"TELL ALL THE TRUTH, BUT TELL IT SLANT": TESTING MODELS OF MEDIA BIAS

COLLIN RAYMOND, UNIVERSITY OF OXFORD

SARAH TAYLOR, DUKE UNIVERSITY

Abstract. Media outlets often bias their news reports. It can be difficult to test the motives behind such biased reporting. We develop models of news provision that detail how the environment and preferences affect demand for, and provision of, news. We examine the implications of these models of media bias, and then test the predictions using weather reports by the *New York Times* in the late 19th century. This period in history provides a natural experiment for testing bias, as the *Times* switched from using government weather reports to producing weather reports in-house. When the *Times* produced its own weather reports it was more accurate at predicting sunny weather and less accurate at predicting rainy weather on days when the local baseball team, the New York Giants, was scheduled to play a home game in Manhattan. The amount of bias is smaller during months that have the least rain, and is higher when the New York Giants are more highly ranked. We find that the empirical results falsify standard neoclassical explanations of media bias. The comparative statics may be better explained by a model of bias with intrinsic preference for information.

1. Introduction

Social and behavioral scientists across several disciplines, including economics, psychology, and political science, are interested in the causes and consequences of media bias. News content can have significant effects on beliefs, actions and outcomes. However, it can often be difficult to assess the existence and impact of bias. First, most measures of bias are by necessity comparative. It can be difficult to ascertain whether the *New York Times'* political reporting is liberally biased relative to the truth, but much easier to ascertain whether it

---

Date: October, 2013.

The first line of the title is due to a poem of Emily Dickinson, who incidentally often wrote in slant verse. We'd like to thank Oliver Cairnes for excellent research assistance. We would also like to thank Yan Chen, Stefano DellaVigna, Rachel Kranton, Yusufcan Masatlioglu, Stephen Salant, Daniel Silverman, Mel Stephens, and Lowell Taylor for helpful comments, as well as seminar participants at the Midwest Economics Association Annual Meeting (2012), University of Oxford, Sabanci University, and the North American and European Summer Econometrics Meetings. Any and all errors are ours alone. Emails: collin.raymond@economics.ox.ac.uk and setay@umich.edu.

is liberally biased relative to Fox News. Second, bias may be caused by individuals' beliefs about likely outcomes, but biased media reports can also lead to changes in outcomes. For example, if the media is biased towards reporting an Obama victory, this may in fact make an Obama victory more likely. Sorting out these complications is a potentially important undertaking; many people believe that media bias can have significant welfare effects, causing rising polarization and mistrust in the media.

Of course, addressing if, and to what extent, consumers are harmed by bias requires a theory of media bias as well as an empirical assessment of key predictions generated by the theory. Moreover, many models of media bias can predict a variety of types of bias depending on the parameters of the models, and so it is necessary to test not just the direction of bias but also how bias changes with the environment. Therefore it is important to find data where there is exogenous variation in the environment in order to test the comparative static predictions of models of media bias.

In this paper, we develop models of news provision and test predictions using novel data, which allow us to identify possible driving mechanisms behind media bias. Such analysis gives further insight into the welfare impact of bias. We document bias in a very simple setting — historical weather reports by the *New York Times*. In the late 1800s Manhattan was home to a professional baseball team, the New York Giants. Before June 16, 1896 the *New York Times* used government-produced weather predictions. The *Times'* accuracies of predicting sunny weather and rainy weather were roughly equal on days when the Giants were playing at home and on other days during the baseball season. On June 16, 1896 the *New York Times* switched to producing in-house weather reports. Our focus is the time frame, June 16, 1896 through the end of the century. For this period, we find an interesting form of bias: the *Times* was more accurate at predicting sunny weather and less accurate at predicting rainy weather on days when the Giants played home games, relative to non-home game days. We use this striking observation to test models of media bias. We find that the excess bias in favor of sunny predictions on game days shrinks as the ex-ante probability of fair weather grows. Furthermore we find that the excess bias in favor of sunny predictions on game days grows as the baseball team does better. These results are inconsistent with standard rational models of bias, including models where the *New York Times* was trying to encourage attendance at Giants' games. However, our comparative statics can be consistent with models of utility over beliefs.

Our research design has several benefits. First, it features a regime shift in the weather reporting by the *New York Times* from the use of Weather Bureau reports to in-house weather reporting. Second, there is a clear exogeneity in realized outcomes — changes in the *Times'* weather predictions do not cause the realized weather to change. Third, we can observe both the predictions and the realized weather, meaning we can develop a measure of absolute bias, not simply comparative bias. Fourth, the type of bias we observe should lead to different posteriors on the part of the consumers, if they are aware of it. This is in contrast to many other papers, which study how different media outlets report the same event using different terms, leading to difficulty in assessing how individual beliefs should respond to biased information. Fifth, we can easily observe how bias changes with the parameters of the model. Last, but not least, the actual environment closely aligns with a tractable model.

A theory of media bias that encompasses all the important facets of the phenomena would be quite extensive. In this paper we set for ourselves a much simpler goal: We seek to identify possible driving mechanisms behind media bias in a stylized setting in which the media provides predictions about two potential states of the weather ("rainy" or "fair"). In Section 2 we consider several models of information provision the media and generates predictions and comparative statics regarding consumer-induced preferences for accuracies. In turn, we develop equilibrium predictions regarding how news-providing firms will respond to such consumer preference — in particular whether firms will provide differently biased news when consumer's actions, priors and payoffs change. We show how the environment and preferences can interact to generate biased signals; different mechanisms yield different predictions regarding the observed bias in news reports.

In Section 3 we test our predictions by evaluating the daily weather reports in the *New York Times* during the late nineteenth century. In the 1890s the *New York Times* published a daily weather report in its morning edition which was extremely simple, typically offering either a rather straightforward prediction for fair weather or for rain. This makes the reports easy to interpret for the purpose of statistical analysis. From 1890 through mid-1986 the *New York Times* used weather reports produced by the U.S. Weather Bureau in Washington

D.C., and its predecessor, the Signal Service.[1] These forecasts were not very accurate (see Nekeber, 1995), but they were produced by a trusted independent source; it seems unlikely that they were slanted so as to please New York baseball fans. On June 16, 1896, the *New York Times* began producing the daily forecast in-house.[2]

The baseline test is to see if the *Times* was more likely to provide biased weather predictions on days when the New York Giants were scheduled to play a baseball game in their home park, the Manhattan Polo Grounds. We observe that reports are indeed biased for the subsample of home game days after June 16, 1896, and that correct reports for fair weather increase while correct reports of rainy weather decrease, i.e., weather reports are "optimistic" about the probability of fair weather on home game days. We run additional tests to understand how the bias changes as the environment changes — in particular what occurs when the priors held by consumers or the payoff of attending a game change. We find that as months have a higher ex-ante probability of rain, the amount of bias grows. Noting that the Giants held various positions in the National League rankings in the years after the switch to in-house forecasting, we also examine how the bias changes with the yearly ranking of the baseball team. We find that the amount of bias grows as the baseball team does better.

In Section 4 we we relate our empirical findings to the theoretical predictions we derive in Section 2. In particular, we find that the evidence does not support the bias being driven by consumers' desire to better match their action to the state nor by the *New York Times* trying to increase attendance at Giants' games. We also consider the predictions of two other popular models of media bias, Gentzkow and Shapiro (2006) and Mullainathan and Shleifer (2005), and show that they are also inconsistent with the data. Because both the evidence on the comparative statics of bias and of the effects of weather predictions on attendance fail to match the predictions of models where information has instrumental value, we believe that the best model to explain our data is where consumers have intrinsic preferences over

[1]Newspaper forecasts were the primary avenue by which people learned about the weather, since radio did not become popular until decades later. In fact, the first radio broadcast for entertainment purposes was not transmitted until 1919.

[2]Evidence for such change comes from the fact that the *New York Times* always cited the Weather Bureau before June 16, 1896. On and after this date, the *Times* did not cite the Weather Bureau. As an additional point, during the period of in-house forecasting, the *New York Times* predictions diverged from the *Wall Street Journal* (as we discuss in more detail below), which sometimes cites the Weather Bureau as a source.

information. We show that a simple model of wishful thinking, where individuals have preferences over information structures even when they do not condition their actions on information, can rationalize the observed data patterns.

Our paper contributes to a burgeoning literature in economics examining media bias. There are several existing models of media bias, including Mullainathan and Shleifer (2005), Suen (2004), Gentzkow and Shapiro (2006), and Baron (2006). In Mullainathan and Shleifer (2005) bias in news reporting arises because readers have a form of confirmation bias; they generally hold views that differ from the true state of the world and get disutility from reading news that is inconsistent with those views. Gentzkow and Shapiro (2006) present a model of news bias with Bayesian consumers who value accuracy in news, and infer the accuracy of reports based on priors. In this setting, news outlets bias toward reader priors as a way of boosting inferred quality. Gentzkow and Shapiro (2008) discuss various mechanisms for bias and show how competition can either enhance or ameliorate bias. Our theoretical model is also closely related to Gentzkow and Kamenica (2011), who look at how a principal can persuade Bayesian agents by utilizing different information structures.

There has also been a great interest in identifying bias in the media. Many papers have focused on determining the extent and causes of media bias. Most of these papers, including Groseclose and Milyo (2005), Gentzkow and Shapiro (2004, 2010, 2011), Gentzkow, Shapiro and Sinkinson (2012), Gentzkow, Petek, Shapiro and Sinkinson (2012), and Durante and Knight (2012) focus on the causes of bias in the context of political ideology. In contrast, DellaVigna and Kaplan (2007) and Gentzkow, Shapiro and Sinkinson (2012) examine the impact of media bias on behavior, again in a political context. DellaVigna and Gentzkow (2010) survey the evidence on persuasion in markets, and consider different forces that could drive persuasion. In a related vein of work, Oster, et al. (2011) also provides a behavioral framework for understanding the motivation behind preference for information.

The rest of the paper is organized as follows. Section 2 develops a model of the market for news and presents the implications for consumer demand and equilibrium provision of information. Section 3 details the data and presents results of the empirical tests. Section 5 relates the evidence to theory, and Section 6 concludes.

2. Models of Information Provision

In this section we present two possible models of information provision. Our setting is stylized; we consider a world with two payoff relevant states and binary signals. Individuals have two actions — a safe action, where the payoff is invariant to the state, and a risky action, in which the payoff depends on the state. Despite this simplicity, we believe it captures the essential details of the empirical work that follows.

In line with the literature, we distinguish between multiple sources of bias; we consider the implications of "supply driven" and "demand driven" bias. If bias is supply driven, firms receive a payoff that varies with the decision maker's action. If bias is demand driven, consumers have preferences over informational structures, and a profit maximizing firm provides signals in line with those preferences. We focus on a setting with a monopolistic provider of information.[3] The firm chooses the information structure to provide, but does not vary the price once a structure is chosen (similar to the fact that we do not observe variation across game days and non-game days in the price of the *New York Times*). Below, we consider the predictions of models where individuals behave according to expected utility, i.e., we consider "neoclassical" models. In our conclusion we discuss the implications of models where consumers have intrinsic preferences over information.

2.1. **Environment.** We assume that there are two states of the world, A (good weather) and B (bad weather). A decision maker takes one of two actions a (attend a game) and b (stay at home). Payoffs to the decision maker depend on action i and state j, and are denoted $u(i, j)$. Action a is the risky action, $u(a, A) > u(a, B)$, while b is safe, $u(b, A) = u(b, B)$.[4] Neither action dominates, and so it is better to match action to state: $u(a, A) > u(b, A) > u(a, B)$. We normalize the payoff of staying at home $u(b, A) = u(b, B) = 0$, and denote the relatively

[3]Focusing on a single firm makes the model more transparent. Moreover, the monopolistic information provided in our model of demand-driven bias provides the consumer an optimal information structure, a finding that would be replicated under competition. Therefore the comparative statics will be equivalent. In the case of supply-driven bias, if all firms receive a benefit from consumers attending the game, then again we would expect to see the exact same comparative statics. If on the other hand some firms do not receive a benefit from consumer actions, then we would not expect to see supply-driven bias in the marketplace at all.

[4]The results below can easily be generalized to accommodate the case where action b is also risky, so long as the difference in payoffs across states given action b is smaller than the difference in payoffs across states given action a.

high payoff of going to the game in good weather $u(a, A) = H > 0$ and the lower payoff of going to the game in bad weather as $u(a, B) = L < 0$.

There is also a news provider which has access to a set of prediction technologies. We characterize a prediction technology as an information structure with two accuracies: If it will be good weather, it generates a signal that indicates *fair* with probability p. If the day will be bad weather, it indicates *rainy* with probability q. Throughout the paper we will assume that both firms and consumers know the chosen technology's values of p and q. In the case of demand-driven bias, since firms are responding to consumer's preferences, consumers have to realize at some level that p and q are changing. Otherwise, firms have no incentive to change p and q. In the case of supply-driven bias, the comparative statics will be exactly the same if consumers are naive and do not realize that the firm is manipulating p and q.

