Media Bias in the Presence of Feedback

David Godes*
University of Maryland

PRELIMINARY AND INCOMPLETE

January 2020

Gentzkow and Shapiro (2006) argue that feedback reduces the incentives for a media outlet to publish biased information. The intuition behind this result is that – assuming consumers place value on finding outlets that will provide true reports with high probability – the outlet should shade reports toward the priors of the public. When feedback – the revelation of the "truth" – is available, there is no reason to shade the report in this way. The authors also provide empirical data in support of this result. In contrast, in this paper, we make the assumption that some consumers seek outlets that provide reports that are biased. In such a setting, we find that feedback actually facilitates biased reports and that, in some cases, the absence of feedback leads to more-truthful news. One important implication of these results is that truthful revelations may not only not reduce the proliferation of "fake news" but may instead increase it. Finally, we provide empirical evidence for bias employing a dataset from a context in which feedback is prevalent: American football (NFL) score predictions.

*Professor of Marketing, Robert H. Smith School of Business, University of Maryland, dgodes@rhsmith.umd.edu. Please do not cite or circulate without permission.

1 Introduction

In this paper, we investigate the presence of bias, or "slant," in the presentation of information to consumers by media outlets. Our primary focus is to present a perspective on the impact of policy options for dealing with such slant. The main argument we make, in contrast to the bulk of the extant literature, is that feedback (truthful revelation) *enhances*, rather than mitigates, the incentives for slant and thus makes it more likely that we observe it. We present a model in which the existence of a segment of consumers seeking media outlets that present directionally-slanted reports (e.g., news favoring a likely, or perceived-likely position) leads firms to adopt slanting strategies. Conditional on having chosen such a slanting policy, a media outlet benefits from consumers' ability to recognize their reports as exhibiting bias concordant with their preferences and stochastically distinct from truthful-reporting outlets. The presence of feedback mechanisms like fact-checking reduces the noise associated with consumers' inferences and, therefore, may increase returns associated with slanted reporting. In this way, we argue that a policy of identifying bias may be counter-productive in the fight against "fake news." After demonstrating this effect analytically, we estimate a difference-in-differences model using a unique dataset of football predictions. This context is particularly appealing for our purposes because it is one in which truthful revelation exists: one eventually knows which team won the game and by how many points. The analysis confirms the existence of bias holding constant team- and game-level differences and, importantly, suggests, consistent with our theory, that the bias subsides as uncertainty is resolved.

The literature on media slant has primarily addressed two sets of questions: (a) Why might media slant exist? and (b) Does media slant exist? If so, how much? With respect to the underlying cause of slant, there are, at a high level, two dominant sets of arguments that have been made. Gentzkow and Shapiro (2006) formulate a model in which consumers attempt to identify the most-accurate provider of news. The media outlet observes a signal of the truth and then decides what to report. In particular, should it report exactly the signal it has received, or should it slant the report in some way? The authors show that incentives for slant exist in this set up because, even as rational Bayesians, consumers may have priors far away from the truth. In such a case, a "true" report (one that simply passes on to the consumer the report received by the outlet) by the outlet would lead the consumer to form a posterior belief that the outlet is not highly-skilled at gathering good reports. In response, then, rational outlets slant news toward consumers' prior beliefs. Again, this happens not because consumers have any preference for a specific news report but because their prior beliefs impact their inferences about the outlet itself. Indeed, the result is driven by the fact that consumers are seeking an outlet that can be inferred to be a likely source of accurate news in the future. An important implication of their theory, to which we return below, is that "feedback" – truthful revelation of the underlying unknown quantity – would remove any incentive for slanting.

Mullainathan and Schleifer (2005), on the other hand, assume explicitly that consumers have a preference for a specific type of news. In this setting, for both monopoly and competitive models, the outlets slant the news towards consumer preferences. While the authors don't explicitly consider the role of feedback, it is clear that truthful revelation of the outlet's signal would not impact the equilibrium outcome since consumers can infer the signal directly from the (slanted) on-path report. Thus, while Gentzkow and Shapiro (2006) would predict that fact-checking media reports would reduce bias, Mullainathan and Schleifer (2005) would predict no impact.[1] In the current paper, on the contrary, we demonstrate that bias may instead be exacerbated by truthful revelation. Theoretical research in this domain has extended the analysis in these papers in a number of important directions. In particular, more-recent studies have focused in large part on assessing the extent to which different contexts may impact the extent of bias, including competition (Gentzkow and Shapiro, 2008; Zhu and Dukes, 2015), consumer heterogeneity (Xiang and Sarvary, 2007), advertising decisions by the media outlet (Ellman and Germano, 2009), and the possibility of user-generated content (Yildirim et al., 2013).

Bias in general and media bias specifically is difficult to identify empirically. Given that bias is typically defined as the systematic and purposeful deviation from the truth, this identification would seem to require that one observe either the truth or a good proxy for it. Thus, much of the empirical literature on media bias has focused on developing novel methods for assessing the presence of bias. Groseclose and Milyo (2005), for example, address this challenge in a three-step

[1] Specifically, outlets announce in advance their "slanting strategy" $s(d)$ where d is their private data signal. Their report is $d + s(d)$ which allows the reader to invert perfectly.

process. First, they collect the Americans for Democratic Action (ADA) scores of left-right lean for legislators.[2] They then use these scores to assess the left-right lean of a set of 200 prominent think tanks by associating each of them with legislators who reference them in the *Congressional Record*. Finally, they assess media outlets' slant based on their references to specific think tanks. Gentzkow and Shapiro (2010) implement a similar three-step process as Groseclose and Milyo (2005) but with important differences. They begin by assessing legislators' slant using their party affiliation as well as their constituents' voting in the 2004 presidential election. Then, they associate these proxies with specific phrases (like "war on terror" or "tax cut") used by the legislators in the *Congressional Record*. Finally, they search the non-opinion content published by the outlet and use the proxies to assess the slant of each. Interestingly, they then incorporate the outlets' slant decisions into a normative model and find that (a) slanting is an element of a profit-maximizing strategy and (b) this is the case because consumers have a preference for news that matches their beliefs. Similar text-analytic approaches have been implemented by other researchers including Gurun and Butler (2012) who identify bias by comparing measures of the negativity of media reports about local companies and national companies. Finally, it is also worth noting that another approach that has been taken to measuring media bias is to compare election endorsements of media outlets. See Chiang and Knight (2011) for an example of this approach.

Distinct from this stream of research, which defines "bias" as referencing topics or ideas that deviate from a constructed "central" legislator or voter, our empirical strategy is instead to compare views which, in the absence of bias, should be identical. Thus, we avoid the need to create explicit proxies for truth using methods like text analysis. Our approach is similar in spirit to that of Dellavigna and Hermle (2017). They also use a difference-in-differences approach in assessing the existence of movie-review bias driven by common ownership of an outlet and a studio. For example, are *Wall Street Journal* reviews of Twentieth-Century Fox-distributed films systematically higher than those for Paramount-distributed films?[3] After controlling for taste correlation using a similarity-based measure, they find no evidence for such bias. On the contrary, we find here evidence of significant bias. Moreover, importantly, our underlying

[2] See https://adaction.org/wp-content/uploads/2018/11/2017.pdf

[3] News Corp owns both the *Wall Street Journal* and Twentieth-Century Fox.

theory is quite different from the conflict-of-interest mechanism proposed by Dellavigna and Hermle (2017). Indeed, we argue that the existence of objective feedback – present in our empirical setting but not, at least explicitly, in theirs – would provide a clean and important explanation for the different results.

