIS, the probability with which a firm will post the incorrect message when the event is* x *is weakly decreasing in the firm's prior* π_0^x.

This result is intuitive. For any given set of leads received, a firm who initially believes that the state is more likely to be x would be less likely to post x'. Therefore, a firm whose posterior places higher probability on the correct state will have a lower *ex-ante* probability of posting an incorrect message.

One should note that the symmetry comes from the fact that the firm's expected payoffs only depend on the firm's posterior. Later, I discuss an extension in which this assumption is relaxed.

3.2 After Internet

Before solving for the editorial standard in the general case, first consider the case in which $\rho = 0$, meaning that retraction is costless. This implies that the firm's decision in the (stochastic) interval between T_k and T_{k+1} has no effect on its future decisions, so the firm will simply choose at each lead arrival time to post whichever message maximizes its expected flow payoff. I state this formally below:

Lemma 2 (**Costless retraction implies BIS standard**). *If retraction is costless, then if the firm's posterior that the event is* 1 *is not equal to* $\frac{1}{2}$*, the firm will post message* 1 *at time* t *iff* $\pi_t^1 \geq \pi^{**}$ *. If* $\pi_t^1 < \pi^{**}$*, it will post message* 0 *if* $\pi_t^0 \geq \pi^{**}$ *and nothing otherwise.*

This implies that the firm uses the same editorial standard as in the BIS. In this situation, the firm's current decision has no effect on its future payoffs, so maximizing its current payoff is the sensible approach. In other words, the problem at each instant can be treated as essentially static.

Moreover, the BIS editorial standard π^{**} plays an important role when $\rho > 0$ as well. In particular, the following proposition holds:

Proposition 3 (**AIS posting decision is bounded by BIS standard**). *If retraction is costly, then the firm will not initially post message* x *at any time* t *at which* $\pi_t^x < \pi^{**}$ *, and it will not retract* x *at any time* t *at which* $\pi_t^x \geq \pi^{**}$.

The intuition behind this result is that if the firm could earn a higher expected flow payoff by not posting or posting a different message, it would not be optimal for the firm to post x even if it did not need to be concerned about retraction. If the firm already has x posted and is earning a higher expected flow profit from having it posted than from its other options, it certainly would not choose to retract. This implies that the range of beliefs over which the firm has no message posted is weakly larger in the AIS than in the BIS.

This implies the following corollary:

Corollary 2. *At any lead arrival time, if the firm does not have an article posted, it will choose a plan in which it changes its post at most once.*

This is true because the firm's posterior does not change between lead arrivals. Therefore, if π_t^x is greater than (less than) π^{**} at the lead arrival time, it will remain above or below the threshold. It is therefore not possible for multiple changes within a given plan to be optimal.

The willingness of the firm to retract depends on the per-consumer cost of having an incorrect article posted, κ, and the retraction cost ρ. Specifically, if $\rho > \kappa$, the firm would never choose to Moreover, as I show below, there exists a time after which the firm will not retract.

Lemma 3 (**Latest retraction time**). *If retraction is costly, there exists a time* $\hat{t} < 1$ *after which the firm will not retract.*

The above analysis hints at the fact that the firm will use cutoff functions that determine its posting decisions. The first function, the editorial standard, $\bar{\pi} : [0, 1] \to [0, 1]$, determines the threshold at which, if $\pi_t^x > \bar{\pi}(t)$, the firm would prefer initially posting message x to posting nothing or $x' \neq x$. The second function, the retraction cutoff, $\underline{\pi} : [0, 1] \to [0, 1]$, determines the threshold at which, if $a_t = x$ and $\pi_t^x < \underline{\pi}(t)$, the firm at time t would prefer to retract. In the proposition below, I show that these cutoffs uniquely determine the optimal policy for the BIS. Prior to stating Proposition 1.4 about the existence and uniqueness of the optimal policy, I state the the intuitive result that if the firm weakly prefers posting at time t to not posting, it strictly prefers doing so with any higher posterior.

Lemma 4 (**Threshold structure**). *Let* $a_t = \varnothing$*. If the firm would post article* x *at time* t *with posterior* π_t^x *that the state is* x*, it would strictly prefer to post with posterior* $\hat{\pi}_t^x > \pi_t^x$.

Given this, it is clear that the optimal policy will have a threshold stucture. As I show below, the thresholds are unique and are time-dependent.

Proposition 4 (**Existence and uniqueness of the optimal policy**). *There exists a unique editorial standard* $\bar{\pi}$ *and retraction threshold* $\underline{\pi}$*, such that the unique optimal policy* *for the firm is to choose a plan at each T_k in which it posts at the earliest t such that $\pi_t^x \geq \bar{\pi}(t)$ for some x, and leaves that article posted until the first time t such that $\pi_t^x < \underline{\pi}(t)$.*

How does the editorial standard and the retraction cutoff vary over time? In the proposition below, I show formally that both cutoffs decrease over time. More precisely, the editorial standard strictly decreases until it reaches the BIS cutoff π^{**}, and the retraction cutoff decreases until it reaches 0.

Proposition 5 (**Editorial standard and retraction cutoff are decreasing over time**). *If retraction is costly, there exist times t^{**} and $\hat{t}$ such that such that the editorial standard is strictly decreasing for all $t < t^{**}$ and equal to π^{**} for all $t \geq t^{**}$ and the retraction cutoff is strictly decreasing for all $t < \hat{t}$ and equal to 0 for all $t \geq \hat{t}$.*

The intuition for this result is that if the firm chooses to initially post at time t it will get a flow payoff, but may receive contradictory information later that would make it wish it had not posted. If $t' > t$, then the probability with which it will receive a contradictory lead is lower, so its threshold for choosing to post is lower. This implies that the firm may choose to delay in order to wait for more information, but if it does not receive additional information it may eventually post using the information that it has. In other words, the timing incentives motivate the firm to be more cautious at first. This does not necessarily mean, however, that a firm who delays before posting has received more information, rather it may recognize that it is now less likely to receive contradictory information later that would make it wish it could retract.