Given a prior $0 < \rho < 1$ on state A (fair weather), along with a particular information structure (p, q), we can easily find a pair of posteriors for a decision maker. After observing a fair signal the posterior for A is

$$\psi_F = \frac{\rho p}{\rho p + (1-\rho)(1-q)}. \tag{1}$$

After observing a rainy signal the posterior is

$$\psi_R = \frac{\rho(1-p)}{\rho(1-p) + (1-\rho)q}. \tag{2}$$

We structure our model so that a fair weather prediction increases beliefs about the weather being good, while a bad weather prediction decreases them. This implies that $p + q \geq 1$, which is without loss of generality within the space of all signals, $(p, q) \in [0, 1] \times [0, 1]$. We show these results, and provide an additional observation in the following lemma.[5]

Lemma 1 *The set $\{(p, q) | (p, q) \in [0, 1] \times [0, 1]$ and $p + q \geq 1\}$ satisfies three properties:*

(1) *Observing a fair signal increases the posterior on state A relative to the prior, and observing a rainy signal decreases the posterior on state A relative to the prior.*
(2) *For any signal structure $(p', q') \in [0, 1] \times [0, 1]$, there exists a $(p, q) \in \{(p, q) | (p, q) \in [0, 1] \times [0, 1]$ and $p + q \geq 1\}$ that generates the same posteriors with the same probabilities as (p', q').*

[5] All proofs are presented in Appendix B.

(3) *For any strict subset S of $\{(p,q)|(p,q) \in [0,1] \times [0,1]$ and $p+q \geq 1\}$, there exists a point $(p',q') \in [0,1] \times [0,1]$ such that there is no element of S that generates the same posteriors as (p',q').*

The first point of the lemma shows that our convention of considering only $p+q \geq 1$ leads to the desired natural interpretation of the "accuracy" probabilities, p and q. The second point establishes our claim about generality. The third point provides a helpful backdrop as we turn to our specification of weather prediction technology.

Our approach to the prediction technology is to assume an inherent tradeoff between the accuracy of predicting fair weather, p, and the accuracy of predicting rain, q. In order to increase accuracy in one state, a news provider must reduce it in the other. To model this, we define a convex set of feasible signals $\Phi(p,q)$.[6] We define the production frontier of the feasible set using $\phi(p)$ — the maximal value of q given p. Therefore $\Phi(p,q) = \{(p,q)|q \leq \phi(p)$ and $p+q \geq 1\}$. Of course $\frac{\partial \phi(p)}{\partial p}$ is negative, and we assume that the absolute value of the tradeoff between p and q is increasing in p, so that $p < p'$ if and only if $\frac{\partial \phi(p)}{\partial p} > \frac{\partial \phi(p')}{\partial p}$ (this is equivalent to $\frac{\partial^2 \phi}{\partial p^2}$ being negative).

Figure 1 shows an example in which the technology set is symmetric around the reflection line $p = q$ (a condition we will continue to adopt).[7] For this case it is natural to define an information structure's *bias* as $\mathbb{B} = p - q$. Here, when $p = q$, bias is zero, and the signal structure is designated *neutral.* The news provider can alternatively choose a signal structures for which $p > q$, in which case we have positive bias, which might also be called *optimistic* (since fair predictions now occur relatively more often than rainy predictions given equal realizations of weather). Signal structures where $q > p$ are *pessimistic.* Notice that in Figure 1, in line with the previous paragraphs, signals below $p+q = 1$ are not feasible (since they are equivalent to signals above that line), and signals above $\{(p,q)|q \leq \phi(p)\}$ are not available. In the figure, the example shown, (p^*, q^*), is a neutral case.

This set up corresponds in a natural way to our empirical work below. We look at bias $\mathbb{B}$ for a particular subset of days, say subset I, and then see if bias is different on an alternative subset J, i.e., we evaluate "differential bias" or "excess bias," in set I relative to J. $\Delta_{I,J} = \mathbb{B}_I - \mathbb{B}_J$. Evaluating excess bias is an important excercise when empirically testing the model.

[6]Because our empirical application entails looking at weather prediction over a relatively short time period, we do not consider changes in the set of feasible signals due to technological advancement.

[7]The symmetry assumption, however, is not necessary for the comparative statics we develop below.

This is because there might be exogenous reasons why there might be bias — e.g., sunny weather might be easier to predict than rainy weather. However, these exogenous reasons should not vary across sets I and J. Thus, for example, we look to see if the *New York Times* weather reports exhibit different biases on home game days than on non-game days.

Intuitively, consumers who use weather predictions to condition decisions will prefer predictions that are more accurate; among structures in the feasible set, they will prefer structures on the boundary of the feasible set (as in the example shown in Figure 1). This idea — about the "informativeness" of the signal structure — is formalized in the following lemma:

Lemma 2 *Assume p and q jointly satisfy $p+q \geq 1$. Then (p', q') is Blackwell sufficient/more informative for (p, q) if and only if $p' \geq \frac{p}{1-q} - \frac{p}{1-q}q'$ and $p' \geq 1 - q'\frac{1-p}{q}$.*

If (p', q') is more Blackwell informative than (p, q) then the posteriors under (p', q') are a mean preserving spread of the posteriors under (p, q) — a result that follows from the law of iterated expectations. Clearly, this will be true if $p' > p$ and $q' > q$, but Lemma 2 shows it can also be true under less stringent conditions. Figure 2 illustrates using an example in which $(p, q) = (0.6, 0.6)$.

With all this in mind, we turn to the time line in our environment. In period 1 the decision-maker and news-provider share a prior ρ on the probability of a future state, in period 3, being A. A monopoly news provider picks an information structure and sets the price of the news. Consumers, knowing the information structure, and the price for information, choose whether to purchase information or not. In period 2, the news-provider receives a signal, and passes it along to their customers who purchases the information structure. Consumers choose an action. In period 3, the state is realized and consumers receive payoffs that depend on their actions and the realized state. Our restriction on the timing of actions is substantive, and we attempt to test this in data. We make this substantive restriction in order to provide the best case possible for a rational explanation of media bias.[8]

[8]Instead of using the assumption of "full commitment," where an agent has to choose an action in period 2, and so information is instrumentally valuable, we could make use of an alternative assumption that the agent makes her choice in period 3, after observing the realization of the state. In this latter case, information obviously has no instrumental value (i.e., it cannot help in decision making). Therefore neoclassical consumers have no value of information. Bias would result only from intrinsic preferences over beliefs and information. These two situations are obviously limiting cases of a more general model where the agent must make decisions in period 2 which are costly (either explicitly or implicitly) to realization in period 3.

2.2. **Preferences over Information Structures.** Because no material payoffs occur before period 3, we model the decision-maker's environment as a compound (2-stage) lottery. Each set of payoffs $(H, L, 0)$, prior probability ρ, and an information structure (p, q) together generate a unique compound lottery. The technology constraint, along with the payoffs and prior, define a feasible set of compound lotteries. The information structure that leads to the highest expected utility for the decision-maker, given parameters, is denoted $(p^*(H, L, \rho), q^*(H, L, \rho))$. (We will typically suppress the dependence on the other parameter values to simplify notation.)

Expected utility maximizers receive no flow utility in periods 1 and 2, and so simply try to maximize the expected utility received in period 3. As noted above, if the agent takes action a, and the realized state is A, the payoff is H; if she takes action a and the realized state is B, the payoff is L; and if she takes action b, the payoff is 0 regardless of the realized state. Figure 3 demonstrates expected payoffs from each of the agent's action in terms of beliefs and utility. The horizontal axis represents the agent's belief in the probability of state A (fair weather). The vertical axis shows the expected payoff from actions conditional on those beliefs.

Using Figure 3 as an example, if the consumer has a prior $\rho = \frac{1}{2}$ and she receives no other information, her optimal action is a, i.e., she goes go to the game (since the average of L and H is greater than 0). (In the figure, her expected utility is designated "$u|$ no information"). Now suppose she receives weather predictions using signal structure S, given by $(p, q) = (1, \frac{1}{2})$. Using (1) and (2) it is easy to confirm that $\psi_R = 0$ and $\psi_F = \frac{2}{3}$, given a rainy or fair prediction, respectively. The consumer now conditions her action on the weather prediction, and this allows her to achieve an higher level of utility than in the absence of news.

Not all signal structures increase expected utility for the consumer. To see this, continue with the example from Figure 3, but consider the signal structure S' given by $(p, q) = (\frac{1}{2}, 1)$. Now posteriors generated are $\psi_R = \frac{1}{3}$ for a rainy signal and $\psi_F = 1$ for a fair signal. Here, regardless of which signal she observes, she will still go to the game. Therefore, her expected payoff from this information structure is the same as with no information. Clearly, then, our consumer prefers information structure S to S', since the former signal structure allows her to condition her actions and thereby increase the expected payoff. We now formalize these intuitions.

To explore circumstances under which a consumer will condition actions on signals, we observe, first of all, that expected utility for a consumer who conditions is

$$u = \rho p H + (1-\rho)(1-q)L. \tag{3}$$

Next notice that if the consumer does *not* condition, then her expected utility will be either (i) $\rho H + (1-\rho)L$ (if she always goes to the game), or it will be (ii) 0 (if she never goes to the game). So the consumer will want to condition behavior if expected utility (3) is greater than both object (i) and object (ii). This gives us two "conditioning constraints," only one of which is typically binding. Constraints (i) and (ii) can be characterized, respectively, as follows: in (p,q) space, signals must be above the lines,

$$q = -\frac{\rho H}{(1-\rho)L} + \frac{\rho H}{(1-\rho)L}p, \tag{4}$$

and

$$q = 1 + \frac{\rho H}{(1-\rho)L}p. \tag{5}$$

Notice that the first of the conditioning constraints passes through the point (1,0) while the second conditioning constraint passes through the point (0,1). Also observe that the constraints have the same slope (which is negative, as $L < 0 < H$), i.e., they are parallel. Thus, as noted, only one of the constraints is typically binding.

As long as there exists an information structure in the feasible set $\Phi(p,q)$ that meets the conditioning constraints, the consumer will have higher utility if she uses that information to condition her behavior. In this case, using (3), the agent has indifference curves defined over (p,q) pairs that satisfy

$$q = \frac{(1-\rho)L - \bar{u}}{(1-\rho)L} + \frac{\rho H}{(1-\rho)L}p \tag{6}$$

for various levels of $\bar{u}$. The slope of an indifference curve is negative. Indeed, the slope of an indifference curve is the same as the slope of each conditioning constraint.

Figure 4 illustrates. In this example, the point of tangency of the indifference curve and the technology boundary is an optimum if the conditioning constraints are met. It is easy to see that this point, (p^*, q^*), lies to the right of the outermost conditioning constraint (which is this example is the constraint that passes through (1,0)). Thus we do indeed have an optimal outcome for the consumer, in which she conditions her behavior on information.

Notice that the slope of the indifference curve becomes steeper as H, L, or ρ increase. The following proposition summarizes the resulting comparative statics.

Proposition 1 *If the decision-maker strictly prefers at least one signal structure to another, then p^* is increasing in H, L, and ρ, and q^* is decreasing in H, L, and ρ.*

2.3. **Demand-Driven Bias.** In order to examine information provision in equilibrium for our environment, suppose that on non-game days consumers have a preference for a neutral (i.e., $p = q$) information structure. On game days, consumers have the action set described in the previous subsections available to them. For the sake of simplicity we will consider a single representative consumer.

A monopolist firm can choose any feasible information structure on a given day (and the chosen information structure can vary by day), which it then sells at price r to consumers. The firm's profits are equal to r if it sells information and 0 if it does not (since the cost of of producing information is 0). Because the firm profits by selling the consumer her optimal signal, the firm will always provide the structure (p^*, q^*) discussed in the previous subsection.

The following proposition summarizes the fact that the information structure will generally be biased on game days (in the sense that only one particular set of parameter values generates a preference for unbiased signals). Furthermore, observed bias will change with the underlying parameters of the model in predictable ways: bias $\mathbb{B} = p^* - q^*$ is increasing in H, L, and ρ.

Proposition 2 *If any information is purchased on non-game days, the information structure produced will exhibit no bias. On game days, if any information is purchased, the information structure will generically be biased, and the bias will be increasing in H, L, and ρ.*

2.4. **Supply-Driven Bias.** In this section, we model bias as being driven by supply-side considerations, i.e., the firm has an incentive to alter consumers' beliefs. We assume that the newspaper receives profits from the number of agents taking action a. The more agents that take action a, the higher the firm's payoff is. For example, a newspaper may be paid by a sports team to increase consumers' beliefs about the likelihood of fair weather. In this case, the newspaper would benefit the higher $\rho p + (1-q)(1-\rho)$ is. We consider a more general benefit function, $b(p, 1-q, L, H, \rho)$, where $b_i \geq 0$ for all $i \in \{p, 1-q, L, H, \rho\}$, $b_{ij} \leq 0$ for all $i, j \in \{p, 1-q, L, H, \rho\}$, and b is concave. The negative cross partials, and overall concavity capture two things: first, the fact that individuals with the lowest willingness to pay require the highest values of (p, q, L, H, ρ) to attend the game; and, second, that the baseball team (and so the firm) would only benefit to the extent that the stadium is not already selling out.[9] The firm also faces a cost of biasing information, $c(p-q)$, where c is convex in the amount of bias, and has a minimum at 0. This captures the fact that there are likely consumers who would prefer that the news not be biased.