The rest of the paper proceeds as follows. In Section 2, we present our theoretical model which predicts that the expectation of future truthful revelation will exacerbate the incentives for bias. In Section 3, we describe the data we have collected for this analysis and discuss some key assumptions we make. In Section 4, we present the model and describe the core results. We conclude in Section 5 with a discussion of limitations and opportunities for future research.

2 Theory

Consider an issue or question around which there exists meaningful uncertainty and let Ψ represent the set of discrete, possible, yet currently unknown, outcomes. The question could relate, for example, to a political election, a scientific question or, perhaps, a football game. The "true" outcome will be drawn or revealed by nature at some point in the future from $\{\Psi\}_{n=1}^{N}$. Consumer beliefs over Ψ are captured by the probability mass function $p\left(n\right)$ such that $\sum_{N} p\left(n\right) = 1$. These correspond with the possible states of the world such as the chosen candidate, the true probability of global warming or the winner of an upcoming game. Media outlets, through their reporting work, gain access to informative signals $\sigma \in \Psi$ of the eventual outcome and decide what announcement $A\left(\sigma\right)$ to present to consumers as a function of the signal received.

Consumers are heterogeneous with respect to their preference for information. A proportion α, whom we'll call "truth seekers," prefer outlets that send truthful signals: $A\left(\sigma\right) = \sigma$. The remaining $1 - \alpha$ consumers, "bias-seekers," seek media outlets that favor disproportionately specific signals, regardless of the media outlet's private signal. Bias-seeking consumers are further heterogeneous with respect to the outcome they prefer to hear about. Specifically, for each $m \in N$, there is a set of consumers of measure $\left(1 - \alpha\right) \beta_m$, $\sum \beta_m = 1$, that seek outlets that send reports consistent with ψ_m being the true state of the world. That is, they prefer outlets that will state that $A(\sigma) = \psi_m$.

Media outlets are also heterogeneous. Let $\theta \in \{\theta^o, \theta^s\}$ represent the type of media outlet, unobserved by consumers, which again can be either honest ($\theta = \theta^o$) or strategic ($\theta = \theta^s$), with equal prior probability. Outlets receive informative, but noisy, signals of the true state of the world ψ^*:

$$\Pr\left[\sigma = n' | \psi^* = \psi_{n'}\right] = \gamma > \frac{1}{2} \tag{1}$$

We further assume that incorrect signals are drawn from the other $N - 1$ possible states with equal probability:

$$\Pr\left[\sigma = n' | \psi^* \neq \psi_{n'}\right] = (1 - \gamma) / (N - 1) \tag{2}$$

An honest outlet always passes through the signal it receives: $A^o(\sigma) = \sigma$. Strategic outlets adopt a strategy to either pass through the signal they receive (i.e., act a an honest outlet) or to select a different, biased report.

*** if the signal is inaccurate, what is the distribution over the other states as to which you receive?

Based on observing a single report, consumers decide whether or not to "adopt" the media outlet at cost c. This cost could be a tangible monetary one or simply an opportunity cost associated with committing time and attention to one outlet to the detriment of another. Each consumer earns utility $U > 0$ if she adopts an outlet that matches her preferences over outlets. That is, truth seekers who adopt an honest outlet and bias seekers who adopt the correct strategic outlet earn U. She earns zero utility otherwise. We further assume that the ratio c/U is uniformly distributed on $[0, 1]$. We assume all distributions are mutually orthogonal. Moreover, we assume an outlet's profit is increasing in "eyeballs" or measure of customers, that adopt them. This could be driven by a number of factors such as subscription fees, advertising revenue or a desire for influence or reputation. We do not model this directly.

To this point, our set up is effectively a discrete version that combines the salient elements of the core models offered by Gentzkow and Shapiro (2006) and Mullainathan and Schleifer (2005). While the former models an outlet concerned about inferences about its ability to deliver accurate information (which we nest here in the form of the honest outlet), the latter incorporates as well consumers who prefer slanted news. Finally, we will add to this the focal element of feedback, or truthful revelation of ψ^*. We see this as informing a policy choice regarding the support of fact checking of media outlet: does revelation of the truth drive bias out of the system? We begin without feedback, in which case the timing of the game is as follows: (1) Nature chooses the type of outlet and the true state of nature ψ^*. The former is revealed only to the outlet. The latter is not revealed to any player in the game. (2) The outlet receives a signal of the true state of nature $\sigma \in \{\Psi\}$ with precision γ. (3) The strategic outlet chooses a strategy (which implies an announcement $A^s(\sigma)$), (4) Consumers decide whether or not to adopt the outlet. (5) Consumer utility and firm profit are recognized. Given the incomplete nature of the information in this game, we adopt the Perfect Bayes Equilibrium (PBE) concept.

While our main focus in the paper is bias in equilibrium and the impact of truthful revelation through some mechanism like fact-checking, we begin by demonstrating that truth is always an equilibrium as well.

Proposition 1 *: For all parameter values, without feedback, there always exists an unbiased PBE in which* $A^o(\sigma) = \sigma$ *and* $A^s(\sigma) = \sigma$*.*

In this pooling equilibrium, on path, it is presumed by all consumers that they receive truthful announcements. Thus, truth-seekers always adopt and biased customers do not. Deviations by the strategic outlet aren't profitable since there is neither an observable distinction to drive a change in consumer behavior (all announcements are presumed truthful) nor a supply-side benefit associated with the deviation. With respect to the former, there is no need need to define here off-path beliefs since off-path outlet behavior is never identified as such. As such, common beliefs-based refinements like the Intuitive Criterion and Divinity have no impact on the existence of this equilibrium.

As a final note, it is worth mentioning that an interesting aspect of this class of models, in which some consumers seek what can be classified as a "bad," is that bias can be, at least without further modeling, welfare enhancing. Within the context of Proposition 1, this implies that for any value of α, the equilibrium exists, but might yield arbitrarily-low levels of welfare if most consumers seek bias. For this reason, we won't apply a pareto comparison to select equilibria. Instead, we'll compare the amount of bias produced as captured by the measure of consumers adopting biased media outlets.

At the same time, we may also observe equilibria, again without feedback, in which the strategic outlet chooses to make biased announcements.

Proposition 2 *Without feedback, there exists a biased-reporting PBE for any n' such that $A^o(\sigma) = \sigma$ and $A^s(\sigma) = n'$ iff*

$$\left[1 + \gamma p(n') + \frac{(1-\gamma)}{N-1}(1 - p(n'))\right]^{-1} \geq \frac{\alpha}{\beta_{n'}(1-\alpha)} \tag{3}$$

A necessary condition for the existence of a biased-reporting equilibrium is, not surprisingly, that there are few truth seekers relative to those preferring to read reports that $\psi^* = \psi_{n'}$. It is also worth noting that the left-hand side of the condition is decreasing in $p(n')$: the more likely the state of the world, the more difficult it will be for a strategic outlet to implement a credible strategy of biased reports touting that state. On the surface, this seems to echo a similar insight by Gentzkow and Shapiro (2006), that one of the reasons for media slant is that honest reporting may lead to negative inferences about outlet quality. To appreciate this insight from their model, consider a context in which the consensus in the market is that the New England Patriots will beat the New York Jets by 50 points. However, imagine that a sportswriter gets information from coaches or players that suggests that the Jets are actually likely to win a close game. If that writer reports her well-founded insight, she is likely to be seen as having poor sources and/or knowledge of the game! As a result, she optimally "slants" her prediction toward the prior. For example, she might predict a 20-point Patriots win. Of course, as we discuss below, this incentive does not persist under feedback.