When is t^{**}? Recall that π^* is the value at which the firm is indifferent in the BIS at between posting message x and posting nothing. This is also the firm's indifference condition at $t = 1$, as there is no longer an option value of waiting. Therefore, if $\kappa \geq 2$, so $\pi^{**} = \pi^*$, then $t^{**} = 1$. Otherwise, $\pi^* < \pi^{**} = \frac{1}{2}$, so t^{**} at the time at which $\bar{\pi}(t) = \frac{1}{2}$, which, given the option value of waiting, is at a time $t^{**} < 1$.

Note that the following corollary follows from Proposition 5.

Corollary 3. *If the firm is following the optimal policy and has a message x posted at time t, it would not retract x if it receives no message or a message equal to x.*

This is true because, if it receives no message, its posterior stays the same, so if its posterior on x is not below $\underline{\pi}(T_k)$ at the time it received the lead, it will not be below it at time $t > T_k$ if it still has not received a lead or received one that increases its posterior.

Note that increasing the editorial standard decreases the probability of posting an incorrect message. However, as the firm is more cautious, it also decreases the probability of posting the correct message. I state this formally in the corollary below.

Corollary 4. *The probability of posting any message at time t is decreasing in $\bar{\pi}(t)$.*

The intuition for this is the same as for Corollary 1 because increasing $\bar{\pi}(t)$ has the same effect as increasing the BIS cutoff in that the firm would have to be more sure of the state in order to post. What this means is that the firm uses stricter standards in the AIS, so this may have the consequence of making it less likely for the firm to release incorrect news.[22] The flip side of this, of course, is that the firm is also less likely to release correct news in a timely manner.

Based on Proposition 5, one can derive the following comparative static result about the effect of the retraction cost on the posting decision. Let ρ' and ρ be different retraction costs, and let $\bar{\pi}'$ and $\underline{\pi}'$ be the editorial standard and the retraction cutoff respectively under ρ'.

Proposition 6 (Retraction-cost comparative statics). *Let $\rho' > \rho$, then then $\overline{\pi}' \geq \overline{\pi}$ and $\underline{\pi}' \leq \underline{\pi}$, with these equalities being strict whenever $\overline{\pi}' > \pi^{**}$ and $\underline{\pi}' > 0$.*

This implies that a higher retraction cost makes a firm more hesitant to initially post, but also more hesitant to retract. One implication of this is that he effect on society of

[22] See the next section for the reason for which this may not result in better news.

having higher penalties for retraction costs is ambiguous. On one hand, it makes the firm hesitate more before posting, while on the other hand, it makes the firm more likely to leave incorrect information posted.

3.3 Comparison

Using the results obtained in section 3.1 and 3.2, it is possible to consider the effect of timing on the editorial standard used by the media.

Note that, when retraction is costless, if the firm would post message x in the BIS, it would post weakly earlier in the AIS. This is true because the firm would choose to post at the first time that $\pi_t^x \geq \pi^{**}$, which is weakly earlier than $\bar{t}$. This implies that the firm will post with weakly less information in the AIS and thus rely more on its prior. The implication of this is that, absent retraction concerns, the internet does lead firms to release information for which they have less support.

When retraction is costly, however, the firm may actually not post at time $\bar{t}$ in the AIS when it would have posted in the BIS. If $\bar{t} < t^{**}$ and $\pi_{\bar{t}}^1 \in (\pi^{**}, \bar{\pi}(\bar{t}))$, the firm would be sure enough of the state to post in the BIS and not in the AIS. Thus, the option value of waiting and retraction cost serve as deterrents against posting information about which the firm is less sure.

Does this mean that, contrary to popular concern, the 24-hour news cycle actually results in less incorrect news being released? Not necessarily. Note that, if the firm's prior about either state is sufficiently high, it will even post with no evidence. This could result in errors like the Fox News misreport of the Quebec City shooting mentioned in the introduction. This implies that the effect of the internet on the correctness of the news depends on the strength of the news firm's prior, which may have implications for consumer welfare. A full comparison would need to take into account consumers preferences for reading news on their own schedules. Would consumers be willing to have news that has a higher probability of being incorrect in order to get to read it sooner? If the retraction cost is zero, this comparison is more ambiguous, because the firm will update its posting whenever needed, so although the initial posting may be done with less information, eventually the posting will be done with more information. A full analysis of this sort is beyond the scope of this article.

4 Extension: Payoffs vary in the message sent

Of course, it is possible for a firm's payoffs to depend not only on whether the message is correct, but on the message itself. In this section, I consider an extension in which the payoffs can vary based on the message sent. In particular, I consider the following two possible scenarios:

1. The firm's revenue from posting message 1 is equal to $b > 1$, whereas it's payoff from posting message 0 remains 1.

2. The firm's cost from getting message 0 incorrect is equal to $\kappa' > \kappa$, whereas it's cost from posting message 1 remains at κ.

In the first case, the firm earns more revenue when a certain message is posted. This could, for example, be due to its customers having a certain political bias, so a larger measure of consumers would want to view an article expressing a certain viewpoint. In the second case, the cost of getting the article incorrect depends on the message reported. This could also be due to political bias, but could also, for example, be due to an incorrect report in one direction being more likely to lead to a lawsuit or reputational concerns.

First suppose that the firm's payoff depends on the message that is sent. There are two straightforward ways through which this could happen: The firm's benefit from posting could vary depending on the message, or the firm's cost of a making a mistake could vary depending on the message. First suppose without loss of generality that the firm's benefit $b > 1$ from posting message 1.

The following result describes the BIS posting rule in this case.

Proposition 7 (Message-dependent firm revenue)**.** *In the BIS with message-dependent firm revenue:*

- *If $\kappa > 1 + b$, the firm will post message 1 if $\pi^1 \geq \frac{\kappa - b}{\kappa}$, post message 0 if $\pi^0 \geq \pi^*$, and post no message otherwise.*
- *Otherwise, the firm will post message 1 if $\pi^1 \geq \frac{1}{2} - \frac{b-1}{2\kappa}$ and post message 0 otherwise*

Given that the firm's payoff from posting message 0 has not changed, the cutoff at which the firm is indifferent between posting message 0 and not posting remains the same, π^*. Unsurprisingly, the cutoff at which the firm is indifferent between posting message 1 and not posting has shifted downwards, as its payoff from not posting is higher. Define this cutoff, $\frac{\kappa - b}{\kappa}$, as π^{b*}.