The firm's marginal benefits are falling in the current amount of bias. Moreover, the marginal benefits to bias are also falling in ρ, H and L. Because the marginal costs are rising in the amount of bias, there will be less bias when ρ, H or L increase. This is true for any optimum of the firm.[10]

Proposition 3 *When bias is supply-driven bias, if information is purchased on non-game days, the information structure produced will exhibit no bias. On game days, if any information is purchased, the information structure will always be positively biased and the bias will be decreasing in H, L, and ρ.*

2.5. **Welfare.** The impact of media bias on welfare can be ambiguous. Bias in and of itself does not harm consumers payoffs. If distortions are occurring for supply-side reasons, then consumers are worse off; they would prefer less bias. But, if distortions are generated on the demand side, then the bias is in fact optimal.

[9]In the attendance records we have available, we find the stadium occasionally sells out.

[10]However, the firm may have multiple optima, not all of which may be on the technology constraint. In fact, only one optimum can be on the technology constraint. Additional restrictions can be imposed so that the optimum is unique, including conditions that imply that the firm always prefers a signal on the technology constraint, but we do not do so here.

2.6. **Testable Implications.** The two models in this section have predictions not only about the direction of bias, but also about how the bias is affected by changes in the parameters. We will use these predictions to test the models in in the next section. The following table summarizes our theoretical results.[11]

Model	Direction of Bias	Change in Bias as H Increases	Change in Bias as L Increases	Change in Bias as ρ Increases
Demand-Driven Bias	Either positive or negative	Increases	Increases	Increases
Supply-Driven Bias	Always positive	Decreases	Decreases	Decreases

In the empirical section we will also consider the predictions of other models of media bias, such as Gentzkow and Shapiro (2006) and Mullainathan and Shleifer (2005).

3. An Empirical Test: Weather Reporting by the *New York Times,* 1890–1899

In order to test predictions about media bias we use a novel data source — the daily weather reports given in the *New York Times* during the late nineteenth century. We test whether the *Times* was more likely to provide biased weather predictions on days when the New York Giants, the local professional baseball team, were scheduled to play a baseball game in the Manhattan Polo Grounds. To establish the plausibility of such bias, we begin with some observations about the historical context.

First, weather reporting, although improving, was still in a relative poor state. By the 1890s the practice of using almanacs and astrology to make weather predictions had practically disappeared, and the telegraph had made synoptic meteorology, or weather maps, a tractable method for making weather prediction. This new technique in weather forecasting heightened the usefulness and popularity of weather predictions in newspapers (Nebeker, 1995). Still, weather reporting was a highly subjective practice that relied on rough empirical rules, with a limited understanding of scientific facts and mathematical modeling. Nebeker

[11]The insights of the particular models presented here are in fact more general. For example, any model that predicts that the information provider wants to encourage attendance at the game will generate comparative static predictions similar to our model of supply driven bias.

indicates that even at the beginning of the 20th century, weather reporting was more of "an art rather than a science."

Second, during this period, the *New York Times* was a relatively small player in the the New York media market. There were over 10 newspapers competing for readers using morning, afternoon and evening editions, with prices for weekday editions typically less than three cents. The largest newspaper had a circulation of around 300,000 during the late 1800s.

During the period of analysis, 1890–1899, the *New York Times*, a morning newspaper, was in a transitional phase. In the spring of 1896, Adolph Ochs took over as owner (from Henry Raymond) of what was, at the time, a struggling and politically motivated newspaper. He instituted major changes, among them the change in weather reporting. He also coined the phrase "All the News That's Fit to Print." Ochs lowered the price of the paper from 3 cents to 1 cent, which perhaps was one reason for an increase in the circulation of the *Times*, which tripled within two years from 26,000 to 76,000.[12] Although we have some information about yearly average circulation, we do not have daily numbers, and so do not use circulation data in our analysis. The growth in circulation would also lead us to suspect a weakening in the motives to bias weather, as the circulation would grow beyond the local Manhattan market (where the Giants were located).[13]

As mentioned, under Ochs' new administration changes were made to weather reporting. Every morning the *Times* would provide the weather report for the rest of the day. Prior to June 16, 1896 the weather report quoted local weather predictions from the U.S. Weather Bureau. Although the Bureau was based in in Washington D.C., the Bureau had an office in Manhattan, and produced weather predictions for most major cities. After June 16, 1896 the *Times* switched its editorial policy and began producing in-house morning weather reports, as is indicated by the introduction of a separate column "Probabilities for the Day" on the front page of the paper. After the switch, forecasts no longer cited the Bureau as a source. We verified that the reports were in general not the same as other New York newspapers. The predictions by the *Times* during this period were simple. There was generally either a prediction of fair weather, or rain of some type (e.g., "showers" or "rain"). No probabilities

[12]The new weekday price of 1 cent is approximately 25 cents in today's terms. The Sunday edition was more expensive.

[13]In fact, bias does indeed empirically disappear. We collected data for 1910 and found no statistically significant evidence for bias in weather reporting associated with home game days.

were given, nor any anticipated amount of precipitation, nor typically any information about the timing of precipitation.

During the late nineteenth century, the *Times* had a sports section that provided extensive coverage of the National League and particularly the New York Giants, a major league baseball team that played in the Polo Grounds, in upper Manhattan. This was a period of rapidly growing popularity of baseball. New York had two major league teams, the Giants and the Brooklyn Bridegrooms (which later became the Dodgers), that were part of the National League, a 12 team league at the time. There were no other professional sports teams in the New York area at the time. The *Times* apparently catered to a Manhattan readership, as they provided extensive coverage of the Giants. The New York Giants were owned by businessmen Andrew Freeman and John Brush. Although we do not have data on price variation across years and games for tickets, we do know that a ticket price was 50 cents at the beginning of the time period we consider.[14] We cannot find any formal association between the *New York Times* and the Giants, nor association between their owners.

3.1. **Data.** To test the comparative statics and level of bias, we use a dataset constructed as follows: First, we coded home games during National League play for the Giants (typically late April to early October) for the 1890 through 1899 seasons.[15] The coding system is binary, with a 1 indicating a home game for the Giants and a 0 representing a non-home game (which includes away games). These records were taken from *Baseball Reference*, which has extensive coverage of historical records, including rankings and win-loss records.

Another part of the dataset is gathered weather reports from the *Times* during the period mentioned above. The *Times* weather report, indicating weather for the day, was published in the morning paper. We recorded an indicator variable — 1 for a fair prediction and 0 for a prediction which mentioned rain. As noted previously, weather reports were sufficiently simple so that there was a clear indication of how to code the data. Consider, as an example, the following weather forecast, given on the 17th of July, 1896:

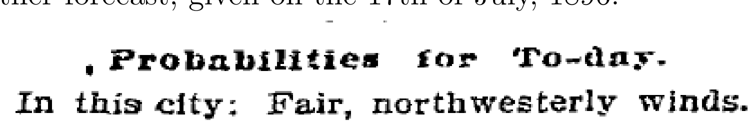
Probabilities for To-day.
In this city: Fair, northwesterly winds.

[14]Approximately 10 dollars in today's terms.

[15]The first game of the season was between April 15 and 28, while the last day was between September 26 and October 15.

Such a report is recorded in the data as a "fair" prediction. We were able to collect these data for all but four of the 1701 days in our study period.

The actual weather outcome was determined from records of the *National Climate Data Center of the National Center for Atmospheric Research*, which holds historical data from the Weather Bureau. We have daily precipitation levels at the Manhattan reporting station. The day was coded as "fair" if there was no rainfall (again, with 1 as an indicator), and "rainy" if precipitation was recorded (with 0 as the indicator).[16] We have data for all but two days.

For a relatively small number of days we were able to find attendance data from the *New York Times* sports page. The *Times* had a column the day after each home game summarizing the results, sometimes listing attendance. Attendance seems to be rounded to the nearest 500. We also collect the ranking of the New York Giants for each season, again using records from *Baseball Reference*.

3.2. **Empirical Strategy and Results.** Given the historical context, the goal is to test the models discussed in the previous section. After we discuss the results of our analysis, we relate them back to the predictions of the models. Doing so involves two steps. First, we ask whether the model can generate the direction of bias consistent with that observed in the data (e.g., positive bias). If so, we then evaluate comparative static predictions from the model. For instance, we look for correlations in the observed bias and variation in payoffs to attending a game or in priors in the likelihood of rain.

Before we turn to our primary analysis, we provide some summary statistics regarding two empirical objects of interest: *weather realizations* (as recorded by the Weather Bureau) and *weather predictions* (from the *New York Times*).

We begin with a description of weather realizations in Table 1. The top row gives statistics for the full sample (baseball seasons from 1890 to 1899), showing the proportion of days for which the realized weather was *fair* for each of the following subsets of days: first, all days; second, days when the Giants played a home game ("home game days"); third, days when there was not a home game ("non-home game days"); and, fourth, for completeness sake, for games played when the Giants were away ("away games"). As we have emphasized, our

[16] The reporting station was at Central Park in Manhattan, very close to the Polo Grounds and also the the offices of the *New York Times*.

analysis focuses on the impact of policy changes in weather reporting at the *New York Times* on June 16, 1894. So, here we show how realized weather varied before and after that date.

As Table 1 shows, the weather is nicer on days with home games than on days that don't have home games. This is true in both the before period and the after period. There are two possible reasons for this. First, there were rainouts; these days, with no home game played, obviously would be exclusively rainy days. Second, it may be the case that the National League tended to schedule relatively fewer home games for the Giants in traditionally rainy months and relatively more games in fair months. In order to account for this fact, we will attempt to control for differences in priors across different days as a robustness check.

As for *weather predictions* from the *New York Times*, Table 2 provides summary statistics. Over the full sample (baseball seasons, 1890–1899), the proportion of *Times* weather reports predicting fair weather was 0.631, slightly lower than the corresponding realizations for fair weather realizations in Table 1. Given that the realized weather was fairer on home game days than on non-home game days, it is not surprising that predictions were on average fairer on home game days than on non-home game days. Below, we will be assessing bias by asking if, in the period after June 16, 1896, the *Times* tended to over-predict fair weather on days with a home game relative to days with no home game. Thus, it is interesting to note here that in the after period, fair weather predictions occurred with a substantially higher probability on home game days than on non-home game days.

We continue to summarize characteristics of our data by considering patterns of autocorrelation in weather outcomes and weather predictions. We do so in three regressions, reported in Table 3. First, Panel A examines possible autocorrelation between realized weather today and realized weather yesterday. This can be important as consumers might base their beliefs about fair weather on their observations about recent weather (i.e., beliefs might be formed using information beyond what the weather report predicts). We note that fair weather yesterday is associated with a significant increase (16 percentage points) in the probability of fair weather today. Panel B considers the same type of autocorrelation but looking at the *predicted weather* by the *New York Times.* Here we find a lower level of autocorrelation, just as we would expect if the weather predictions are a noisy guess about realized weather. Finally, Panel C provides information about the informativeness of the weather predictions above and beyond yesterday's weather realization. Put another way, if an individual observed yesterday's weather, would she still learn something from reading the *New York Times*? We find that both yesterday's weather and the weather prediction are informative about today's weather.

With these characteristics of the realized weather and predicted weather in mind, we now turn to our primary analyses, in which we evaluate the accuracy of predicted weather, conditional on the realized weather. Our goal here is to estimate the accuracy of *New York Times* weather predictions within the framework of the theory posited above. Thus we are interested in an empirical assessment of the probability of correctly predicting fair weather when in fact the weather will be fair (p), and the probability of correctly predicting rainy weather given the weather will be rainy (q).

Table 4 presents the observed values of p and q for various subsets of days. We note three features of the data. First, over the entire sample, the accuracy of predicting fair weather is somewhat higher than the accuracy of predicting rainy weather (0.754 compared to 0.624). This is true before the change to in-house weather predictions, on June 16, 1896, as well as after, as can be seen by comparing the second and fifth rows of the Table 4. Second, in the before period, when the Weather Bureau was producing weather reports, predicting accuracies p and q were quite similar on non-home game days and game days. Third, in contrast, in the after period p and q follow a pattern that is strikingly consistent with our model of optimistic bias on home game days relative to non-home game days. Thus p is higher on home game days than on non-home game days, and the converse is true for q. We note, in addition, that on non-home game days the *Times* used a somewhat different weather prediction policy (i.e., with lower p and higher q) than the Weather Bureau used in the before period.[17]

Figure 5 shows the information structures we observe: "before home" and "before non-home" (i.e., home game days and non-home game days before June 16, 1896) information structures look similar, while the corresponding "after home" and "after non-home" structures look very different. Consistent with the observations we have just made, this figure

[17]We are not certain why the *Times* tended to adopt a more pessimistic weather reporting policy on non-home game days than the Weather Bureau. Our theory focuses only on how the *Times* might adopt differing prediction policies on home game days compared to non-home game days, and that also is the focus of our empirical work below.

shows that although the before period weather reports for home game days were more Blackwell informative than the before period non-home game days reports, these differences were very slight. Moreover, no other rankings by Blackwell informativeness exist. It does not seem to be the case that the *New York Times* was receiving more Blackwell informative signals and then garbling them in two different ways on different days after the switch. Also, there appears to be a clear trade-off between p and q; high values of one accuracy are associated with lower values of the other accuracy. Finally, the figure shows that in the after period, weather reports on home game days are more optimistic than on non-home game days.