Here, however, the mechanism is somewhat different in that the sportswriter, in equilibrium, may attempt to purposefully attract the bias-seeking Patriots fan who just wants to hear positive news about his team. As $p(n')$ gets high, it becomes too difficult to signal bias because one looks like the honest outlet. Sticking with our example, being overly positive about the Patriots may be consistently correct. Thus, it is easier to attract bias-seekers for outcomes with *lower probability*. Combined with the comparative static on $p(n)$, this suggests that the most-likely context for equilibrium biased reporting is when a high proportion of bias seekers prefer news about a very low-likelihood states of the world. On the surface, this seems consistent with the prevalence of "fake news" about conspiracy theories, anti-vaccination misinformation and urban legends.

The proposition also shows that the likelihood of a biased-reporting equilibrium is increasing in N, the dimensionality of the state variable. The intuition for this is driven by "false positives" by truthful outlets who receive incorrect signals. As the dimensionality grows, it is easier to identify the biased outlet because it is less likely that the truthful outlet will be saying the exact same thing due to an incorrect signal. As an example, consider again a football context. A bias-seeking Jets fan is looking for an outlet that will give her favorable coverage about her team. What she *doesn't* want is an honest outlet. A biased equilibrium is possible as long as such an outlet can be effective enough at distinguishing itself from an honest outlet. When an honest outlet receives an incorrect signal, it might produce content positive about the Jets. However, this is more likely in, say a five-team league than it would be in a 32-team league. This would also suggest that biased equilibria are more-easily sustainable on continuous variables like probabilities, as modeled by Gentzkow and Shapiro (2006) than in discrete settings like those we allow for here at low N.

Finally, it is straightforward to demonstrate that, when dimensionality N is high enough, as we assume is likely in most cases, the likelihood of a biased-reporting equilibrium is decreasing in γ, the precision of the outlet's signal of the state of the world. The bias-seeker, upon seeing a report that $A = \psi_{n'}$ want to be sure she doesn't adopt a truth-telling outlet. The posterior probability that $\theta = \theta^o$ conditional on observing the announcement $A = \psi_{n'}$, at high N is inversely proportional to the likelihood $\gamma p(n')$. Intuitively, if one sees a report from θ^o then (again at high N) it will be $\psi_{n'}$ with probability $\gamma p(n')$. As precision increases, then, it is more likely that the outlet is the honest type and thus the biased equilibrium is more difficult to sustain.

We now state our main result which addresses the impact of feedback and, by extension, informs the potential efficacy of a policy of fact-checking in an effort to decrease bias. As we demonstrate, such policies may *increase*, rather than decrease, the likelihood of biased media reports. Following Gentzkow and Shapiro (2006), we incorporate feedback by assuming that the game is amended by the addition of a truthful revelation step between (3) and (4) in which ψ^*

is revealed to all players following the media report but preceding adoption.[4]

Proposition 3 *With feedback, there exists a biased-reporting PBE for outcome n' in which $A^o(\sigma) = \sigma$ and $A^s(\sigma) = n'$ iff*

$$\left[\frac{p'(n')}{1+\gamma} + \frac{1-p'(n')}{1+\frac{1-\gamma}{N-1}}\right]\beta_{n'} \geq \frac{\alpha}{(1-\alpha)} \tag{4}$$

Moreover, the likelihood of biased reporting is strictly higher with feedback as compared with the case with no feedback.

In Gentzkow and Shapiro (2006), feedback or truthful revelation allows a consumer to dis-associate inferences about an outlet's quality, or news-gathering ability, from those about the state of the world ψ. In particular, with feedback, reporting about an "unlikely" state doesn't rationally impugn their reporting skills. Returning to our example above, when the reporter is confident that the outcome will be observed by readers, she worries much less about the long-term effect of inferences about her ability; she's confident she'll be proven correct. As a result, it mitigates bias. Here, however, owing to our different model set up – in particular, our consideration of bias-seeking consumers – feedback has a very different effect. In particular, when the state revelation does not support the outlet's report, this provides the strategic outlet with an additional means of differentiating from the honest outlet. This is an important distinction from the existing literature and, again, suggests clearly that such revelation, or 'fact-checking," may actually exacerbate the returns to biased reporting, particularly in the face of a significant proportion of bias-seeking consumers.

Most of the characteristics of the equilibrium posited in Proposition 2 persist here. The likelihood of biased reporting decreases in the precision of their signal (γ) when the dimensionality N is high. Similarly, at high N, the likelihood of biased equilibria is decreasing in the probability associated with the state of the world: less-likely states are easier to sustain in a biased equilibrium. These predictions will play an important role in our empirical analysis to follow.

In summary, our key theoretical result is that, contrary to the finding by Gentzkow and Shapiro (2006), the prospect of future revelation of the true state of the world does not mitigate

[4] The same result would of course be obtained under alternative timing in which revelation occus after adoption as long as consumers were allowed do "un-adopt."

media bias. On the contrary, it exacerbates it. The key implication of this is that even (or, perhaps, especially) in contexts in which the truth will ultimately be known, we would expect to see significant biased reporting in equilibrium. In the next section, we consider one such setting and investigate empirically the existence of bias. Moreover, in order to establish the validity of the model, we test several other implications of the model including (a) whether bias attenuates as signals are more precise and (b) which states (as captured by their prior probability) are better fodder for biased reporting.

3 Setting, Data and Model-Free Analysis

Our empirical context is National Football League (NFL) score predictions. This setting is attractive for at least three reasons. First, this is one of the settings chosen by Gentzkow and Shapiro (2006) to provide evidence for the mitigating effect on bias of feedback. They make use of data collected from predictions made by the *New York Times* sports editor, arguing that bias, if it were to exist, would be observed in the form of overly-positive forecasts of the winning probabilities for the New York Jets and New York Giants. Consistent with their theoretical model, they find no systematic bias in favor of the New York teams. The second reason we focus on football predictions is that, compared with a decade ago, there are now available rich sources of data, allowing for more-granular and more-rigorous empirical analysis than what has been heretofore possible. Finally, and most important, of course, is the fact that this setting is characterized by fairly rapid feedback as well as non-trivial proportions of consumers who are likely to prefer slanted news (i.e., there are likely to be high values of β_n).

The specific data source we make use of here is ESPN.com's weekly "NFL Nation" predictions. NFL Nation provides an aggregation of football content along with a dedicated blog for each team in the NFL. Associated with these blogs are writers, or sets of writers, focusing on each local team.[5] Starting in the 2015 NFL season and continuing through the 2018 season, a weekly feature was added in which, for each game to be played during the upcoming weekend, two writers (one each from the two teams in each game) offer their score prediction. Thus, for each

[5]Note that it is in fact a team and not a market in the sense that the New York Giants and New York Jets have their own blog and writers.

game g, we have four predicted values, the home team writer's prediction of the number of points to be scored by the home team and the away team, as well as the away team's writer's prediction of the same quantities. See Figure 1 for an example of the data and format. Note that both predictions are provided as well as the writer's (i.e., the predictor's) name. We make no use of the associated text in this study.