Given this, there are a wider range of parameters over which the firm would prefer having either message posted to no message posted, so the firm will always post a message when $\kappa \geq 1 + b > 2$.

As in the baseline case considered before, if retraction is costless, the firm will follow the same posting policy in the AIS as in the BIS. The case where the costs differ works similarly, in the sense that the firm will be more reluctant to post the story in which the cost of being incorrect is higher. Let $\kappa^0 \geq \kappa$ be the cost for incorrectly reporting message 0. In this case the result is as follows:

Proposition 8 (Message-dependent firm cost)**.** *In the BIS with message-dependent firm cost:*

- *If* $\frac{\kappa-1}{\kappa}+\frac{\kappa^0-1}{\kappa^0}\leq 1$*, the firm will post message* 1 *if* $\pi^1\geq\frac{\kappa-1}{\kappa}$*, post message* 0 *if* $\pi^0\geq\frac{\kappa^0-1}{\kappa^0}$*, and post no message otherwise.*

- *If* $\frac{\kappa-1}{\kappa}+\frac{\kappa^0-1}{\kappa^0}>1$*, the firm will post message* 1 *if* $\pi^1\geq\frac{\kappa}{\kappa+\kappa^0}$*, and post message* 0 *otherwise.*

The above results are for the BIS. Note, however, that in both situations, the effect of the payoff variation is to change the cutoff at which the firm would be indifferent for a given message, when there is no possibility of (or cost of) retraction. Given this, as before, if the retraction cost in zero there the firm uses the same cutoff in the AIS as in the BIS, and the other results for the AIS should remain qualitatively the same.

However, in the AIS, there is another possibility to consider, namely that the retraction cost could be message-dependent. There could be multiple reasons why this could occur. For example, it could be the case that consumers would be more upset if the firm retracted news that supported their political opinions. Alternatively, in the case in which an event either occurred or did not, retracting a claim that it did happen may be more problematic than retracting a claim that it did not. Let ρ^1 be the cost of retracting article 1, and ρ^0 be the cost of retracting article 0, and assume that $\rho^1>\rho^0$. Let $\bar{\pi}^k$ and $\underline{\pi}^k$ be their respective editorial standards and retraction thresholds. Then from Proposition 6, we know that $\overline{\pi}^1\geq\overline{\pi}^0$ and $\underline{\pi}^1\leq\underline{\pi}^0$, with these equalities being strict whenever $\overline{\pi}^1>\pi^{**}$ and $\underline{\pi}^1\geq 0$. This means that the firm would be more hesitant to initially post message 1, but also less likely to retract it.

5 Conclusion

The above analysis shows that having a "24-hour news cycle" as opposed to fixed posting times can have a significant effect on the content produced by the media. Although allowing for constant posting and updating of news may lead the firm to post earlier, after receiving fewer leads, if retraction is costly, it might actually lead a firm to delay posting slightly in order to receive more information. This implies that the ability to post any time may provide the incentive for a firm to wait a little longer for additional information before releasing news, thereby improving quality.

The implication of this depends strongly on the prior belief of the firm. If the firm has a moderate prior, the internet is likely to lead to it being more cautious and less likely to release a news story that would be later shown to be false. However, if it has a strong prior, it would be inclined to release news sooner now that it has the option to do so. While this article does not model competition, it is likely that this could lead to increased polarization of the news media. If firms with strong priors are more likely to release news—assuming that it agrees with their priors—then more polarized news would likely result. In the past all firms *had* to wait for more information, but now it is a choice.

Of course, the internet has had more effects on news production than just timing. In particular, the internet has increased the number of news providers. Although competition likely would change the editorial standards used by firms, the key factor leading to the time-dependent editorial standard—the option value of waiting—would still apply if there is competition. Although the effect of competition on the news produced by heterogeneous firms is interesting, it is essentially a separate issue from the timing aspect and is therefore beyond the scope of this article. It may also be useful to consider what would happen if firm's had the option to improve the accuracy of their reports by verifying the information that they receive. Verifying this information would be costly and take time, but it eliminates the risk of posting false information.

Figure

Figure 1: Firm's Posting Decision in the BIS

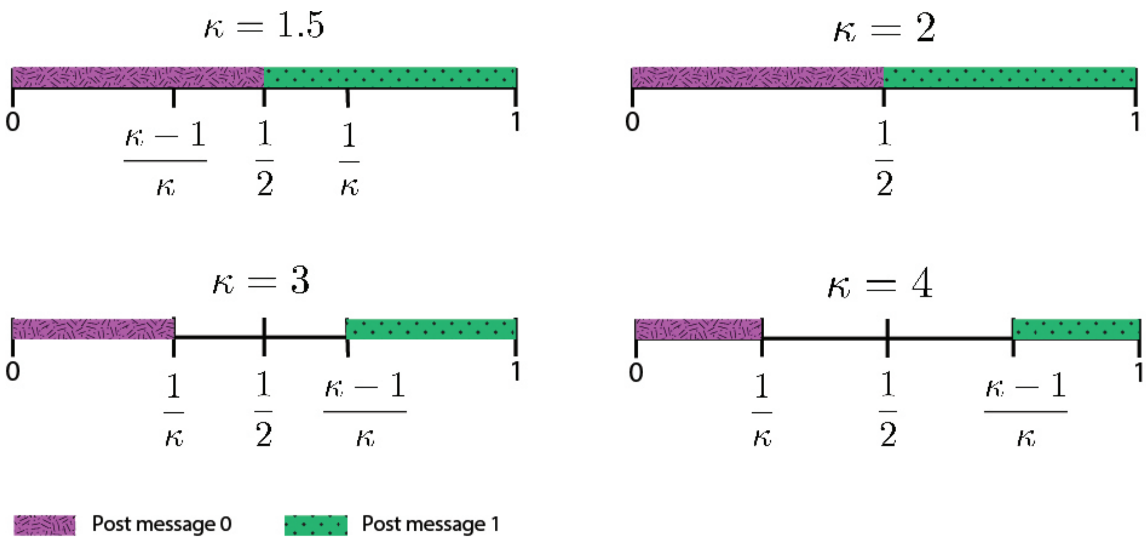

A Proofs

Proof of Lemma 1

Lemma (Restatement). *The functional $M : C([0,1]^2 \times A) \to C([0,1]^2 \times A)$, where, for each $k \in \mathbb{Z}_+$:*