We now turn to a statistical test of our theory. Our primary question is: When the *New York Times* produced its own weather reports, was there "excess bias" on home game days relative to non-home game days, i.e., is $\Delta_{\text{After;Home,Non-Home}} = \mathbb{B}_{\text{Home}} - \mathbb{B}_{\text{Non-Home}} = (p_{\text{Home}} - q_{\text{Home}}) - (p_{\text{Non-Home}} - q_{\text{Non-Home}})$ positive? Using estimates from Table 4 we find

$$\Delta_{\text{After;Home,Non-Home}} = 0.291. \tag{7}$$

The standard error for this estimate of 0.087.[18] The t statistic is 3.48. We reject that $\Delta_{\text{After;Home,Non-Home}} = 0$ at the 0.001 level. Thus we have clear evidence of excess bias in the after period. In terms of our model, the game-day bias is optimistic; the *Times* tended to over-predict fair weather on home game days.

Next we try a counter-factual analysis. A potential concern with our finding of excess game-day bias in the after period is that it is being driven the fact that weather on home game days generally differed somewhat from weather on non-home game days (a fact documented in Table 1 for both the before and after period). This concern would be ameliorated if excess bias does not appear in the before period, when the *Times'* weather predictions came from the Weather Bureau.[19] In fact, we find that

$$\Delta_{\text{Before;Home,Non-Home}} = -0.002. \tag{8}$$

[18]Given that outcomes shown in Table 4 are binomial, the standard error is calculated as follows: Sample sizes to estimate p_{Home}, q_{Home}, $p_{\text{Non-Home}}$, and $q_{\text{Non-Home}}$, respectively, are 185, 58, 233, and 155. So the standard error is

$$\sqrt{\frac{0.800(1-0.800)}{185} + \frac{0.500(1-0.500)}{58} + \frac{0.674(1-0.674)}{233} + \frac{0.665(1-0.665)}{155}} = 0.087.$$

[19]To be clear, we believe our theory potentially pertains when the *Times* makes in-house predictions, but make no claims one way or another about it's applicability in the before period.

Excess home game bias is estimated to be very close to zero. This estimate is reasonably precise; the standard error is 0.065 (using the same approach as in footnote 17).

We can shed further light on our main results with regression analyses. Our baseline regressions, presented in the first column of Table 5 follows a differences-in-differences approach. Specifically, we estimate

$$\mathbb{D}_{\text{PredictCorrect}} = \beta_0 + \beta_1\mathbb{D}_{\text{Fair}} + \beta_2\mathbb{D}_{\text{Home}} + \beta_3\mathbb{D}_{\text{Fair,Home}} + \epsilon. \tag{9}$$

In this regression, $\mathbb{D}$ are dummy variables. The dependent variable, $\mathbb{D}_{\text{PredictCorrect}}$, is 1 if the predicted weather was correct and 0 otherwise. $\mathbb{D}_{\text{Fair}}$ is 1 if the day was fair and 0 otherwise; $\mathbb{D}_{\text{Home}}$ is 1 if the Giants were playing at home, 0 otherwise; and $\mathbb{D}_{\text{Fair,Home}}$ is 1 for a fair day with a home game and 0 otherwise. The estimated coefficients in this regression are related to the p and q estimates given in Table 4: β_0 is an estimate of $q_{\text{Non-Home}}$, $\beta_0+\beta_1$ is an estimate of $p_{\text{Non-Home}}$, and so forth. It can be shown that $\beta_3 = (p_{\text{Home}} - q_{\text{Home}}) - (p_{\text{Non-Home}} - q_{\text{Non-Home}})$; coefficient estimates of β_3 for the after and before periods correspond with estimates of excess bias given in (7) and (8) respectively.

Column (2) of Table 5 shows coefficient estimates of these same regressions also including year, month, and weekend fixed effects. The year effects allow for the possibility that accuracy varies over time. We have seen that the probability of fair weather varies by month, so we include month fixed effects in case this affects results. The inclusion of weekend indicator variables allows for the possibility that the bias was primarily for weekends and not home games more generally. Panel B shows that our key inference about excess bias in the after period is unchanged when we include these fixed effects; estimated excess bias (the coefficient on $\mathbb{D}_{\text{Fair,Home}}$) is about the same as in our baseline regression. For completeness sake, Panel A shows the result of this same regression in the before period — showing that inclusion of fixed effects does not alter our basic inference that there is no excess bias in the before period.[20]

[20]Finally, we estimated our baseline regression, reported in column (1) of Table 5, Panel B, but use a definition of "home" that includes only home games that were scheduled at the beginning of the season. (National League schedules were typically released in February or March prior to the season, and were published in the *New York Times* and other newspapers.) This definition thus includes home games that were rained out and excludes home games that were make-up games due to rain-outs. The basic inference is quite similar as with games actually played: The estimated coefficient on $\mathbb{D}_{\text{Fair,Home}}$ in the after period is 0.240 with a standard error of 0.079. In contrast, in the before period the estimated coefficient on $\mathbb{D}_{\text{Fair,Home}}$ is 0.005 with a standard error of 0.060.

We next try an alternative approach, in which we specify a regression in which the dependent variable is the *Time's* weather *prediction* for a given day — 1 for fair and 0 for rain. We continue to use a differences-in-differences design:

$$\mathbb{D}_{\text{PredictFair}} = \alpha_0 + \alpha_1 \mathbb{D}_{\text{Fair}} + \alpha_2 \mathbb{D}_{\text{Home}} + \alpha_3 \mathbb{D}_{\text{Fair,Home}} + \epsilon. \tag{10}$$

Again, regression coefficients correspond to p and q for various subsets of days (as reported in Table 4), although the relationship is different than in our previous regression. Here α_0 estimates the probability that the *Times* predicts fair weather on a non-home day that is rainy, so $\alpha_0 = 1 - q_{\text{Non-Home}}$. Similarly, $\alpha_0 + \alpha_1 = p_{\text{Non-Home}}$, $\alpha_0 + \alpha_2 = 1 - q_{\text{Home}}$, and $\alpha_0 + \alpha_1 + \alpha_2 + \alpha_3 = p_{\text{Home}}$. It can be shown that excess bias is estimated by $2\alpha_2 + \alpha_3$. Column (1) of Panel B of Table 6 gives our baseline results; as expected we estimate $2\hat{\alpha_2} + \hat{\alpha_3} = 0.291$ (and as indicated by the F statistic reported in the note to Table 6, we reject the hypothesis that excess bias is 0).

Our reason for estimating regression (10) is that it allows us to ask what happens to our inference when we attempt to control for individuals' prior beliefs about the probability of fair weather on any given day. To do so, we first form estimates of individuals' prior beliefs by estimating a probit regression, in which the dependent variable is the weather realization, fair (1) or rainy (0) for a given day, and the independent variables are the amount of precipitation each of the previous five days. We then use results from this equation to generate the prior probability of fair weather, denoted *Daily Prior for Fair Weather*, using the information that would have been available to individuals before reading the weather report. One plausible theory of bias would be that there is an issue with prediction technology such that it forces p higher and q lower as the prior increases. This could account for the bias we observe. So we include this estimated prior as a control in regression (10).

Table 6 shows the coefficients from the probit specification used to construct priors. Column (1) is our basic specification. Column (2) adds week fixed effects for the baseball season, as a flexible way of picking up a possible persistent seasonal pattern in rain from late Spring through early Fall.[21] Finally, as a simple check, in column (3) we add a dummy variable for a game being scheduled that day, and discover that it is not statistically significant. We also

[21]Specifically, we formed 26 dummy variables: for April 15-21, April 22-28, and so forth, up through the second week of October.

notice that estimated coefficients on lagged daily rainfall are extremely similar across the specifications. The first of our specifications is the simplest, and it has the advantage that it is formed strictly with information that was available before individuals would be forming priors about a given day's weather.[22]

Estimates of regression (10) with and without the priors are reported in Table 7. Panel A gives results for the before period, simply as a counter-factual. Our key interest is Panel B, which gives analysis for the after period. We see that coefficients in our regression change very little across the two columns (other than the intercept) when we include the prior. When we include the prior in the regression, estimated excess bias is $2\hat{\alpha_2} + \hat{\alpha_3} = 0.248$ (and we reject the hypothesis that excess bias is 0). In short, it seems that our key inferences about *Times* weather reporting biases are not driven by differences in daily priors about the weather. As for Panel A—the counter-factual analysis for the before period—in both specifications we fail to reject the hypothesis that $2\alpha_2 + \alpha_3 = 0$.[23]

We can use our constructed prior to address two additional questions: First, whether weather predictions were informative beyond priors. Second, whether the game schedulers had information in excess of that of consumers regarding the weather, and if this could account for the fact that home game days were sunnier on average than non-home game days. (Alternatively, it may be the case that the National League tended to schedule relatively fewer home games for the Giants in traditionally rainy months and relatively more games in fair months.) Given our constructed daily priors for fair weather, we ran a regression estimating whether there is additional information provided by weather predictions, and the fact that a home game may be scheduled, in addition to the daily prior. We regressed the realized weather on the prior, the weather prediction, and dummy variables to indicate whether or not there was a home game scheduled. We find that the coefficients on the daily prior, 0.564 (s.e. = 0.164), and the weather prediction, 0.332 (s.e. = 0.028), are significant, indicating that both provide information regarding weather realizations. However, the coefficient on

[22]In contrast, in the other two specifications, priors are formed using weather outcomes that have not yet occurred.

[23]Finally, we tried our analysis but with priors formed with the other two specifications given in Table 6. When we use the specification from column (2) to construct our prior we estimate bias to 0.226, and $F(1, 611) = 7.14$. When we use the specification from column (3) estimated bias is 0.190 and $F(1, 611) = 4.52$. In both specifications we reject the null hypothesis of no bias.

scheduled games is insignificant, indicating that knowing whether it was a home game day or not would not have further changed individuals' beliefs.

As a final robustness check, we undertake an additional counterfactual analysis, using weather reports from a different New York newspaper, the *Wall Street Journal.*[24] In the late nineteenth century, the *Journal* primarily covered financial topics, including American and international business and economic news; the paper's claim to fame was the construction and reporting of the Dow Jones "Average," one of the first index measures of stock prices on the New York Exchange. The paper did not typically cover sports, and so probably did not attract readers who wanted to learn news about the New York Giants. Like the *Times,* the *Wall Street Journal* provided a rudimentary weather forecast.

With this in mind, we coded the weather forecasts from the *Wall Street Journal* in the same fashion as for the *New York Times.* Our assumption is that the typical reader did not read the weather reports with the intention of using the information to attend games. Thus we would not expect that the *Journal's* reports to exhibit the same bias as the *Times*.

The *Journal* was an afternoon paper, which produced forecasts covering the next day; this was taken into account when analyzing predictions. We assembled data for the baseball seasons during the after period, 1896–1899. We collected *Journal* data for a total of 236 predictions; this is all that is available in digitized records, and certainly doesn't include the entire after period. For those 236 dates, one observation is missing in the *Times* records. We drop this day to focus only on dates when both papers made predictions, leaving a sample of 235 days.[25] For these dates we find that the *Times* produced excess bias of $\Delta = 0.365$ (which is not too different than for the entire sample), while the *Journal* had excess bias $\Delta = 0.017$, which is close to 0, as we would have expected.[26]

We next turn to some of the other predictions from our theory. Recall that our theory provides some reason to think that the bias will depend on the prior for fair weather (ρ) in the model.[27] We can look for such a relationship by looking at month-to-month variation in

[24]This paper was established in 1882 by reporters Charles Dow, Edward Jones and Charles Bergstresser as a letter that provided financial news, and was converted by Dow Jones and Company into the *Wall Street Journal*, which distributed it's first issue in 1889.

[25]The two papers agree on 166 observations (70.6% of the 235).

[26]These estimates of the *Times'* excess bias and *Journal's* excess bias are not statistically different at the 0.05 level. (The sample size is small, so the test is under-powered.)

[27]If individuals' beliefs are rational, we would expect a one percent increase in beliefs about the probability of fair weather to lead to a one percent increase in the probability of fair weather being realized. To be clear,

the prior. We start by using the realized weather in the sample to construct the probability of fair weather in any given month (as in Table 1). We denote these monthly priors $\hat{P}_{\text{Fair,Month}}$. (We drop April and October for the purpose of constructing these priors as those months have few observations.) Figure 6 shows that for the post-June 1896 period, there is a clear negative relationship between these priors and the average excess game day bias in that month, $\Delta_{\text{AfterHome,AfterNon-Home,Month}}$.[28] In Figure 6 the relationship appears to be roughly linear; we thus fit a line to these observations,

$$\Delta_{\text{AfterHome,AfterNon-Home,Month}} = \beta_0 + \beta_1 \hat{P}_{\text{Fair,Month}} + \epsilon,$$

where $\Delta_{\text{AfterHome,AfterNon-Home,Month}}$ is, for a given month, the difference in bias between home game days and non-home game days. Note that this regression is not simply picking up effects of summer relative to non-summer months, nor months early in the season relative to late in season (since fair weather is not perfectly correlated with either of these). We estimate β_1 to be -3.68 (with a standard error of 1.02), which is a statistically significant relationship at the 0.05 level. A similar analys, but conducted using the constructed daily priors discussed above, generates qualitatively the same results — $\Delta_{\text{AfterHome,AfterNon-Home}}$ is falling in the prior.