BILLS AT DOLPHINS

PickCenter

Buffalo Bills

In five games against the Dolphins since he became the Bills' starter in 2015, Tyrod Taylor has an 83.8 Total QBR, his fourth highest against any team. He has completed 65 percent of his passes for 8.99 yards per attempt, nine touchdowns, no interceptions and a 115.6 passer rating against the Dolphins. If Taylor continues that trend, the Bills can awaken from their offensive slump. But will they get the necessary losses from the Ravens, or the Chargers and Titans, in order to make the playoffs? **Bills 21, Dolphins 10** -- *Mike Rodak*

Bills' path to the playoffs

The Buffalo Bills can reach the postseason if one of the following scenarios plays out:

1. Win vs. Dolphins + Ravens loss

2. Win vs. Dolphins + Titans and Chargers losses

• Also see: NFL Playoff Machine

Miami Dolphins

The Dolphins have nothing but pride to play for and already lost to Buffalo earlier this year. But Miami is 2-0 this season in the second game when facing division opponents. Similar to second meetings against the Patriots and Jets, head coach Adam Gase will make the proper adjustments to finish with a 7-9 record. **Dolphins 23, Bills 20** -- *James Walker*

Figure 1: Example of ESPN.com NFL Nation raw data

We collected the predictions for each of the regular-season games for which predictions were available in all four seasons from 2015-2018 inclusive. In each season, each team plays 16 games over seventeen weeks, for a total of 256 scheduled games. Games are played on Thursdays, Saturdays, Sundays and Mondays. A team's schedule is neither random nor identical to other teams, even those in the same division. With respect to a given team's *opponents* in a given season, the team's success in the previous year as well as a fixed rotation are factors. The factors that go into deciding the specific timing of each game, conditional on opponents, is not publicly known, though the league states that they account for equity considerations to ensure that teams face roughly-equal levels of disutility associated with travel.[6] As will be discussed in the next section, heterogeneity associated with scheduling and game difficulty will be handled in the model.

Of the 1,024 games played over these four seasons, we have usable predictions for 959 of them, or 94%. For most of the games scheduled on Thursdays, predictions were not available, presumably due to timing associated with publishing them online. These represent 60 of the 65 unavailable observations. In addition, four observations were dropped because no identifiable writer was associated with the prediction. In the remaining dropped observation, the writer made a conditional prediction dependent on who would start at quarterback. In all, our data from 959 games give us 3,836 individual predictions (two writers x two teams for each game) from 45 different writers. With multiple observations for each prediction, we thus have a unique opportunity to investigate the existence of bias in a context in which our theory predicts it should exist. While most writers provided predictions for a single team over over the 2015-2018 period, it is not a 1-to-1 mapping. See Table 1 for detailed information on the panel of writers and how it evolved over the four seasons for which we have data. One caveat: one shouldn't infer anything about bias from a comparison of a writer's own-team average vs other-team average predictions in Table 1. There is substantial unobserved heterogeneity across teams in terms of their skill level, schedule and expected weather conditions. It is precisely for this reason that we specify and estimate a model in the next section.

Before presenting the formal model, we review some model-free evidence regarding bias. Table 2 presents both the average prediction of points scored by a team as well as the selection of a winner as a function of whether it is the predictor's own team or the team's adversary. As is clear, for both measures, predictors are more favorable toward a predictor's own team than toward her team's opponent.[7] An important consideration in our work here is whether the

[6]See https://operations.nfl.com/the-game/creating-the-nfl-schedule/ for more information on the scheduling process

[7]Recall that, for a given game, there is a prediction by two writers, one each for their own team and for the opponent. Thus, both the set of teams and games is held constant in these averages. The realized values are

Writer	Overall Average	Own Team Average	Other Team Average	Total Games Predicted	Games in 2015	Games in 2016	Games in 2017	Games in 2018
Archer	23.31	24.29	22.32	59	16	15	14	14
Astleford	21.17	20.44	21.89	9	9			
Barshop	19.69	19.86	19.51	43		15	15	13
Brown	21.38	26.25	16.50	4	4			
Cimini	21.18	19.24	23.12	59	15	15	14	15
Cronin	21.94	24.94	18.94	31			16	15
Davenport	20.82	23.06	18.59	17				17
Demovsky	23.80	24.75	22.85	60	15	15	15	15
DiRocco	18.79	17.65	19.93	60	14	15	16	15
Dickerson	20.95	19.74	22.16	61	16	15	15	15
Fowler	24.15	26.74	21.56	57	12	16	14	15
Ganguli	17.75	19.38	16.13	16	16			
Goessling	18.92	21.27	16.57	30	15	15		
Gonzalez	21.95	22.63	21.27	30		15	15	
Graziano	26.03	26.13	25.93	15	15			
Gutierrez	23.20	21.85	24.55	60	15	15	15	15
Harvey	21.30	24.67	17.93	15	15			
Henderson	22.90	24.67	21.13	30			15	15
Hensley	20.22	20.92	19.52	60	15	15	15	15
Kapadia	20.47	23.77	17.17	30	15	15		
Keim	21.44	21.20	21.69	61	15	16	15	15
Kuharsky	20.98	20.37	21.60	30	15	15		
Laine	23.88	22.80	24.96	46		15	15	16
Legwold	21.64	22.64	20.64	59	15	14	15	15
McClure	24.02	26.72	21.32	57	14	15	14	14
McManamon	19.75	15.97	23.54	59	14	15	15	15
McManus	23.13	24.14	22.11	44		15	15	14
Newton	21.75	25.56	17.95	61	17	14	15	15
Raanan	21.53	20.54	22.52	46		15	16	15
Reiss	24.07	28.90	19.24	58	14	15	14	15
Rodak	19.92	18.23	21.61	61	15	15	15	16
Rothstein	22.86	22.77	22.95	62	15	16	16	15
Sheridan	25.13	25.50	24.75	16	16			
Teicher	23.97	25.72	22.23	57	14	15	13	15
Terrell	20.80	20.47	21.13	45		15	15	15
Thiry	26.13	31.13	21.13	15				15
Triplett	25.26	26.17	24.36	59	15	15	15	14
Wagoner	21.17	19.30	23.03	60	15	15	15	15
Walker	18.74	18.29	19.20	45	15	15	15	
Weinfuss	23.30	23.52	23.08	60	15	15	15	15
Wells	22.75	22.38	23.13	61	15	16	15	15
Williams	23.85	24.89	22.82	62	16	15	16	15
Williamson	22.69	22.63	22.75	16	16			
Wolfe	21.23	21.43	21.03	30			15	15
Yasinskas	19.75	19.00	20.50	2	2			
Overall	21.99	22.72	21.26	1918	480	482	478	478

Table 1: Writers in Dataset

predicted differences reflect different information rather than bias. There are two reasons one might suspect that this is not driving the results. First, while a team's dedicated writer may be expected to have better information about her team, it would necessarily be the case that this would lead to both higher and lower predictions, as compared with a less well-informed writer. As we show here, the directionality is systematically skewed toward one's own team. Second, as we also show in Table 2, the average writer is statistically no better at predicting the points scored by her team than by her team's opponents. This does not prove, of course, that there is no information asymmetry whatsoever. However, it suggests that the bias is strong enough that any informational advantage does not overwhelm it.