$$M(V(T_k, \pi^1_{T_k}, a_{T_k})) = \sup_{p_k \in \mathcal{P}(a_{T_k})} \left\{ \Pi(T_k, p_k, \pi^1_{T_k}, a_{T_k}) + \mathbb{E}V(T_{k+1}, \pi^1_{T_{k+1}}, a_{T_{k+1}}) \right\} \qquad \text{(A.1)}$$

is a contraction mapping, and therefore M has a unique fixed point V^* which is equal to V^*.

Proof. As $C([0,1]^2 \times A)$ is a Banach space of bounded functions, to show that this is a contraction mapping, it is sufficient to show that the value function satisfies Blackwell's sufficient conditions of monotonicity and discounting.

1. Let $W(T_k, \pi^1_{T_k}, a_{T_k}) < V(T_k, \pi^1_{T_k}, a_{T_k})$ for all $(T_k, \pi^1_{T_k}, a_{T_k}) \in [0,1]^2 \times A$.

   Then $\sup_{p_k \in \mathcal{P}(a_{T_k})} \Pi(T_k, p_k, \pi^1_{T_k}, a_{T_k}) + \mathbb{E}W(T_{k+1}, \pi^1_{T_{k+1}}, a_{T_{k+1}}) < \sup_{p_k \in \mathcal{P}(a_{T_k})} \Pi(T_k, p_k, \pi^1_{T_k}, a_{T_k}) + \mathbb{E}V(T_{k+1}, \pi^1_{T_{k+1}}, a_{T_{k+1}})$, so monotonicity is satisfied.

2. For any constant c, $(V+c)(T_k, \pi^1_{T_k}, a_{T_k}) = \sup_{p_k \in \mathcal{P}(a_{T_k})} \Pi(T_k, p_k, \pi^1_{T_k}, a_{T_k}) + \mathbb{E}[V(T_{k+1}, \pi^1_{T_{k+1}}, a_{T_{k+1}}) + c] = \sup_{p_k \in \mathcal{P}(a_{T_k})} \Pi(T_k, p_k, \pi^1_{T_k}, a_{T_k}) + \mathbb{E}V(T_{k+1}, \pi^1_{T_{k+1}}, a_{T_{k+1}}) + \int_{T_k}^1 cf(s)ds$. As the exponential distribution has support on $(0, \infty)$, $c\int_{T_k}^1 f(s)ds \leq \beta c$ for some $\beta < 1$.

Therefore, this is a contraction mapping. The fact that this is equivalent to the sequential definition of V^* follows immediately from the fact that $V^*(T_k, \pi^1_{T_k}, a_{T_k})$ is defined as the sum of all of the disjoint (future) stochastic intervals. □

Proof of Proposition 1

Proposition (Restatement)**.** *If the firm's posterior that the event is* 1 *is not equal to* $\frac{1}{2}$*, the firm will post message* 1 *iff* $\pi^1(d) \geq \pi^{**}$ *. If* $\pi^1(d) < \pi^{**}$*, it will post message* 0 *if* $\pi^0(d) \geq \pi^{**}$ *and nothing otherwise.*

Proof. Let $\pi^x(d)$ be the firm's posterior that the state is x after receiving lead difference d. The revenue from posting x equals one and the expected cost equals $(1 - \pi^x(d))\kappa$, so the revenue exceeds the cost when $\pi^x(d) \geq \frac{\kappa - 1}{\kappa} = \pi^*$. Therefore, it will post a message if $\pi^x(d) \geq \pi^*$ for some x. If $\pi^1(d) \geq \pi^0(d) \geq \pi^*$, then its expected profit from posting 1 is higher than its expected profit from posting 0 so the firm will post 1, and the reverse is true when $\pi^0(d) > \pi^1(d)$. □

Proof of Proposition 2

Proposition (Restatement)**.** *In the BIS, the probability with which a firm will post the incorrect message when the event is* x *is decreasing in the firm's prior* π^x_0*.*

Proof. Assume that the event is 1. All of the results hold symmetrically for the case where the event is 0.

For the firm to post message 0 after receiving n_1 type-1 leads and n_0 type-0 leads, it must be the case that, for any d, $\frac{q^d \pi_0^1}{q^d \pi_0^1 + (1-q)^d (1-\pi_0^1)} \leq \min\{\frac{1}{\kappa}, \frac{1}{2}\}$. As $\frac{q^d \pi_0^1}{q^d \pi_0^1 + (1-q)^d (1-\pi_0^1)}$ is increasing in π_0^1 for all d, the probability with which this posterior is less than $\min\{\frac{1}{\kappa}, \frac{1}{2}\}$ is decreasing in π_0^1. □

Proof of Lemma 2

Lemma (Restatement)**.** *If retraction is costless, then if the firm's posterior that the event is* 1 *is not equal to* $\frac{1}{2}$*, the firm will post message* 1 *at time* t *iff* $\pi_t^1 \geq \pi^{**}$ *. If* $\pi_t^1 < \pi^{**}$*, it will post message* 0 *if* $\pi_t^0 \geq \pi^{**}$ *and nothing otherwise.*

Proof. The only restriction that the plan chosen at T_k places on the plan chosen at T_{k+1} is that $p_{k+1}(T_{k+1}) = p_k(T_{k+1})$. Therefore, given that retraction is costless, the action within each stochastic interval $[[T_k, T_{k+1}[[$, where $[[T_k, T_{k+1}[[:= \{(\omega, t) | T_k(\omega) \leq t < T_{k+1}(\omega)\}$, does not affect the firm's profit in the next stochastic interval. Therefore, the firm will post x iff posting maximizes its flow profit, which occurs whenever $\pi_t^x \geq \pi^{**}$, and as before will post message 1 if indifferent. □

Proof of Proposition 3

Proposition (Restatement)**.** *If retraction is costly, then the firm will not initially post message* x *at any time* t *at which* $\pi_t^x < \pi^{**}$ *, and it will not retract* x *at any time* t *at which* $\pi_t^x \geq \pi^{**}$.