Next recall that our theory suggests that bias will be related to the value fans place on being at a game when the weather is fair (H) and on loss from attending a game when it is rainy (L). Intuitively, these values might be related to the relative success of the team. As seen in Figure 7, in fact relative bias was highest in 1897, when the Giants had a relatively good season; was lower in 1898 and 1896, when the Giants finished 5th (of the 12 teams in the league); and was lower yet in 1899 when the team finished near the bottom of the

these beliefs would be individuals' priors and posteriors, not the weather predictions themselves (although posteriors would be constructed using priors and the weather predictions). We conduct such a test, first regressing realized weather on our constructed daily priors, and second regressing realized weather on a constructed measure of a daily posterior, which we create by applying Bayes Rule, using the prior and weather predictions, allowing for different values of p and q across different subsets of days (the interaction of before, after, home and not home). Both regression coefficients are tightly estimated to be 1, in line with the rational expectations literature.

[28]One might be concerned that individuals' priors for game days and for non-game days within a month differed. Instead of lumping game days and non-game days within a month together, one could instead compute priors for game days and non-game days for each month separately, and then estimate the effects of priors on excess bias. Such an analysis again indicates that excess bias is falling (at a significant level) in the prior.

standings.[29] Since in the figure the relationship appears to be roughly linear, we attempt to fit this relationship (for the period post June 16, 1896) using the regression,

$$\Delta_{\text{Year,AfterHome,AfterNon-Home}} = \beta_0 + \beta_1 R_{\text{Year}} + \epsilon,$$

in which $\Delta_{\text{Year,AfterHome,AfterNon-Home}}$ is the difference in bias in a given year between home game days and non-home game days, while R_{Year} is the end of season rank of New York Giants in that year. We estimate β_1 to be $-.092$ (with a standard error of .007), which is a statistically significant relationship at the 0.01 level. A similar analys, but conducted using the the daily rankings of the New York Giantsgenerates qualitatively the same results — $\Delta_{\text{AfterHome,AfterNon-Home}}$ is falling in the rank, or in other words, increasing in the relative ability of the baseball team.

3.3. **Further Analysis: Attendance.** In order to get a better sense of how predicted and realized weather might affect the actions of readers and fans, we analyze the effects of weather predictions and realizations on the New York Giant's home game attendance. It also seems plausible that the performance of the team should affect attendance of the game when the team is doing well, and so we also look at the impact of team ranking on attendance. We estimate the following regression for the period after June 16, 1896:

$$A = \beta_0 + \beta_1 \mathbb{D}_{\text{PredictFair}} + \beta_2 \mathbb{D}_{\text{Fair}} + \sum_i \beta_i \mathbb{D}_{\text{Rank}} + \epsilon,$$

where A is the game attendance, $\mathbb{D}_{Rank}$ are dummy variables for the ranking of the team (either rank 3, 5, or 10 depending on the year) and the rest of the variables are as described previously. One potential issue with this analysis is that we only have attendance records for a subset of games.[30] However, we have no reason to believe that the subset of games available to us is biased. Table 8 reports the results of this regression (with rank 5 being the omitted category). Surprisingly, predicted weather is not a significant predictor of attendance. Even more surprisingly, realized weather is not correlated with attendance either. The only significant variable is the yearly rank of the New York Giants. It seems to be the

[29]The difference in estimated excess bias for the year with the team finished 3rd (1897) and year it finished 10th (1899) is statistically significant at the 0.05 level. (No other differences between years were statistically significant.) Of course the year-to-year differences in excess bias may have had little to do with the ranking. Our only point here is that this evidence is not what we would have predicted with our model.

[30]42.0% of the attendance records are complete. The *Times* provided a report the day after every home game, only some of which included attendance records. Efforts are being made to collect more data.

case individuals did not condition their actions either on information regarding the weather nor the weather itself.[31]

4. Discussion

4.1. Relating Evidence and Theory. In Section 2 we derived some testable implications of two models of media bias. We revisit them here:

Model	Direction of Bias	Change in Bias as H Increases	Change in Bias as L Increases	Change in Bias as ρ Increases
Demand-Driven Bias	Either positive or negative	Increases	Increases	Increases
Supply-Driven Bias	Always positive	Decreases	Decreases	Decreases

We believe that it is reasonable to use the monthly average probability of a rainy day as a proxy for consumers' priors about fair weather in a given month. Furthermore, we believe that if the New York Giants were doing better, it represented an increase to H, L or both (i.e., it is better to watch a good baseball team rather than a bad baseball team when the weather is sunny).

Given these interpretations, recall that we found that as the prior increases the bias shrinks. This is consistent with supply-driven bias, but not with demand-driven bias. In contrast, we also found that as the team did better, the bias grew, which is inconsistent with supply-driven bias, but is consistent with demand-driven bias. In fact, neither of the models correctly predict both observed comparative statics.

Casting further doubt on both the demand- and supply-driven stories is our evidence that the predicted weather does not influence attendance. For example, if the team was paying the newspaper to bias the weather reports in hopes of boosting attendance, they would have presumably quickly learned that this was not effective.

In light of this evidence, it would appear that any alternative explanation for the media bias should not center around models where individuals condition their actions based on the

[31]One might be concerned that our findings are affected by correlation between predicted and realized weather. However, even if we include both regressors separately in the regression they are still insignificant.

weather predictions, but rather where bias occurs even when individuals do not condition their actions on predictions.

4.2. **Alternative Explanations.** Of course, the two models developed above are not the only models of media bias. Here we consider some alternative explanations that could possibly rational the data we observe.

Before we discuss alternative models that already exist in the literature, it is instructive to think about the effects of relaxing our two major assumptions regarding the rationality of consumers. First, that individuals condition their actions on the predictions if optimal. Second, that individuals know (at some level) the values of p and q on non-home game days and home game days, and change their updating rules accordingly. Violating the first assumption would eliminate the prediction of bias in the two models we derive in this paper, not changing our overall conclusions. Violating the second assumption would immediately imply that there would be no benefit for the *New York Times* to biasing their information in order to improve consumers' ability to condition their actions, and so we would not expect to see bias on game days if bias is demand-driven. Moreover, violating the second assumption would not change the comparative static results in the case of supply-driven bias. Thus our conclusions regarding the falsification of both models would remain unchanged.

It is also important to consider explanations that could result from the weather reports being about different 'things' on different days.[32] It could be that on game days the forecast is meant to better predict the weather only during the baseball game, but on non-game days the prediction is for the entire day. This could generate positive bias as we observe in the data. However, if days were particularly rainy, then the bias should shrink relative to non-game days, as it is more likely that on the rainiest days it is also raining during the game. We observe the opposite relationship in the data.

We now turn to other models of media bias in the literature. Gentzkow and Shapiro (2006) develop a model where firms bias news for reputational reasons, i.e., in order to appear informed. This model predicts that firms try to bias their news to match consumers' priors. Of course, as Gentzkow and Shapiro point out, their model is generally more appropriate for situations where individuals cannot quickly and easily verify predictions, unlike weather

[32]We thank David Gill for initially raising this point.

forecasts. And in fact, we do not observe patterns in our data consistent with their model. Gentzkow and Shapiro's (2006) model would predict that in our setting, as consumers' priors increase, the bias should grow. This is because as consumers' priors on fair weather increase, they expect (all else equal) to actually observe more fair weather predictions. In fact, we observe the opposite comparative static in the data.

Mullainathan and Shleifer (2005) model media bias as being driven by consumers' desires to see their beliefs confirmed. Their particular modelling approach is disjoint from our setting, but we can still draw from Mullainathan and Shleifer's intuitions, and consdier a simple adaptation that is suitable in our environment. A consumer's utility is increasing in the overall informativeness of the signal, $p + q$. Utility is falling in $|p - q|$, so that consumers dislike biased news. However, a consumer's utility is also falling in $|\rho - \mathbb{D}_\alpha|$, where $\mathbb{D}_\alpha = 1$ if the realized signal was α (i.e. fair weather) and 0 otherwise. This means that consumers like to see their prior beliefs confirmed by the prediction. Ex-ante, the consumer computes the probability of an α signal using p, q and ρ. Therefore, utility is equal to: $U(p,q) = (p+q) - |p-q| - [p\rho + (1-q)(1-\rho)](1-\rho) - [(1-p)\rho + q\rho](\rho)$.[33] It is relatively easy to check that the slopes of the indifference curves fall (i.e. become less positive or more negative) as priors increase. For example, if $p \geq q$ the slopes of the indifference curves are $\frac{3+4(\rho-3)\rho}{(3+\rho(2\rho-3))^2}$, which negative and falling in ρ. Recall that this is inconsistent with the data.

Drawing on the evidence regarding attendance, it seems likely that bias could be driven by intrinsic preferences over information. In this framework, unlike that developed in Section 2, there are no actions available to the agent. The utility realized in state A is higher than that of state B. Otherwise, the setup is the same. Many models predict a demand for particular types of information even in the absence of the conditioning of actions. These models try to capture some form of belief-based utility, and could possibly rationalize the patterns we observe in the data. Examples include Koszegi and Rabin (2009), Artstein-Avidan and Dillenberger (2010) and Kreps-Porteus (1978).[34] However, it needs to be the case that any model needs to jointly capture: (i) a preference for learning something about

[33] Mullainathan and Shleifer (2005) do not allow for varying accuracies of signals, and so neglect the first term. Moreover, they assume losses are quadratic in $(p - q)$ and in $(\rho - \mathbb{D}_\alpha)$. However, these differences do not change the essential insights in our setting.

[34] Brunnermeier and Parker (2005) have a model of belief based utility, but their results depend on information also having instrumental value.

the weather (i.e. a demand for weather predictions); (ii) a falling demand for bias as the prior increases (at least for some range of priors greater than .5), and;(iii) a rising demand for bias as the payoff to either the high or low state (or both) increase. However, even models with intrinsic preferences for information can have trouble rationalizing the observed data. Existing models, including both those that feature an endogeous reference point and Kreps-Porteus preferences, generate informational demand where the optimal signal structure has a bias that is monotonically increasing in the prior.

Instead, we propose a different model of belief based utility which is a modification of Mullainathan and Shleifer (2005). We refer to our model as wishful thinking. As in Mullainathan and Shleifer (2005) individuals dislike bias; utility is falling in $(p-q)^2$. Individuals like to hear that the weather will be good, but the benefits are concave: utility is increasing in $(p\rho + (1-q)(1-\rho))^{\frac{1}{2}}$. Lastly, individuals prefer to hear the weather is good more if the payoff from good weather is higher (i.e. H is higher). We show that this type of model can generate the patterns we observe in the data (given reasonable parameterizations of ϕ). Individuals' utility, given a signal structure p, q and a prior ρ is $U(p, q|\rho) = H(p\rho + (1-q)(1-\rho))^{\frac{1}{2}} - (p-q)^2$ (we normalize the low payoff to 0) The set of available signals is, as before, $\Phi(p, q) = \{(p, q)|q \leq \phi(p) \text{ and } p + q \geq 1\}$. For simplicity, we will assume $\phi(p) = k - p$, where $2 \geq k \geq 1$, where k represents a technology constraint. Proposition 4 demonstrates that wishful thinking preferences can match the observed patterns in the data.[35]

Proposition 4 *If an individuals has wishful thinking preferences then the optimal signal always has positive bias. Moreover the bias in their optimal signal is decreasing in ρ and increasing in H.*

5. Conclusion

Our data, consisting of local weather predictions from the *New York Times* and realized weather in Manhattan during the years 1890-1899, provide a new setting to try to understand media bias. We find evidence that the *New York Times* biased its weather reports on home

[35]We restrict ourselves here to trying to understand what might happen on game days. Moreover, we only consider the optimal consumer signal. An easy extension, repeating the method used in the proof of Proposition 2, shows that the results extend to information provided in equilibrium.

baseball game days relative to non-home baseball game days when it produced its own weather reports. We find that the excess bias on home game days may fall when the ex-ante probability of sunny weather grows. We also find that the excess bias on home game days may be higher when the baseball team does better. These comparative statics are inconsistent with the rational models of media bias that we set out in our paper. Moreover, in the limited subset of data for which we have attendance data, it does not appear that predicted weather influences attendance, which would indicate that individuals do not condition their actions on the information provided by the weather prediction.

The data has several advantageous features. First of all, because we know both the weather prediction and weather realization, we can construct a measure of absolute (rather than comparative) bias. Second, the realized weather is unaffected by the weather report, and so there are no issues with endogeneity. Third, our data allow us to measure how individuals' posteriors should change as the bias grows or falls.

Although we attempt to test the predictions of some models of media bias, we caution that it is not clear how far the intuitions and results in this paper extend. Our environment features information about outcomes that are easily verifiable within a short span of time. In many situations outcomes will not be verified by individuals receiving information (for example, information about climate change). There are many other situations where the information is about causal claims that involve a stochastic data generating process (for example, that smoking causes lung cancer). In both cases, it is easier for individuals to hold distorted beliefs. In these situations other motivations for media bias may dominate.