	Prediction for Own Team	Prediction for Opponent	
Predicted Points	22.54	21.61	***
Realized Points	22.61	22.61	
Predicted Win Proportion	58.56%	41.44%	***
Realized Win Proportion	49.79%	49.79%	
MSE of Point Prediction	104.02	103.77	n.s.

*** = p < .001

Table 2: Model-Free Evidence of Bias

An important prediction of our theory is that bias will decline in the precision of the available pre-announcement signal σ. Here, this would imply that the better informed the writers are, in the sense that the information they've acquired is more precise, the less-biased predictions we expect to see. As a proxy for time variation in signal quality, we use the **?** in the season. The assumption behind this decision is that, at the beginning of the season, a team's strengths and weaknesses are less well known. However, after observing the team's performance over successive weeks, one is better able to judge their ability. Of course, a downside of this approach is that it precludes a consideration of heterogeneity in these learning dynamics.

According to our model in Section 2, this improving precision should make signaling via biased predictions more difficult to implement and, thus, less visible in the data. A simple, and possibly-naive, test of this is to look at the dynamics of the difference in predictions of one's

identical across cases because, again, they reflect the same underlying games and teams.

own team vs the opponent. As shown in Figure 2, we see exactly the trend we'd expect as the difference systematically trends towards zero. We again offer the same caveat: this isn't probative and we need to carefully model the various sources of heterogeneity to be confident we've identified "bias."

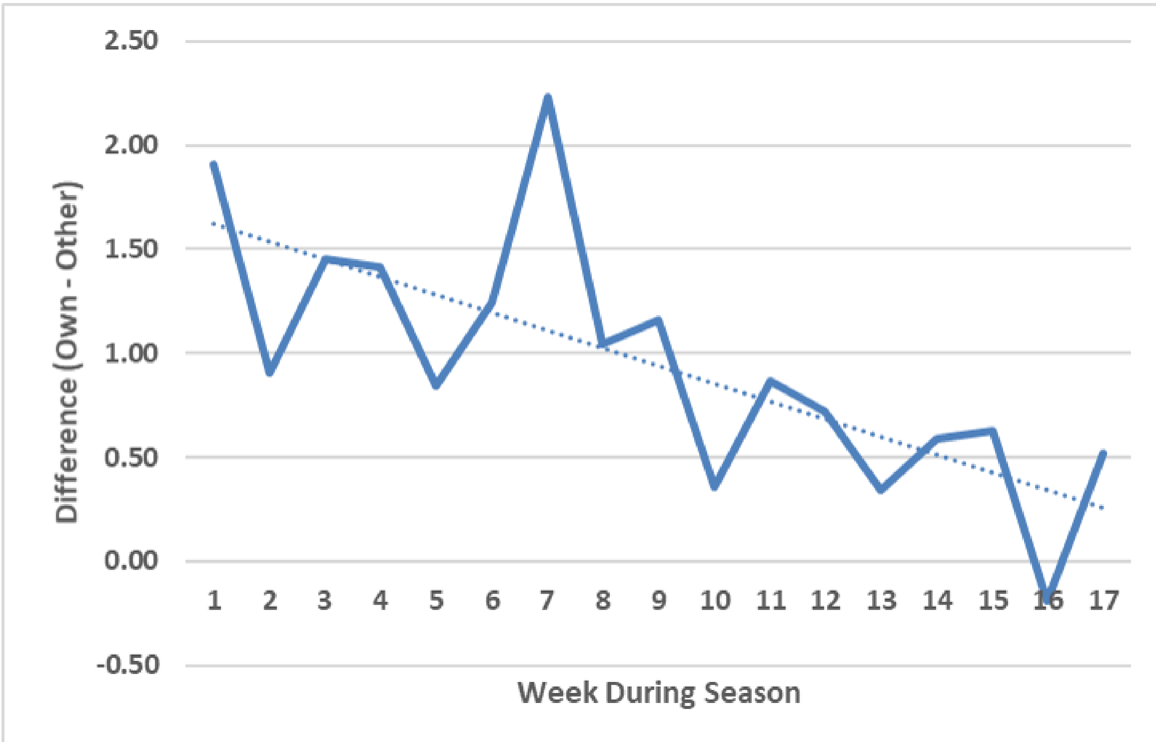

Figure 2: Difference Between Own and Other Predictions Abates Over Time

In order to check whether our assumption that week in a given season is a valid proxy for precision - i.e., that precision improves from week-to-week – is a valid and reasonable one, we perform several analyses. Since the predictions in our dataset are (we claim) corrupted by bias, for these analyses, we collect additional data that is less likely to have the same drawback. Specifically, we collect game predictions from a separate-and-distinct set of experts on ESPN.com for the 2018 season.[8] The key difference between these panelists and those in our primary dataset is that they are not, at least explicitly, associated with any specific team and are, instead, positioned as experts. We would, therefore, expect that there would be less concern about a bias-signal confound here. For each game in each week of the 2018 season, each of the

[8]We were not able to locate similar data from earlier seasons.

ten experts predicts who will win.[9] Here, we first assess the precision of the (common) signal by looking, again, at the agreement across experts over time. Specifically, we calculate for each game in each week the proportion of the panelists who predict the home team will win. We then transform this to a probability and, finally to a measure of standard deviation across the panelists for a given game in a given week.[10] Figure 3a presents, over the 17 weeks of the season, the average across-expert standard deviation in predictions. We can see a marked decline in variation over time. That is, these experts are agreeing with systematically-higher likelihood as time goes on, suggesting that the uncertainty is being resolved.

In Figure 3b, we see that accuracy of the predictions are also improving. While this is quite noisy, the increase from about 60% to well over 70% in week 17 suggests, again, that the quality of the signals one can acquire is improving. Finally, in Figure 3a, we see that the variation (standard deviation) in experts' success in prediction is also improving (i.e., declining). Overall, taken in the aggregate, we argue that the data presented in Figure 3 are highly suggestive that the dynamic week variable is a reasonable proxy for the quality of information available to writers over the course of the season In summary, the model-free evidence suggests that there is reason to suspect the existence of bias in the sense that writers systematically predict higher scores for their own teams relative to others. Of course, this could plausibly be driven by differences in teams' skill levels, writers' access to private information, or other unobservable factors. Thus, in the next section, we build a formal model that addresses these challenges. In addition, we will be able to use the model to test other implications of our theory, including the impact of changing precision (parameter γ in the signaling model) and the sensitivity of the prior beliefs about teams (parameter $p(n)$ in the model).

[9] These predictions are not made against the "point spread," a handicapping system used by gamblers.

[10] To be precise, let $\epsilon^i_{gw} \in \{0,1\}$ represent expert i's assessment as to whether the home team would win ($\epsilon^i_{gw} = 1$) or lose ($\epsilon^i_{gw} = 0$) game g in week w. Then $\pi(w) \equiv N_g^{-1} \sum_G \left[N_e^{-1} \sum_E \epsilon^i_{gw}\right]$ represents the average home-win predicted probability. There are N_g games and N_e experts. Finally, $SD(w) \equiv \sqrt{\pi(w)(1-\pi(w))}$ is thus the weekly standard deviation across games and experts, reflecting the average level of variation in a given week across predictions.