Proof. If $\pi_t^x < \pi^{**}$, then $\pi_t^x < \pi^*$ or $\pi_t^x < \frac{1}{2}$. If $\pi_t^x < \pi^*$, the firm would get a negative flow payoff from posting x, so it would prefer not to post. If $\pi_t^x < \frac{1}{2}$, the firm would prefer to post the other message so it would not post x.

If the firm already has message x posted at time t, and $\pi_t^x \geq \pi^{**}$, its flow profit from posting x is higher than its flow profit from any other action, so it would not want to pay a retraction cost to switch. □

Proof of Lemma 3

Lemma (Restatement)**.** *If retraction is costly, there exists a time $\hat{t} < 1$ after which the firm will not retract.*

Proof. If $\kappa \leq \rho$, the firm will never retract as the retraction cost is greater than the cost of leaving the wrong article posted for the entire time, so $\hat{t} = 0$. Otherwise, suppose the firm has message x posted at time t. If it chooses not to retract at time t, the maximal cost it could incur is $\int_t^1 \kappa ds$. As $t \to 1$, $\int_t^1 \kappa ds \to 0$, which implies that there exists a $\hat{t} < 1$ such that $\int_t^1 \kappa ds = \rho$, and after this time the firm would choose not to retract. □

Proof of Lemma 4

Lemma (Restatement)**.** *Let $a_t = \varnothing$. If the firm would post article x at time t with posterior π_t^x that the state is x, it would strictly prefer to post with posterior $\hat{\pi}_t^x > \pi_t^x$.*

Proof. Suppose without loss of generality that $\pi_t^1 \geq \frac{1}{2}$. This implies that the firm would not post message 0. From Corollary 1.3, we know that if the firm posts message 1, it would not retract it unless it receives another lead. Therefore, for the firm to post an message 1 at time t, it must be the case that, for all $t' > t$:

$$\int_t^1 \left[\int_t^s (1 - (1 - \pi_t^1)\kappa) d\tau + \mathbb{E}V(s, \pi_s^1, 1 | \pi_t^1) \right] f(s) ds \geq \tag{A.2}$$
$$\int_t^{t'} \mathbb{E}V(s, \pi_s^1, \varnothing | \pi_t^1) f(s) ds + \int_{t'}^1 \left[\int_{t'}^s (1 - (1 - \pi_t^1)\kappa) d\tau + \mathbb{E}V(s, \pi_s^1, 1|, \pi_t^1) \right] f(s) ds,$$

which is equivalent to the requirement that:

$$\int_t^{t'} \left[\int_t^s (1-(1-\pi_t^1)\kappa) d\tau \right] f(s)ds \geq \int_t^{t'} \left[\mathbb{E}V(s,\pi_s^1,\varnothing|\pi_t^1) - \mathbb{E}V(s,\pi_s^1,1|\pi_t^1) \right] f(s)ds. \quad \text{(A.3)}$$

If the firm would strictly prefer to post message 1, then this inequality is strict.

As $1-(1-\pi_t^1)\kappa$ is strictly increasing in π_t^1, it is sufficient to show that the right-hand side of this equation is non-increasing in π_t^1. Note that $V(s,\pi_s^1,\varnothing|\pi_t^1) > V(s,\pi_s^1,1|\pi_t^1)$ only when the firm would prefer not to have article 1 posted at time s. Otherwise they are equal. As the instantaneous payoff from having the message posted at time s is increasing in π_s^1, the the only way that $\int_t^{t'} [\mathbb{E}V(s,\pi_s^1,\varnothing|\pi_t^1) - \mathbb{E}V(s,\pi_s^1,1|\pi_t^1)] f(s)ds$ could increase in π_t^1 would be if increasing π_t^1 increase the probability that π_s^1 would fall below π^{**} at some point in the future. As the expectation of π_s^1 at any time $s > t$ is increasing in π_t^1, this cannot occur. The analogous argument holds when the firm is considering posting message 0. □

Proof of Proposition 4

Proposition (Restatement)**.** *There exists a unique editorial standard $\bar{\pi}$ and retraction threshold $\underline{\pi}$, such that the unique optimal policy for the firm is to choose a plan at each T_k in which it posts at the earliest t such that $\pi_t^x \geq \bar{\pi}(t)$ for some x, and leaves that article posted until the first time t such that $\pi_t^x < \underline{\pi}(t)$.*

Proof. Let T_k be such that that $a_{T_k} = \varnothing$. If the firm posts at any time t, we know from Corollary 2 that it will leave the article posted at least until the next lead arrives.

Assume without loss of generality that $\pi_{T_k}^1 \geq \frac{1}{2}$. At any T_k, the optimal p_k therefore involves posting nothing on $[T_k, \tilde{t}]$ and posting message 1 on $(\tilde{t}, 1]$ for some $\tilde{t} \in [T_k 1]$, with $\tilde{t} = 1$ if the firm would never post message 1 with that posterior, where time $\tilde{t}$ is such that:

$$\tilde{t} \in \arg\max_{T_k \le t \le 1} \int_{T_k}^{t} \mathbb{E}V(s, \pi_0^1, \varnothing | \pi_{T_k}^1) f(s) ds + \int_t^1 \left[\int_t^s (1 - (1 - \pi_{T_k}^1)\kappa) d\tau + \mathbb{E}V(s, \pi_s^1, 1 | \pi_{T_k}^1) \right] f(s) ds.$$

Let $\bar{\pi}(T_k)$ be the smallest $\pi_{T_k}^1$ such that the firm would post message 1 immediately at T_k. As Proposition 3 shows, the firm would not post if $\pi_{T_k}^1 < \pi^{**}$, and the firm would certainly post if $\pi_{T_k}^1 = 1$. As the expected profit from posting at time T_k varies continuously in $\pi_{T_k}^1$, this $\bar{\pi}(T_k)$ exists. Lemma 4 further implies that, if for a given $\pi_{T_k}^1$ the firm would post at time $\tilde{t}$ with $\pi_{T_k}^1$, it would do so for all $\hat{\pi}_{T_k}^1 > \pi_{T_k}^1$, so this implies that $\bar{\pi}(T_k)$ forms a unique cutoff.