We would like to conclude by reflecting on the welfare implications of media bias. If individuals were conditioning their actions on the weather prediction, bias on game days could be welfare improving for individuals who may want to attend the game. In contrast, if media bias was being driven by the *New York Times* wanting to increase attendance, then bias could reduce welfare for fans. Of course, we do not find support for either of these models. In contrast, our evidence points towards belief-based utility being the driving factor behind bias. Our proposed model implies that the observed bias is welfare improving for consumers, even though they do not condition their actions on signal realizations. This is important, because even when consumers can condition their actions, we might observe them preferring signals which are sub-optimal in terms of instrumental value (and so lead to worse material payoffs), but are still utility maximizing.

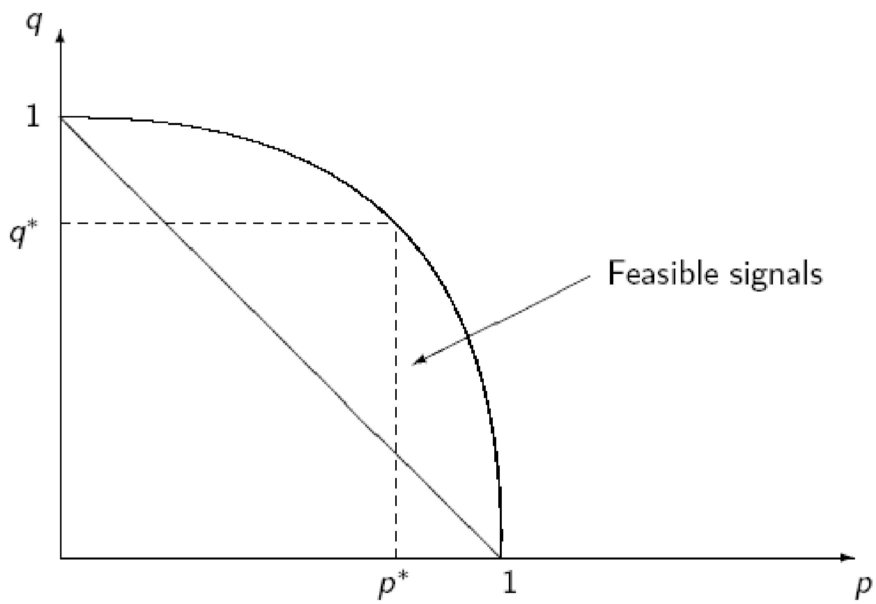

FIGURE 1. The Feasible Set of Information Structures

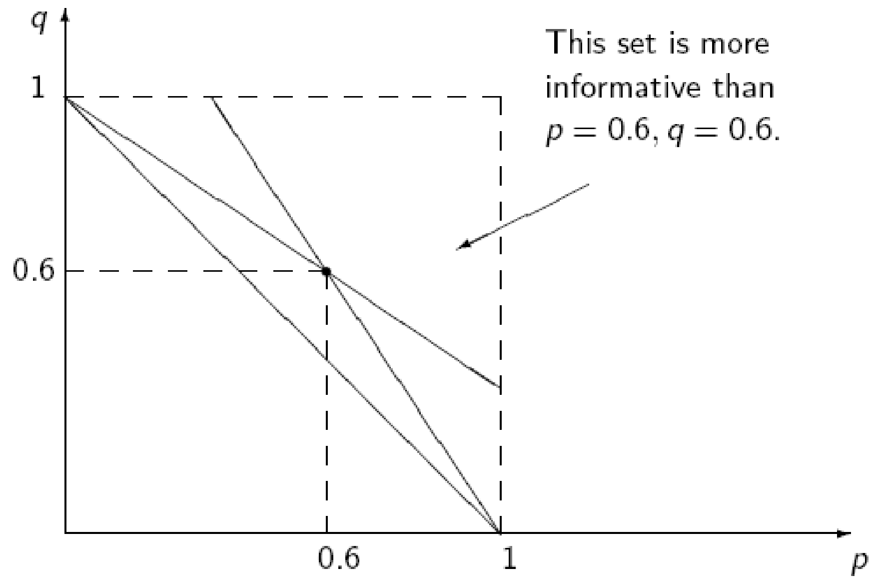

FIGURE 2. Signals that are Blackwell More Informative than (0.6,0.6)

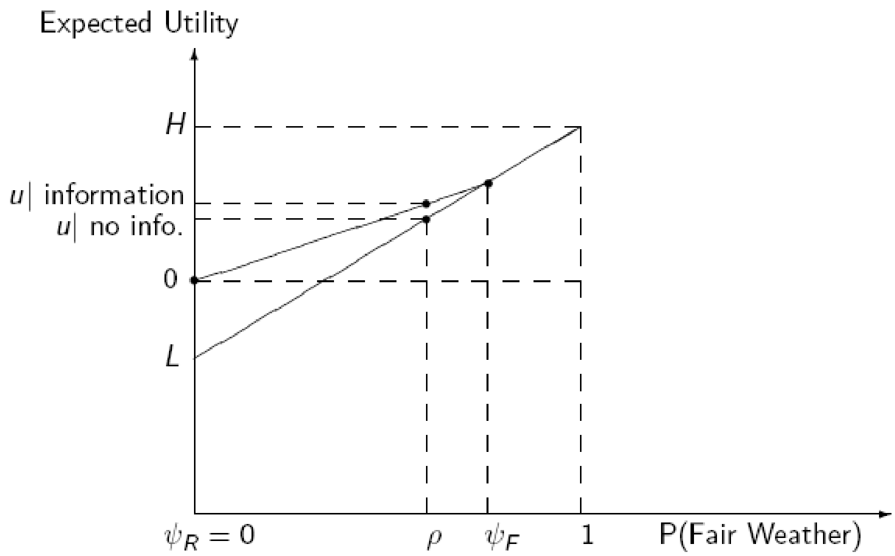

FIGURE 3. Expected Utility Payoffs with and without Information

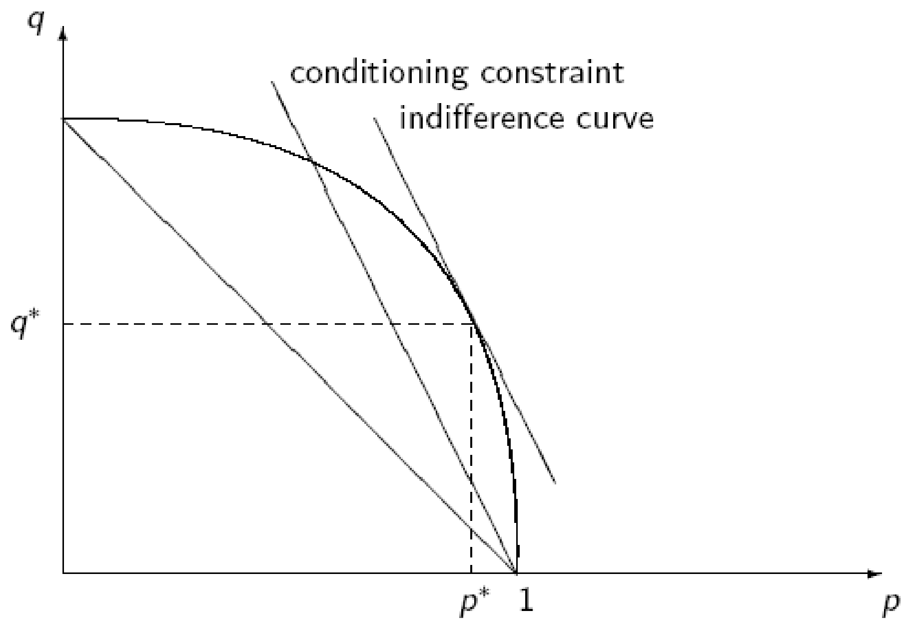

Figure 4. An Example of a Utility-Maximizing News Structure

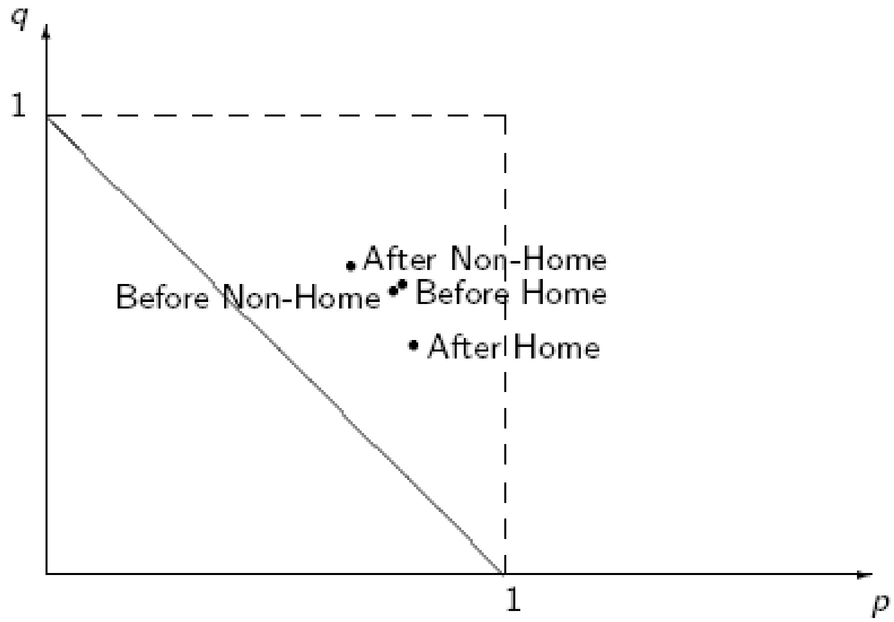

FIGURE 5. Empirically Observed Information Structures

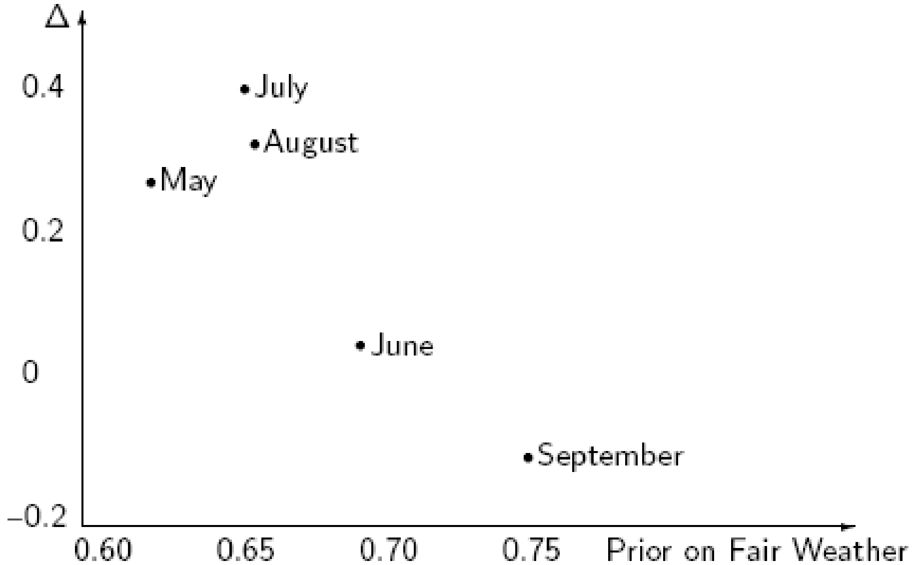

FIGURE 6. Relationship Between Monthly Priors for Fair Weather and Excess Bias Δ

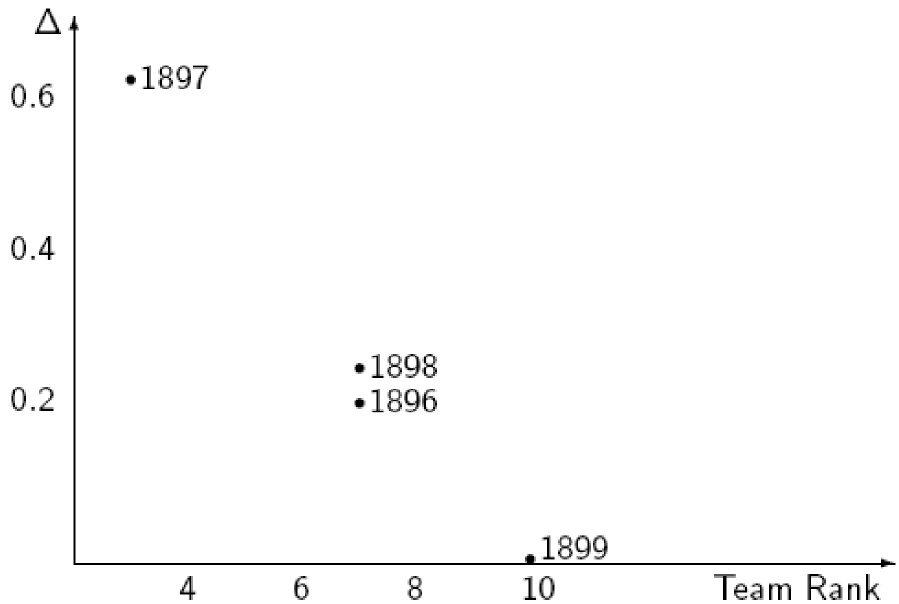

FIGURE 7. Relationship Between Team Ranking and Excess Bias Δ

Appendix A: Empirical Tables

Table 1. Summary Statistics: *Realized Fair Weather*

	All Days	Days with Home Game	Days with No Home Game	Days with Away Game
Full Sample, 1890 to 1899	0.676	0.755	0.626	0.678
Number of Days (n)	(1695)	(650)	(1045)	(642)
Dates before June 16, 1896	0.683	0.752	0.641	0.700
(n)	(1064)	(407)	(657)	(406)
June 16, 1896 and After	0.662	0.761	0.601	0.640
(n)	(631)	(243)	(388)	(236)
By Month:				
April	0.732	0.775	0.711	0.864
(n)	(123)	(40)	(83)	(44)
May	0.612	0.678	0.586	0.608
(n)	(309)	(87)	(222)	(143)
June	0.692	0.772	0.603	0.667
(n)	(299)	(158)	(141)	(81)
July	0.643	0.756	0.612	0.627
(n)	(308)	(82)	(226)	(150)
August	0.650	0.692	0.614	0.673
(n)	(309)	(143)	(166)	(104)
September	0.752	0.845	0.684	0.763
(n)	(294)	(123)	(171)	(97)
October	0.736	0.824	0.694	0.783
(n)	(53)	(17)	(36)	(23)

Note: Authors' calculations, data collected from the *National Climate Data Center of the National Center for Atmospheric Research* for the baseball season, 1890–1899.