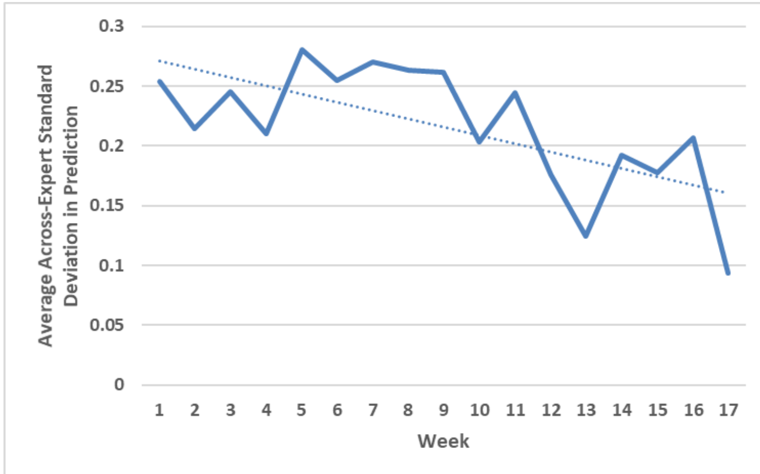

(a) Variation in Predictions Across Experts by Week

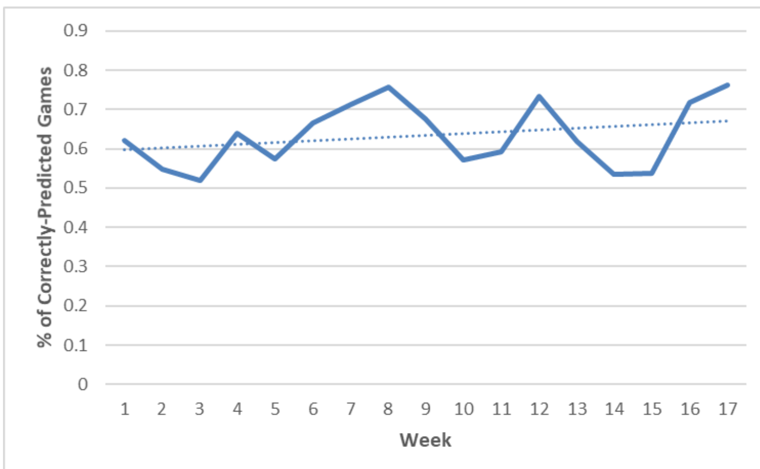

(b) Prediction Accuracy by Experts by Week

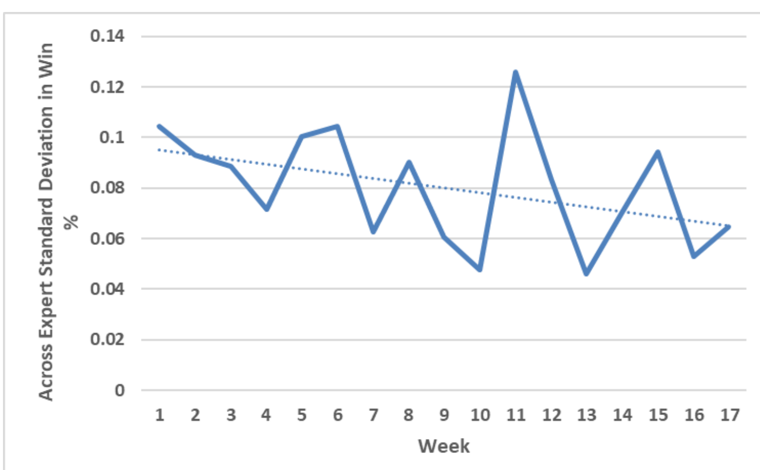

(c) Standard Deviation of Prediction Success Across Experts by Week

Figure 3: The Use of Week as a Proxy for Signal Precision

4 Difference-in-Differences Model

We begin by specifying our difference-in-differences model. At a high level, we model a writer's prediction of a the number of points a given team will score in a given game. Our focal test will assess the extent to which *ceteris paribus* writers makes predictions that are systematically higher for their own teams than for opponents. We first assess the average level of bias in the entire set of writers and then attempt to do so at the writer level.

4.1 Model Specification

Let M represent the set of 32 teams in the NFL, W the set of 45 writers and G the set of 959 games in our dataset. Let p^g_{mw} be the prediction by writer w of the number of points for team $m \in M$ in game $g \in G$. For a given game, each writer $w \in W$ is associated with a team m. according to the mapping $t(w,g) : (W \times G) \rightarrow M$.[11] We will make use of the following:

$$h^g_{mw} \equiv 1\left(m = t(w,g)\right) \tag{5}$$

where $1(\cdot)\cdot$ is the indicator function. Thus, $h = 1$ reflects a situation in which the writer is associated with the team – that is, the "hometown" writer – while $h = 0$ implies she is not.[12] We begin by assuming that each prediction can be described by the following simple linear model:

$$p^g_{mw} \quad = \quad \alpha + \Omega_m + \Omega^g + \Omega^g_m + d_w + h \cdot b^1_w + (1-h) \cdot b^0_w + \varepsilon^g_{mw} \tag{6}$$

Below, we discuss in more detail our assumption on the error structure. In (6), the set of Ω terms represents the "objective" factors that impact w's predictions. We see these as factors that are available to all (at least to all writers) and are evaluated symmetrically by all. We divide these into three distinct sets in order to be as clear as possible. All team-level, game-independent factors are captured in Ω_m. These would include, in particular, team m's persistent skill or coaching expertise associated with scoring points. Team-invariant game-level factors are

[11] As shown in Table 1, most writers are associated with the same team throughout the dataset. However, this is not true for all, which requires this seemingly unnecessarily-complicated notation.

[12] For parsimony, we will drop the super- and subscripts on δ whenever it seems clear without them.

reflected in Ω^g. These would include, for example, the day/time of the game or possibly the weather. Finally, Ω^g_m are game factors that might affect the teams differently. For example, one team might be fighting for a playoff spot in game g while the other is not, or one team may have a number of injuries in game g. Stressing again, the assumption here is that the evaluation of these factors is symmetric across writers. The d_w term reflect individual writer-level (game- and team-invariant) differences in prediction. For example, some writers may just predict higher-scoring games than others, *ceteris paribus*. In order to control for all of these important factors, we implement a difference-in-differences approach, described in more detail below. Our goal is to estimate the magnitude of the parameters b^1_w and b^0_w, which represent, respectively, writer w's subjective bias in predicting scores for her own team and her team's adversary.