If $a_t = 1$, the firm would retract message 1 if its benefit from retracting is greater than the retraction cost.

Note that, as the firm is considering retraction, we know from Proposition 3 that $\pi_t^1 < \pi^{**}$. If $\pi_t^1 \le 1 - \pi^{**}$, meaning that $\pi_t^0 \ge \pi^{**}$, then:

$$V(t, \pi_t^1, \varnothing | \pi_t^1) = \max_{t' \in [t,1]} \int_t^{t'} \mathbb{E}V(s, \pi_s^1, \varnothing | \pi_t^1) f(s) ds + \int_{t'}^1 [\int_{t'}^s (1 - \pi_t^1 \kappa) d\tau + \mathbb{E}V(s, \pi_s^1, 0 | \pi_{T_k}^1)] f(s) ds,$$

which is decreasing in π_t^1. As $V(t, \pi_t^1, 1 | \pi_t^1)$ is increasing in π_t^1, $V(t, \pi_t^1, \varnothing | \pi_t^1) - V(t, \pi_t^1, 1 | \pi_t^1)$ is decreasing in this region. If $\pi_t^1 \in (1 - \pi^{**}, \pi^{**})$, then if the firm retracts, it will for sure not post another message until a lead arrives, and $V(t, \pi_t^1, \varnothing | \pi_t^1) = \int_t^1 \mathbb{E}V(s, \pi_s^1, \varnothing | \pi_t^1) f(s) ds$, whereas

$$\begin{aligned} V(t, \pi_t^1, 1 | \pi_t^1) = \max_{t' \in [t,1]} & \int_t^{t'} [(1 - (1 - \pi_{T_k}^1)\kappa) d\tau + \mathbb{E}V(s, \pi_s^1, 1 | \pi_{T_k}^1)] f(s) ds \\ & + \int_{t'}^1 \mathbb{E}V(s, \pi_s^1, \varnothing | \pi_{T_k}^1)] f(s) ds, \end{aligned}$$

note that as the firm is earning a negative flow payoff from having the message posted, learns no information between lead arrivals, and could costlessly repost article 1 if it learns otherwise, if it is going to retract, it would be better off doing so immediately at time t. Therefore, $V(t, \pi_t^1, \varnothing|\pi_t^1) - V(t, \pi_t^1, 1|\pi_t^1) = \int_t^1 (\mathbb{E}V(s, \pi_s^1, \varnothing|\pi_t^1) - [(1-(1-\pi_{T_k}^1)\kappa)d\tau + \mathbb{E}V(s, \pi_s^1, 1|\pi_{T_k}^1)])f(s)ds$, which is decreasing in π_t^1

Therefore, there is a well-defined unique cutoff such that the firm would retract iff $\pi_t^1 < \underline{\pi}(t)$, and the optimal policy can be uniquely described by the proposition, under the continued assumption that the firm posts message 1 if indifferent. □

Proof of Proposition 5

Proposition (Restatement)**.** *If retraction is costly, there exist times t^{**} and $\hat{t}$ such that such that the editorial standard is strictly decreasing for all $t < t^{**}$ and equal to π^{**} for all $t \geq t^{**}$ and the retraction cutoff is strictly decreasing for all $t < \hat{t}$ and equal to* 0 *for all $t \geq \hat{t}$.*

Proof. First I consider the editorial standard, and assume without loss of generality that $\pi_0^1 \geq \frac{1}{2}$. Note that, given proposition 3, the firm will never post message 1 (in the absence of more leads) unless $\pi_0^1 \geq \pi^{**}$, so assume that $\pi_0^1 \geq \pi^{**}$. Without loss of generality, consider what the firm would do at a lead arrival time T_k.

For any given time T_k, the editorial standard $\bar{\pi}(T_k)$ satisfies the following indifference condition:

$$V(T_k, \bar{\pi}(T_k), \varnothing) = V(T_k, \bar{\pi}(T_k), 1),$$

meaning that the firm's optimal posting time is that moment. Given that Proposition 4 states that the optimal posting time is unique, the firm solves:

$$T_k = \arg\max_{t\in[T_k,1]} \int_{T_k}^{t} \mathbb{E}[V(s,\pi_s^1,\varnothing)|\bar{\pi}(T_k)]f(s)ds$$
$$+ \int_t^1 \left[\int_t^s (1-(1-\bar{\pi}(T_k))\kappa)d\tau + \mathbb{E}[V(s,\pi_s^1,1)|\bar{\pi}(T_k)]\right] f(s)ds.$$

Given Proposition 3, for any $\pi_{T_k}^1 \geq \pi^{**}$,$V(T_k,\pi_{T_k}^1,1)$ is equal to:

$$\int_{T_k}^1 \left[\int_{T_k}^s (1-(1-\pi_{T_k}^1)\kappa)d\tau + \mathbb{E}[V(s,\pi_s^1,1|\pi_{T_k}^1)]\right] f(s)ds.$$

First note that, since the firm would certainly prefer posting at time 0 if $\pi_0^1 = 1$, by continuity there must exist a π_0^1 such that the firm is indifferent about whether to post immediately at time 0. Call this value $\overline{\pi}(0)$. This implies that for all $\pi_0^{1\prime} > \pi_0^1$, the firm will prefer to post immediately at time 0. At $t = 1$, because retraction is no longer a consideration, this is equivalent to the case where there is no retraction cost and $\overline{\pi}(t) = \pi^{**}$.