TABLE 2. Summary Statistics: *Predicted Fair Weather*

	All Days	Days with Home Game	Days with No Home Game	Days with Away Game
Full Sample	0.631	0.695	0.591	0.615
(n)	(1695)	(650)	(1045)	(642)
Before June 16, 1896	0.643	0.676	0.623	0.643
(n)	(1064)	(407)	(657)	(406)
June 16, 1896 and After	0.612	0.728	0.539	0.568
(n)	(631)	(243)	(388)	(236)

Note: Authors' calculations, data collected from the *New York Times* for the baseball season, 1890–1899. **significance = 0.01. *significance = 0.05.

Table 3. Regression Results: *Relationships Between Realized Weather, Past Weather, and Predicted Weather*

A. Dependent Variable: *Realized Fair Weather* (1 if Fair, 0 Else)	
Constant	0.564**
	(0.020)
Realized Fair Weather Yesterday	0.164**
	(0.024)
n	1685
B. Dependent Variable: *Predicted Fair Weather* (1 if Fair, 0 Else)	
Constant	0.555**
	(0.019)
Predicted Fair Weather Yesterday	0.119**
	(0.024)
n	1685
C. Dependent Variable: *Realized Fair Weather* (1 if Fair, 0 Else)	
Constant	0.385**
	(0.026)
Realized Fair Weather Yesterday	0.116**
	(0.035)
Predicted Fair Weather	0.327**
× Realized Fair Weather Yesterday	(0.028)
Predicted Fair Weather	0.358**
× Realized Rainy Weather Yesterday	(0.037)
n	1685

Note: Authors' calculations, data collected from the *National Climate Data Center of the National Center for Atmospheric Research* for the baseball season, 1890–1899 for Panels A, and also from the *New York Times* for Panels B and C. **significance = 0.01. *significance = 0.05.

Table 4. Summary Statistics: *Times Accuracy of Weather Reports*

	p	q
Total	0.754 (1145)	0.624 (550)
Before June 16, 1896		
All Days	0.768 (727)	0.626 (337)
Non-Home Games	0.758 (421)	0.619 (236)
Home Games	0.781 (306)	0.644 (101)
June 16, 1896 and After		
All Days	0.730 (418)	0.620 (213)
Non-Home Games	0.674 (233)	0.665 (155)
Home Games	0.800 (185)	0.500 (58)

Note: Authors' calculations, data collected from Baseball Reference, *New York Times*, and the *National Climate Data Center of the National Center for Atmospheric Research* for the baseball season, 1890–1899.

Table 5. Regression Results: *Times Correct Weather Prediction*

A. Before June 16, 1896	(1) No Fixed Effects	(2) With Fixed Effects
Constant	0.619**	—
	(0.029)	
Fair	0.139**	0.142**
	(0.036)	(0.036)
Home	0.025	0.034
	(0.053)	(0.054)
Fair × Home	-0.002	-0.006
	(0.063)	(0.063)
Year, Month, Weekend F.E.?	no	yes
n	1064	1064
R^2	0.022	0.038

B. June 16, 1896 and After	(1) No Fixed Effects	(2) With Fixed Effects
Constant	0.664**	—
	(0.037)	
Fair	0.009	0.006
	(0.047)	(0.047)
Home	-0.165*	-0.200**
	(0.070)	(0.071)
Fair × Home	0.291**	0.294**
	(0.083)	(0.083)
Year, Month, Weekend F.E.?	no	yes
n	631	631
R^2	0.033	0.060

Note: Authors' calculations, data collected from Baseball Reference, *New York Times*, and the *National Climate Data Center of the National Center for Atmospheric Research* for the baseball season, 1890–1899. **significance = 0.01. *significance = 0.05.

Table 6. Probit Regression Results: *Construction of Daily Priors for Fair Weather*

	(1) Without Week Fixed Effects	(2) With Week Fixed Effects	(3) With Week Fixed Effects
Constant	0.468 (0.041)	–	–
Rain in Day $t-1$	-0.443** (0.097)	-0.424** (0.099)	-0.420** (0.099)
Rain in Day $t-2$	0.206 (0.109)	0.218* (0.111)	0.222* (0.111)
Rain in Day $t-3$	0.008 (0.102)	0.014 (0.104)	0.020 (0.104)
Rain in Day $t-4$	0.140 (0.107)	0.140 (0.110)	0.147 (0.110)
Rain in Day $t-5$	-0.011 (0.099)	-0.027 (0.101)	-0.025 (0.101)
Home Game Scheduled	–	–	0.131 (0.072)
n	1635	1635	1635
Pseudo R^2	0.012	0.030	0.031

Note: Author's calculations, data collected from *New York Times* (for the scheduled home games), and the *National Climate Data Center of the National Center for Atmospheric Research* for the baseball season, 1890–1899. 26 week fixed effects are included in regression (2). **significance = 0.01. *significance = 0.05.

Table 7. Regression Results: *Times Predict Fair Weather*

A. Before June 16, 1896	(1) Without Prior	(2) With Prior
Constant	-0.025**	-0.440**
	(0.053)	(0.167)
Fair	0.376**	0.353**
	(0.036)	(0.037)
Home	$-0.025^{\dagger}$	$-0.040^{\ddagger}$
	(0.053)	(0.053)
Fair × Home	$0.048^{\dagger}$	$0.070^{\ddagger}$
	(0.062)	(0.063)
Daily Prior for Fair Weather	–	1.241**
		(0.250)
n	1064	1019
R^2	0.129	0.150

B. June 16, 1896 and After	(1) Without Prior	(2) With Prior
Constant	0.335**	-0.277
	(0.037)	(0.192)
Fair	0.338**	0.320**
	(0.047)	(0.048)
Home	$0.165^{*\dagger\dagger}$	$0.121^{\ddagger\ddagger}$
	(0.070)	(0.071)
Fair × Home	$-0.038^{\dagger\dagger}$	$0.006^{\ddagger\ddagger}$
	(0.083)	(0.083)
Daily Prior for Fair Weather	–	0.933**
		(0.285)
n	631	616
R^2	0.129	0.150

Note: Author's calculations, data collected from Baseball Reference, *New York Times*, and the *National Climate Data Center of the National Center for Atmospheric Research* for the baseball season, 1890–1899. **significance = 0.01. *significance = 0.05. Hypothesis test, $2\alpha_2 + \alpha_3 = 0$: $^{\dagger}F(1,1060) = 0.00$ (sig. = 0.98), $^{\ddagger}F(1,1052) = 0.03$ (sig. = 0.87), $^{\dagger\dagger}F(1,627) = 12.22$ (sig. = 0.0005), and $^{\ddagger\ddagger}F(1,623) = 8.77$ (sig. = 0.003).

Table 8. Regression Results: *Attendance, for the Period after June 16, 1896*

Constant	4711.43**
	(1239.50)
Predict Fair Weather	341.68
	(1068.31)
Realized Fair Weather	-463.36
	(1026.86)
Year when the Rank was 3	2289.51*
	(1003.01)
Year when the Rank was 10	-1015.84
	(1238.75)
n	102
R^2	0.102

Note: Authors' calculations, data collected from Baseball Reference, *New York Times*, and the *National Climate Data Center of the National Center for Atmospheric Research* for the baseball season from June 16, 1896 through 1899. **significance = 0.01. *significance = 0.05.

Appendix B: Proofs

Lemma 1 *The set* $\{(p,q)|(p,q) \in [0,1] \times [0,1]$ *and* $p+q \geq 1\}$ *satisfies three properties:*

(1) *Observing a fair signal increases the posterior on state* A *relative to the prior, and observing a rainy signal decreases the posterior on state* A *relative to the prior.*
(2) *For any signal structure* $(p',q') \in [0,1] \times [0,1]$*, there exists a* $(p,q) \in \{(p,q)|(p,q) \in [0,1] \times [0,1]$ *and* $p+q \geq 1\}$ *that generates the same posteriors with the same probabilities as* (p',q')*.*
(3) *For any strict subset* S *of* $\{(p,q)|(p,q) \in [0,1] \times [0,1]$ *and* $p+q \geq 1\}$*, there exists a point* $(p',q') \in [0,1] \times [0,1]$ *such that there is no element of* S *that generates the same posteriors as* (p',q')*.*

Proof We will prove each part of the Lemma in turn. First we prove the first part. Recall that for a given prior $0 < \rho < 1$ on state A (fair weather) and information structure (p,q), the posterior for A given the fair signal is

$$\psi_F = \frac{\rho p}{\rho p + (1-\rho)(1-q)}.$$

Now $\psi_F > \rho$ if and only if

$$\psi_F = \frac{\rho p}{\rho p + (1-\rho)(1-q)} > \rho,$$

which holds if and only if

$$(1-\rho)p > (1-\rho) - (1-\rho)q,$$

which is the same as

$$p + q > 1.$$

An analogous series of steps establishes the result for the posterior after observing a rainy signal,

$$\psi_R = \frac{\rho(1-p)}{\rho(1-p) + (1-\rho)q}.$$

To prove the second part, assume that $p+q \leq 1$. In this case, denote $p' = 1-p$ and $q' = 1-q$. We will simply work with likelihood ratios. Under (p,q), likelihood ratio $\frac{p}{1-q}$ occurs with probability $\rho p + (1-\rho)(1-q)$ and likelihood ratio $\frac{1-p}{q}$ occurs with probability $\rho(1-p) + (1-\rho)q$.

Under (p',q') likelihood ratio $\frac{1-p'}{q'} = \frac{p}{1-q}$ occurs with probability $\rho(1-p') + (1-\rho)q' = \rho p + (1-\rho)(1-q)$. Likelihood ratio $\frac{p'}{1-q'} = \frac{1-p}{q}$ occurs with probability $\rho p' + (1-\rho)(1-q') = \rho(1-p) + (1-\rho)q$. Therefore (p',q') generates the same posterior distribution as (p,q). Moreover, $p'+q' = (1-p)+(1-q) = 2-p-q \geq 1$ since $p+q \leq 1$. So therefore, instead of considering some (p,q) we can always instead consider the corresponding $p' = 1-p, q' = 1-q$. This proves the second part.

To prove the third part, observe that in order for two signal structures (p,q) and (p',q') to generate the same posteriors (so that for both signal structures, a fair weather prediction increases the posterior relative to the prior *and* a rainy weather prediction decreases it) it must be the case that $\frac{p'}{1-q'} = \frac{p}{1-q}$ and $\frac{1-p'}{q'} = \frac{1-p}{q}$.

Therefore $p'-p'q = p-pq'$ and $q-p'q = q'-pq'$, which is equivalent to $q = \frac{-p+pq'+p'}{p'}$ and $q = \frac{q'-pq'}{1-p'}$. Simplifying, we have $\frac{-p+pq'+p'}{p'} = \frac{q'-pq'}{1-p'}$, or $p'q'-pq'p' = -p+pq'+p'+pp'-pp'q'-p'^2$. This implies that $p'q' = -p+pq'+p'+pp'-p'^2$, and so $p(1-q'-p') = -p'q'+p'-p'^2 = p'(1-q'-p')$, which holds only if and only if $p = p'$. This proves the third part.

□

Lemma 2 *Assume p and q jointly satisfy $q \geq 1-p$. Then (p', q') is Blackwell sufficient/more informative for (p, q) if and only if $p' \geq \frac{p}{1-q} - \frac{p}{1-q}q'$ and $p' \geq 1 - q'\frac{1-p}{q}$.*

Proof Recall that one signal structure (p', q') is Blackwell more informative than another (p, q) if and only if the distribution of posteriors induced by (p', q') is a mean preserving spread of the distribution induced by (p, q). By the law of iterated expectations, the expected posterior under (p', q') and (p, q) must be the same — the prior. Because there are only 2 signals, there will be only 2 posteriors. So we just have to show that the posteriors under (p', q') are more extreme (in the sense that they are farther from the prior) than the posteriors under (p, q). In order to simplify the proofs, we will show an equivalent result — that the likelihood ratios under (p', q') are more extreme (farther from 1) than the likelihood ratios under (p, q).