Recall that for each game g, we have four team-level score predictions. We begin by differencing within writer w across her predictions for teams s and t:

$$\Delta p^g_w \equiv p^g_{sw} - p^g_{tw} = (\Omega_s - \Omega_t) + (\Omega^g_s - \Omega^g_t) + b^1_w - b^0_w + \varepsilon^g_w, \quad \varepsilon^g_w \equiv \varepsilon^g_{sw} - \varepsilon^g_{tw} \tag{7}$$

Again, a distinctive element of this dataset is that, for each game g, there is a second writer $y \in W/w$ for whom we perform the same differencing. Notice that we will continue to maintain the same ordering of teams[13] and thus take the difference $p^g_{sk} - p^g_{tk}$.

$$\Delta p^g_y \equiv p^g_{sy} - p^g_{ty} = (\Omega_s - \Omega_t) + (\Omega^g_s - \Omega^g_t) + b^0_y - b^1_y + \varepsilon^g_y$$

Notice the difference in the bias parameters in that writer $y's$ bias is the inverse of writer w's. Because, according to the theory, each of the writers in a given game is biased in favor of her own team and against its adversary, we are able to identify the bias even while removing individual scale effects d_w. Indeed, without this, it would be impossible to identify bias independent from scale-usage differences.

We now take the second difference across writers w and y within game g, yielding:

$$\Delta\Delta p^g \equiv \Delta p^g_w - \Delta p^g_y = \left(b^1_w - b^0_w\right) - \left(b^0_y - b^1_y\right) + \varepsilon^g \tag{8}$$

[13] In terms of logistical details, we always difference home team minus the away team

To repeat, the critical point here is that the bias parameters don't get differenced either within writer or across writers because the (a) each writer is assumed to be biased differently with respect to each team; and (b) each writer is assumed to be biased in opposite directions from each other. Our luck runs out beyond this, however, as we are unable to identify separately b_w^1 and b_w^0. Thus, we define the following:

$$b_w \equiv \left(b_w^1 - b_w^0\right) \tag{9}$$

which represents the total bias that writer w has for her assigned team, combining the *upward* bias associated with her prediction for her own team and the (assumed) *downward* bias for her team's adversary. Based on this, we can rewrite (8) as follows:

$$\Delta\Delta p^g \equiv b_w + b_y + \varepsilon^g \tag{10}$$

Again, note that this is, indeed $b_w + b_y$ and *not* $b_w - b_y$ due to the fact that the biases are in the opposite direction. We begin by assuming homogeneous bias across writers such that $b = b_w$ $\forall w$. In this case, then, our estimation equation becomes:

$$\Delta\Delta p^g = 2b + \varepsilon^g \tag{11}$$

Equation (11) is a simple model through which we can test our core theoretical prediction that – even in a context characterized by rapid and highly-precise feedback, and counter to the extant literature – significant bias exists.

To this basic equation, we also add important interactions to test other aspects of the model. First, we consider the precision of the writers' private information γ. Our theory in Section 2 predicts that the extent of bias will decline as the available private information becomes more precise. We do so, as explained in Section 3, by incorporating as a proxy for precision the week during the given season in which the game is played. This serves as a useful proxy under the assumption, supported by the analysis above, that it becomes less and less uncertain as the season goes by to forecast which team will outscore the other and by how much.

We also add a consideration of the prior beliefs associated with a team. In any given contest, a team may have a very high or very low probability of winning. As explained in our analysis in Section 2, this prior plays an important role in whether or not the writer sees an opportunity to gain via biased predictions. Recall that the likelihood of bias in support of team m is decreasing in the prior belief on team m's ablity. Since our model yields an estimate of the *sum* of the biases, we need to consider jointly the signaling power of both writers. As we did with our analysis in Section 3, we seek here a source of information free from the bias in our (or similar) data. For this, we use the closing point spread for each game.[14] According to Proposition 3, for us to observe maximum bias – for both writers to choose to slant their predictions in favor of their teams – it must be the case that:

$$Min\left\{\frac{p_w}{1+\gamma}+1-p_w, \frac{1-p_w}{1+\gamma}+p_w\right\} \geq K \tag{12}$$

Here, we've reflected the symmetry in the process of forecasting what amounts to the score differential in a game. If one team has a high prior – that is, "favored" to win by many points – then the other team has a low prior and is favored to lose by many points.[15] Thus, in order for the minimum to exceed some fixed quantity, it must be the case that the prior is at a "moderate" level. Thus, in order to incorporate this into the model, we include an interaction term as well as its square.

The foregoing gives rise to the following, our main estimation equation:

$$\Delta\Delta p^g = \quad \beta_0 + \beta_1 \cdot WEEK^g + \beta_2 \cdot SPREAD^g + \beta_3 \cdot SPREAD^{g^2} + \varepsilon^g \tag{13}$$

$$\beta_0 \equiv 2b,\ \ \beta_1 \equiv 2b\gamma,\ \ \beta_2 \equiv 2bp,\ \ \beta_3 \equiv 2bp^2 \tag{14}$$

4.2 Results

See Table 3 for the main results. Our estimates reflect a range of assumptions on the error structure. We present both the OLS results as well as those specifying cluster-robust standard errors. In this estimation, the errors allow for correlation among shocks to forecasts by pairs of

[14] We capture this data from vegasinsider.com

[15] Note also that in (12) we've assumed that the size of each segment seeking slanted news from each writer is the same. This is not at all critical to our specification.

Difference-in-Differences Model (d.v. = ΔΔp)					
		Model (1)	Model (2)	Model (3)	Model (4)
Constant (2 x bias)		1.860	3.410	2.780	3.890
	OLS	0.230 ***	0.496 ***	0.283 ***	0.505 ***
	Cluster	0.241 ***	0.515 ***	0.289 ***	0.521 ***
2 x bias x Week			-0.170		-0.130
	OLS		0.047 ***		0.046 **
	Cluster		0.048 ***		0.048 **
2 x bias x line				0.120	0.111
	OLS			0.049 *	0.049 *
	Cluster			0.054 *	0.055 *
2 x bias x line2				-0.033	-0.030
	OLS			0.005 ***	0.005 ***
	Cluster			0.006 ***	0.006 ***
N		959	959	959	959

* p<0.05, ** p<0.01, *** p<0.001

Table 3: Main Results

writers. As is clear, the results are not qualitatively impacted by the choice of error model.

The results are quite consistent with our theory. Centrally, bias exists in this setting with rapid and precise feedback. This suggests that policies addressing fake, or slanted, news by correct or "fact-checking" erroneous reports may not be successful. In fact, our model suggests they may actually exacerbate the problem. Also consistent with our theory, this bias attenuates over the course of the season. The extent of estimated bias is not trivial: at the beginning of the season, writers on average forecast a score differential that is nearly two points higher than it should be. This reflects more than 25% of the average predicted margin of 7.7 points over the four seasons in our data. The impact of precision is also important: the bias declines by about .03 points in each week implying that the bias is cut by about 25% by the end of the season. Finally, consistent with our theory, the predictions seem to associated with the prior in such a way that the bias is highest in the interior. Very high or very low expected differentials are not as closely associated with bias in this dataset.

Finally, we relax the assumption of homogeneous bias across writers by estimating a variation of (13) in which we estimate a separate bias parameter β_{0w} for each writer. We continue to maintain the assumption of homogeneous β_1-β_3. A few comments on the results, which we do not provide in their entirety, First, as one would expect, this model, by virtue of capturing individual differences in incentives (and proclivity) for slanting predictions, explains a great deal more of variance in the data. While the homogeneous model explains less than 5% of the (twice-differenced) data, the individual model explains over 20%. Of the 45 writers in our sample, 21 of them exhibited a significant ($p < .05$) bias and all but one of them were in the "correct," that is hypothesized, direction. It is important to note as well that the estimates on β_1-β_3 are qualitatively identical to those in Table 3.