Note also that $V(\cdot,\pi_{T_k}^1,\varnothing) - V(\cdot,\pi_{T_k}^1,1)$ is weakly decreasing in T_k. This is true because these expressions only differ over the period in which a message is not posted, and this difference is always weakly decreasing in time. This implies that it must be the case that, if $T_k < T_k'$, then $\bar{\pi}(T_k) \geq \bar{\pi}(T_k')$. This implies that if $\overline{\pi}(T_k) = \pi^{**}$ for some T_k, then for all $T_k' > T_k$, $\overline{\pi}(T_k) = \pi^{**}$.

Now consider the times in-between. Over this interval, the $\overline{\pi}(\cdot)$ is strictly decreasing. The instantaneous benefit that the consumer gets from posting at time T_k is $1-(1-\pi_{T_k}^1)\kappa$, while the instantaneous cost is $\lambda(\mathbb{E}[V(T_{k+1},\pi_{T_k+1}^1,\varnothing|\pi_{T_k}^1)] - \mathbb{E}[V(T_{k+1},\pi_{k+1}^1,1|\pi_{T_k}^1)])$. As the instantaneous benefit is independent of T_k, it is sufficient to show that $\mathbb{E}[V(T_k,\pi_s^1,\varnothing|\pi_{T_k}^1)] - \mathbb{E}[V(T_k,\pi_s^1,1|\pi_{T_k}^1)]$ is strictly decreasing in T_k. As shown previously, for all $\pi_{T_k} \geq \pi^{**}$, $V(\cdot,\pi_{T_k}^1,\varnothing) - V(\cdot,\pi_{T_k}^1,1)$ is weakly decreasing in T_k. If after receiving either lead, the firm would not post immediately, then this difference is strictly decreasing in T_{k+1}. If the firm would post immediately after receiving either lead, however, than it would be strictly better off posting earlier, which violates the indifference condition. .

Next, consider the retraction threshold $\underline{\pi}(t)$. Assume without loss of generality that message 1 is posted, meaning $a_t = 1$. It would retract at time t if its benefit from retracting is greater than the retraction cost. Suppose that the firm is indifferent between retracting and leaving message 1 posted at t. Given that we know that if the firm would retract, it would do immediately.

If the firm retracts, it's payoff is $V(t, \pi_t^1, \varnothing)$, while if it it does not its payoff is:

$$\int_{T_k}^{1} \left[\int_{T_k}^{s} (1 - (1 - \pi_{T_k}^1)\kappa) d\tau - \mathbb{E}[V(s, \pi_s^1, 1 | \pi_{T_k}^1)] \right] f(s) ds$$

.

If the firm is considering retraction, then there are two possibilities. Either the firm is earning a negative expected flow payoff, or $\kappa < 2$ and $\pi_{T_k}^1 \in [\pi^*, \frac{1}{2}]$. If the firm is earning a negative expected flow payoff and does not retract, it earns this until the next lead arrives, while if it retracts, it either will not post and earn a flow payoff of 0, or will eventually post and get a positive flow payoff. The expected negative payoff from having an article posted is strictly decreasing in T_k. Given that the posterior at the next state is independent of the time, the difference in the firm's payoff from retracting and posting is decreasing over time for a given $\pi_{T_k}^1$. If the firm is earning a positive flow payoff from posting and chooses to retract, it would only do so if it would immediately post the other article. The benefit from making this switch is also decreasing in time.

Therefore, in order for the firm to be indifferent between retracting and keeping the lead posted, it must be the case that $\underline{\pi}(\cdot)$ is strictly decreasing whenever it is greater than 0. $\square$

Proof of Proposition 6

Proposition (Restatement)**.** *Let $\rho' > \rho$, then then $\overline{\pi}' \geq \overline{\pi}$ and $\underline{\pi}' \leq \underline{\pi}$, with these equalities being strict whenever $\overline{\pi}' > \pi^{**}$ and $\underline{\pi}' > 0$.*

Proof. First consider the retraction threshold. As shown in the proof of Proposition 4, if the firm has message 1 posted at time t, it will retract iff $V(t, \pi_t^1, \varnothing) - V(t, \pi_t^1, 1) \geq \rho$, so if $\rho' > \rho$, this difference would need to be greater under ρ' for each time t. As this difference is strictly decreasing in π_t^1, this implies that the retraction threshold is lower.

From Proposition 5, we know that the optimal posting time is monotonically decreasing in the prior. For any arrival time T_k at which an article has not been posted and $\pi^1_{T_k} \geq \pi^{**}$, the firm will post at T_k without waiting iff:

$$T_k \in \underset{t \in [T_k, 1]}{\arg\max} \int_{T_k}^{t} \mathbb{E}V(s, \pi_s^1, \varnothing | \pi^1_{T_k}) f(s) ds + \int_t^1 \left[\int_t^s (1 - (1 - \pi^1_{T_k})\kappa) d\tau + \mathbb{E}[V(s, \pi_s^1, 1 | \pi^1_{T_k})] \right] f(s) ds.$$

Taking the first-order condition of the maximization problem, one gets that: $\mathbb{E}V(t, \pi_t^1, \varnothing | \pi^1_{T_k}) f(t) - \mathbb{E}V(t, \pi_t^1, 1 | \pi^1_{T_k}) f(t) = (1 - (1 - \pi^1_{T_k})\kappa) \int_t^1 f(s) ds$

The left-hand side is increasing in ρ, while the right-hand side does not depend on ρ, so for equality to hold, $\pi^1_{T_K}$ must also increase, which means that for all times T_k, $\bar{\pi}'(T_k) > \bar{\pi}(T_k)$. □

Proof of Proposition 7

Proposition (Restatement)**.** *In the BIS with message-dependent firm revenue:*

- *If $\kappa > 1 + b$, the firm will post message 1 if $\pi^1 \geq \frac{\kappa - b}{\kappa}$, post message 0 if $\pi^0 \geq \pi^*$, and post no message otherwise.*
- *Otherwise, the firm will post message 1 if $\pi^1 \geq \frac{1}{2} - \frac{b-1}{2\kappa}$ and post message 0 otherwise*

Proof. The firm is indifferent between posting message 1 and no message when $b-(1-\pi^1)\kappa \geq 0$, or equivalently when $\pi^1 \geq \frac{\kappa-b}{\kappa}$. Denote this cutoff as π^{b*}. The cutoff for indifference between posting message 0 and no message is the same as in the regular BIS case, which is π^*. It is possible for the firm to exceed both of these thresholds simultaneously when $\pi^1 \geq \pi^{b*}$ and $\pi^0 \geq \pi^*$. Because $\pi^0 = 1 - \pi^1$, this condition is equivalently that $\frac{1}{\kappa} \geq \pi^1 \geq \frac{\kappa-b}{\kappa}$. This interval is only nonempty when $\frac{1}{\kappa} > \frac{\kappa-b}{\kappa}$, or equivalently when $\kappa < 1+b$. In these cases, the firm will always post a message, and it will post the message that gives it the highest expected payoff. This means that it will post message 1 iff $b-(1-\pi^1)\kappa \geq 1-\pi^1\kappa$, which occurs iff $\pi^1 \geq \frac{1}{2} - \frac{b-1}{2\kappa}$ □

Proof of Proposition 8

Proposition (Restatement)**.** *In the BIS with message-dependent firm cost:*

- *If* $\frac{\kappa-1}{\kappa}+\frac{\kappa^0-1}{\kappa^0} \leq 1$, *the firm will post message* 1 *if* $\pi^1 \geq \frac{\kappa-1}{\kappa}$, *post message* 0 *if* $\pi^0 \geq \frac{\kappa^0-1}{\kappa^0}$, *and post no message otherwise.*

- *If* $\frac{\kappa-1}{\kappa}+\frac{\kappa^0-1}{\kappa^0} > 1$, *the firm will post message* 1 *if* $\pi^1 \geq \frac{\kappa}{\kappa+\kappa^0}$, *and post message* 0 *otherwise.*

Proof. As before, the firm would be indifferent between posting message 1 and no message when $\pi^1 \geq \pi^* = \frac{\kappa-1}{\kappa}$. The firm is similarly indifferent between posting message 0 and nothing when $\pi^0 \geq \frac{\kappa^0-1}{\kappa^0}$. Note that $\pi^0 \geq \frac{\kappa^0-1}{\kappa^0}$ is equivalent to the statement that $\pi^1 \leq 1-\frac{\kappa^0-1}{\kappa^0}$ The region in which the firm will not both strictly prefer posting message 1 to posting nothing and posting message 0 to posting nothing is the region in which $\frac{\kappa-1}{\kappa} \geq 1-\frac{\kappa^0-1}{\kappa^0}$, or $\frac{\kappa-1}{\kappa}+\frac{\kappa^0-1}{\kappa^0} \geq 1$. If this condition does not hold, then there is a region in which posting either message is preferable to posting nothing. In this case, the firm posts message 1 iff its expected payoff from doing so if higher than that of posting message 0. This is the case when $1-(1-\pi^1)\kappa \geq 1-\pi^1\kappa^0$, which occurs iff $\pi^1 \geq \frac{\kappa}{\kappa+\kappa^0}$. $\square$

References

Allcott, H. and M. Gentzkow (2017). Social media and fake news in the 2016 election. *Journal of economic perspectives 31*(2), 211–36.

Anderson, S. P. and J. McLaren (2012). Media mergers and media bias with rational consumers. *Journal of the European Economic Association 10*(4), 831–859.

Bäuerle, N. and U. Rieder (2010). Optimal control of piecewise deterministic markov processes with finite time horizon. *Modern Trends in Controlled Stochastic Processes: Theory and Applications 123*, 143.

Brocas, I. and J. D. Carrillo (2007). Influence through ignorance. *RAND Journal of Economics 38*(4), 931–947.

Cancian, M., A. Bills, and T. Bergstrom (1995). Hotelling location problems and television scheduling. *Journal of Industrial Economics 73*(1), 121–124.

Chan, J. and W. Suen (2008). A spatial theory of news consumption and electoral competition. *Review of Economic Studies 75*(3), 699–728.

Duggan, J. and C. Martinelli (2011). A spatial theory of media slant and voter choice. *Review of Economic Studies 78*(2), 640–666.

Ellman, M. and F. Germano (2009). What do the papers sell? a model of advertising and media bias. *The Economic Journal 119*(537), 680–704.

Gentzkow, M., J. Shapiro, and D. Stone (2016). Media bias in the marketplace: theory. In S. P. Anderson, J. Waldfogel, and D. Strömberg (Eds.), *Handbook of Media Economics*, Volume 1b, Chapter 14. North Holland: Elsevier.

Gentzkow, M. and J. M. Shapiro (2006). Media bias and reputation. *Journal of Political Economy 114*(2), 280–316.

Germano, F. and M. Meier (2013). Concentration and self-censorship in commercial media. *Journal of Public Economics 97*, 117–130.

Grossman, G. M. and E. Helpman (2019). Electoral competition with fake news. Working Paper 26409, National Bureau of Economic Research.

Khan, U. and M. B. Stinchcombe (2015). The virtues of hesitation: Optimal timing in a non-stationary world. *The American Economic Review 105*(3), 1147–1176.

Lin, C. (2013). Speed versus accuracy in the news.

Ozerturk, S. (2019). Media access, bias and public opinion.

Pant, A. and F. Trombetta (2019). The newsroom dilemma. *Available at SSRN 3447908*.

Puglisi, R. and J. M. Snyder, Jr. (2016). Empirical studies of media bias. In S. P. Anderson, J. Waldfogel, and D. Strömberg (Eds.), *Handbook of Media Economics*, Volume 1b, Chapter 15. North Holland: Elsevier.

Shahanaghi, S. (2021). Competition and errors in breaking news.

Simon, L. K. and M. B. Stinchcombe (1989). Extensive form games in continuous time: Pure strategies. *Econometrica: Journal of the Econometric Society*, 1171–1214.

Strömberg, D. (2004). Mass media competition, political competition, and public policy. *Review of Economic Studies 71*(1), 265–284.