5 Conclusion

In this paper, we re-visit the seminal analysis by Gentzkow and Shapiro (2006) and, in particular, the implication that the prospect of feedback can attenuate the incentives for media outlets to slant their reporting. We extend their model by adding a feature that is popular in the broader literature on media bias: bias-seeking consumers. With this, we derive results in a signaling framework that suggest that feedback – which we associate with "fact-checking" – may not reduce media bias. In fact, as we show, it may exacerbate it. The intuition behind this result is that the feedback provides bias-seekers an additional signal on the media outlet's type. In particular, a "false" fact-check can be interpreted by these consumers as a strong signal that the outlet is, in fact, not only slanted but slanted in the direction of interest to them.

We then test our model in an empirical context characterized by rapid and precise feedback: American football. We construct a unique dataset of football predictions and demonstrate via a Difference-in-Differences framework that significant bias exists. On average, writers slant their predictions of the score differential about 2 points per game in favor of their team. We also find, consistent with our theory, that the bias declines over the course of the season, which we argue is due to an improvement in the precision of the information available to writers. Finally, we show that the bias is highest when the two teams are more closely matched.

Our work could be extended in a number of interesting directions. First, of course, this is a single study in a single (and possible idiosyncratic) context. Testing the extent to which the same results are produced in other contexts would be very valuable. Another direction one might consider would be to extend the methodological approach from the quasi-experimental to lab or field experimental setting. This would allow for tighter control over the inquiry and would, in particular help in isolating the impact of priors on the incentives to slant as well as the dimensionality, which is of course held constant here.

REFERENCES

Chiang, C.-F. and Knight, B. (2011). Media bias and influence: Evidence from newspaper endorsements. *Review of Economic Studies*, 78:795–820.

Dellavigna, S. and Hermle, J. (2017). Does conflict of interest lead to biased coverage? evidence from movie reviews. *Review of Economic Studies*, 84(4):1510–1550.

Ellman, M. and Germano, F. (2009). What do that papers sell? a model of advertising and media bias. *Marketing Science*, 119(537):680–704.

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

Gentzkow, M. and Shapiro, J. M. (2008). Competition and truth in the market for news. *Journal of Economic Perspectives*, 22(2):133–154.

Gentzkow, M. and Shapiro, J. M. (2010). What drives media slant? evidence from u.s. newspapers. *Econometrica*, 78(1):35–71.

Groseclose, T. and Milyo, J. (2005). A measure of media bias. *The Quarterly Journal of Economics*, 120:1191–1237.

Gurun, U. and Butler, A. (2012). Don't believe the hype: Local media slant, local advertising, and firm value. *Journal of Finance*, 67:561–598.

Mullainathan, S. and Schleifer, A. (2005). The market for news. *The American Economic Review*, 95(1):1031–53.

Xiang, Y. and Sarvary, M. (2007). News consumption and media bias. *Marketing Science*, 26:611–628.

Yildirim, P., Gal-Or, E., and Geylani, T. (2013). User-generated content and bias in news media. *Management Science*, 59(12):2655–2666.

Zhu, Y. and Dukes, A. (2015). Selective reporting of factual content by commercial media. *Journal of Marketing Research*, 52(1):56–76.

Proof of Proposition 1: In an unbiased equilibrium, every outlet is presumed to be telling the truth. Thus, truth-seekers always adopt. Biased consumers never adopt. Deviations by the firm have no effect on these decisions.□

Proof of Proposition 2: Conditional on observing $A = n'$, the posterior belief over θ is given by:

$$\Pr\left[\theta^s | A = n'\right] = \frac{1}{1 + p\left(n'\right)\gamma + \frac{(1-\gamma)}{N-1}\left(1 - p\left(n'\right)\right)} \tag{15}$$

so the bias measure of bias seekers who adopt is:

$$\left[1 + \gamma p\left(n'\right) + \frac{(1-\gamma)}{N-1}\left(1 - p\left(n'\right)\right)\right]^{-1} \tag{16}$$

Conditional on observing an off-path announcement $n'' \neq n'$, the inference is that it is that $\theta = \theta^o$ with probability 1. Therefore, assuming condition (i) is met, the strategic type will deviate to n'' if:

$$\alpha > (1 - \alpha)\beta_{n'} \tag{17}$$

□

Proof of Proposition 3: Conditional on the revelation of ψ^*, the on-path posterior belief by consumers are:

$$\Pr\left[\theta^s | A = \psi_{n'}, \psi^* = \psi_{n'}\right] = \frac{1}{1+\gamma} \tag{18}$$

$$\Pr\left[\theta^s | A = \psi_{n'}, \psi^* \neq \psi_{n'}\right] = \frac{1}{1 + \frac{1-\gamma}{N-1}} \tag{19}$$

The *ex ante* expected measure of biased-n' customers adopting is thus

$$\frac{p\left(n'\right)}{1+\gamma} + \frac{1 - p\left(n'\right)}{1 + \frac{1-\gamma}{N-1}} \tag{20}$$

To see that this is represents a higher return, note that satisfying the claim requires that:

$$\frac{p\left(n'\right)}{1+\gamma} + \frac{1 - p\left(n'\right)}{1 + \frac{1-\gamma}{N-1}} - \frac{1}{1 + \gamma p\left(n'\right) + \frac{(1-\gamma)}{N-1}\left(1 - p\left(n'\right)\right)} > 0 \tag{21}$$

We begin by differentiating this expression with respect to N, which yields:

$$(1-p(n'))\left[1+\frac{1-\gamma}{N-1}\right]^{-2}\frac{1-\gamma}{(N-1)^2}-\left[1+\gamma p(n')+\frac{(1-\gamma)}{N-1}(1-p(n'))\right]^{-2}(1-\gamma)(1-p(n'))(N-1)$$
$$=K\left[\left[1+\frac{1-\gamma}{N-1}\right]^{-1}-\left[1+\gamma p(n')+\frac{(1-\gamma)}{N-1}(1-p(n'))\right]^{-1}\right], K>0$$

Since $\left[1+\frac{1-\gamma}{N-1}\right]-\left[1+\gamma p(n')+\frac{(1-\gamma)}{N-1}(1-p(n'))\right]=-\gamma p(n')+\frac{p(n')(1-\gamma)}{N-1}<0$, this implies that (21) is increasing in N. Thus:

$$\frac{p(n')}{1+\gamma}+\frac{1-p(n')}{1+\frac{1-\gamma}{N-1}}-\frac{1}{1+\gamma p(n')+\frac{(1-\gamma)}{N-1}(1-p(n'))}$$
$$>\frac{p(n')}{1+\gamma}+\frac{1-p(n')}{2-\gamma}-\frac{1}{1+\gamma p(n')+(1-\gamma)(1-p(n'))}$$

Now, differentiating this with respect to $p(n')$ yields:

$$\frac{1}{1+\gamma}-\frac{1}{2-\gamma}+\frac{2\gamma-1}{(1+\gamma p(n')+(1-\gamma)(1-p(n')))^2}$$

Which is clearly monotonically declining in $p(n')$. This concavity implies a unique internal maximum and thus, in order to prove our claim, we need only check the endpoints $p(n')=0$ and $p(n')=1$, As these are both zero, this proves the claim. $\square$