The likelihood ratios after observing a fair signal under (p', q') and (p, q) are (respectively) $\frac{p'}{1-q'}$ and $\frac{p}{1-q}$ while the likelihood ratios after observing a rainy signal are $\frac{1-p'}{q'}$ and $\frac{1-p}{q}$.

In order for the ratios under (p', q') to be farther from 1 than (p, q), then $\frac{p'}{1-q'} \geq \frac{p}{1-q}$ and $\frac{1-p'}{q'} \leq \frac{1-p}{q}$. This is equivalent to $p' \geq \frac{p}{1-q} - \frac{p}{1-q}q'$ and $p' \geq 1 - q'\frac{1-p}{q}$.

□

Proposition 1 *If the decision-maker strictly prefers at least one signal structure to another, then p^* is increasing in H, L, and ρ and q^* is decreasing in H, L, and ρ.*

Proof Because any feasible signal structure is Blackwell dominated by a signal structure on the technology constraint, the agent will always want to choose an information structure along the technology constraint, if any signal feasible signal structure allows her to condition her action. If the agent finds it optimal to condition, then she must be getting a higher payoff than by not conditioning.

First we deal with cases in which the consumer conditions her actions. This occurs when neither conditioning constraint, (4) or (5), is binding. Recall that indifference curves have slope $\frac{\rho H}{(1-\rho)L}$. This slope is negative (because $L < 0$), and is decreasing in H, L, and ρ. If the agent conditions her actions, and chooses an interior solution (p^*, q^*) at the point of tangency to the boundary of the technologically-feasible set (as in Figure 4), then standard comparative statics lead to the result that p^* is increasing in H, L, and ρ, while q^* is decreasing in H, L and ρ.

Next, we consider cases that lead the consumer to *not* condition her actions. As discussed in the text, this happens when one of the two conditioning constraints, (4) or (5), holds with equality (i.e., is binding). Notice that the slope of the conditioning constraints is the same as the slope of the indifference curves for a consumer who conditions her actions, and recall that the binding conditioning constraint passes either through $(p, q) = (1, 0)$ or

$(p, q) = (0, 1)$. So if the consumer is not conditioning her actions, she must be at one of these two corners; either the consumer's indifference curve is so steep that her optimal signal structure, among feasible signals, is $(1, 0)$, or so flat that her optimal signal structure is $(0, 1)$. Either way, the conditioning constraint holds with equality, so the consumer places zero value on any information structure available in the feasible set. Such a consumer is thus trivially indifferent over signal structures.

□

Proposition 2 *On non-game days the information structure produced will exhibit no bias. On game days, if any information is purchased, the information structure will generically be biased, and the bias will be increasing in H, L, and ρ.*

Proof By construction, on non-home game days, the consumer has a preference for unbiased signals, and so if information is provided it will be unbiased.

Recall that our measure of bias is $\mathbb{B} = p - q$. On game days, if the consumer is willing to purchase a signal for positive amounts, then the firm will provide her the optimal feasible signal structure (p^*, q^*) because this maximizes the price at which the information can be sold, and will charge willingness to pay, $r = p^*\rho H + (1 - q^*)(1 - \rho)L - \bar{u}$. In general the optimal signal structure differs on game days because payoffs to actions (e.g., H and L) differ. So we expect game-day bias, and comparative statics from Proposition 1 pertain. Thus, an increase in H, in L, or in ρ, leads to an increase in p^* and a decrease in q^*, and thus to an increase in $\mathbb{B} = p^* - q^*$.

□

Proposition 3 *When bias is supply-driven bias, if information is purchased on non-game days, the information structure produced will exhibit no bias. On game days, if any information is purchased, the information structure will always be positively biased and the bias will be decreasing in H, L, and ρ.*

Proof As in Proposition 2, by construction, on non-home game days the consumer has a preference for unbiased signals, and so if information is provided it will be unbiased.

As for game days, the firm's problem is to maximize it's payoff. We can write this maximization problem as

$$\max_{p,q}\{b(p, 1 - q, H, L, \rho) - c(p - q)\}, \tag{11}$$

subject to the technology constraint,

$$q \leq \phi(p). \tag{12}$$

If the technology constraint is binding, the following first order conditions characterize the resulting optimal signal structure, (p^*, q^*) :

$$b_1(p^*, 1 - q^*, H, L, \rho) - c'(p^* - q^*) + \lambda\phi'(p^*) = 0$$

and

$$-b_2(p^*, 1 - q^*, H, L, \rho) + c'(p^* - q^*) - \lambda = 0.$$

In turn, these first order conditions imply that

$$b_1(p^*, 1 - q^*, H, L, \rho) - c'(p^* - q^*) = \phi'(p^*)[b_2(p^*, 1 - q^*, H, L, \rho) - c'(p^* - q^*)],$$

which in turn can be rewritten,

$$b_1(p^*, 1-q^*, H, L, \rho) - \phi'(p^*) b_2(p^*, 1-q^*, H, L, \rho) = [1-\phi'(p^*)]c'(p^*-q^*). \tag{13}$$

Note that in this last equation each of these terms is positive: $b_1(p^*, 1-q^*, H, L, \rho)$, $-\phi'(p^*)$, $b_2(p^*, 1-q^*, H, L, \rho)$, $[1-\phi'(p^*)]$, and $c'(p^*-q^*)$.

Now we turn to properties of the maximum.

The first property concerns *positive bias.* From inspection of the firm's payoff (7), it is clear that the firm will always choose a positively-biased signal structure, i.e., set $p \geq q$. To see this, consider some signal structure that is not positively biased. If $p < q$ then the firm could instead choose an alternative structure, $p' = q$ and $q' = p$, and $c(p-q)$ would be the same, but $b(p, 1-q, H, L, \rho)$ would be strictly larger.

Next, we need to consider two possibilities for the maximum — the optimum is on the the interior of the feasible set, $\Phi(p, q)$, or on boundary. We start with the latter case.

Suppose the maximum is on a boundary. It cannot be on the lower boundary ($p + q = 1$) because if it is, it has the same properties, from the consumer's perspective, as $p = q = 0.5$ (as is clear from posteriors, (1) and (2) in the text). But as we have seen, it is always optimal for the firm to introduce a small amount of bias. Thus, an optimum on the boundary of $\Phi(p, q)$ must lie on the technology constraint, so (8) is binding. We have already seen that solution in this case is characterized by (9).

The comparative statics exercises are standard. Consider, for example, an increase in H. Since we are on a boundary, we substitute $q^* = \phi(p^*)$ into (9), and then evaluate that expression using the implicit function theorem. This gives

$$\frac{\partial p^*}{\partial H} = -\frac{b_{11} - \phi'(p^*)b_{12} - \phi''(p^*)b_2 - \phi'(p^*)b_{21} + \phi'(p^*)^2 b_{22} + \phi''(p^*)c'(\cdot) - [1-\phi'(p^*)]^2 c''(\cdot)}{b_{13} - \phi'(p^*)b_{23}},$$

which is negative (given assumptions about the signs of cross partials and the convexity of $c(\cdot)$. Similar exercises lead to the conclusions,

$$\frac{\partial p^*}{\partial H} < 0, \quad \frac{\partial p^*}{\partial L} < 0, \quad \text{and} \quad \frac{\partial p^*}{\partial \rho} < 0, \tag{14}$$

and q^* must in each case move the opposite direction. This gives us the desired results concerning bias, $\mathbb{B} = p^* - q^*$.

Given our assumptions, second order conditions pertain, and the optimum on the technology constraint must be unique. We note, for completeness sake, that there cannot be a corner solution on the technology constraint given our assumptions that both $\phi'(p)$ and $\phi''(p)$ are strictly negative.

Suppose that instead the the optimum is in the interior of the feasible set. In this instance first order conditions for (7) simplify to

$$b_1(p^*, 1-q^*, H, L, \rho) - c'(p^*-q^*) = 0,$$

and

$$-b_2(p^*, 1-q^*, H, L, \rho) + c'(p^*-q^*) = 0.$$

Here, comparative static exercises analogous to those just presented lead us to the same conclusion just presented.

□

Proposition 4 *If an individuals has wishful thinking preferences then the optimal signal always has positive bias. Moreover the bias in their optimal signal is decreasing in ρ and increasing in H.*

Proof The Lagrangian is $H(p\rho + (1-q)(1-\rho))^{\frac{1}{2}} - (p-q)^2 + \lambda(\phi(p) - q)$. The optimum is always on the technology frontier. Observe that if it is not then they individual could increase both p and q by ϵ, not change the cost, but increase the benefit.

The first order conditions are then:

$$H.5\rho(p\rho + (1-q)(1-\rho))^{-.5} - 2(p-q) + \lambda\phi'(p) = 0$$

and

$$-H.5(1-\rho)(pr + (1-q)(1-\rho))^{-.5} + 2(p-q) - \lambda = 0$$

Combining them gives:

$$\frac{H}{2}(p\rho + (1-k+p)(1-\rho))^{-.5} - 4(2p-k) = 0$$

A necessary condition for the solution is that $4(2p-k) > 0$, or $2p > k > 1$, so $p > q$.

To derive the comparative statics results we use the implicit function theorem. Denoting $f(p, \rho, H) = \frac{H}{2}(p\rho + (1-k+p)(1-\rho))^{-.5} - 4(2p-k)$ we obtain:

$$\begin{aligned}\frac{\partial f}{\partial p} &= -.25H(\rho + (1-\rho))(p\rho + (1-k+p)(1-\rho))^{-1.5} - 8 \\ &= -.25H(p\rho + (1-k+p)(1-\rho))^{-1.5} - 8\end{aligned}$$

$$\begin{aligned}\frac{\partial f}{\partial \rho} &= -.25(p-1+k-p)H(p\rho + (1-k+p)(1-\rho))^{-1.5} \\ &= -.25(k-1)H(p\rho + (1-k+p)(1-\rho))^{-1.5}\end{aligned}$$

$$\frac{\partial f}{\partial H} = (p\rho + (1-k+p)(1-\rho))^{-.5}$$

Therefore

$$\frac{\partial p}{\partial r} = \frac{.25(k-1)H(p\rho + (1-k+p)(1-\rho))^{-1.5}}{-.25H(p\rho + (1-k+p)(1-\rho))^{-1.5} - 8}$$

The numerator and denominator have difference signs, so this is negative.

Moreover,

$$\frac{\partial p}{\partial H} = \frac{-(p\rho + (1-k+p)(1-\rho))^{-.5}}{-.25H(p\rho + (1-k+p)(1-\rho))^{-1.5} - 8}$$

which, because both the numerator and denominator have the same sign, is positive. □

References

Artstein-Avidan, S. and Dillenberger, D., 2010. "Dynamic Disappointment Aversion: Don't Tell Me Anything Until You Know For Sure," PIER Working Paper Archive, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.

Baron, D.P., 2006. "Persistent Media Bias," *Journal of Public Economics*, 90(1), 1–36.

Brunnermeier, M.K. and Parker, J.A., 2005. "Optimal Expectations," *American Economic Review*, 95(4).

DellaVigna, S. and Gentzkow, M., 2010. "Persuasion: Empirical Evidence," *Annual Review of Economics*, 2(1), 643–669.

DellaVigna, S. and Kaplan, E., 2007. "The Fox News Effect: Media Bias and Voting," *Quarterly Journal of Economics,* 122(3), 1187–1234.

Durante, R. and Knight, B. 2012. "Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy," *Journal of the European Economic Association*, 10(3), 451–481.

Gentzkow, M. and Kamenica, E., 2011. "Bayesian Persuasion," *American Economic Review*, 101(6).

Gentzkow, M., Petek, M., Shapiro, J., and Sinkinson M. 2012. "Do Newspapers Serve the State? Incumbent Party Influence on the US Press, 1869-1928."

Gentzkow, M.A. and Shapiro, J.M., 2004. "Media, Education and Anti-Americanism in the Muslim World," *Journal of Economic Perspectives*, 18, 117–134.

Gentzkow, M.A. and Shapiro, J.M., 2006. "Media Bias and Reputation," *Journal of Political Economy*, 116, 280–316.

Gentzkow, M.A. and Shapiro, J.M., 2008. "Competition and Truth in the Market for News," *Journal of Economic Perspectives*, 22(2), 133–154.

Gentzkow, M.A. and Shapiro, J.M., 2010. "What Drives Media Slant? Evidence From US Daily Newspapers," *Econometrica,* 78(1), 35–71.

Gentzkow, M. and Shapiro, J., 2011. "Ideological Segregation Online and Offline," *Quarterly Journal of Economics*, 126(4).

Gentzkow, M., Shapiro, J., and Sinkinson M., 2010. "Competition and Ideological Diversity: Historical Evidence from US Newspapers."

Gentzkow, M., Shapiro, J., and Sinkinson, M., 2011. "The Effect of Newspaper Entry and Exit on Electoral Politics," *American Economic Review*, 101(7).

Groseclose, T. and Milyo, J., 2005. "A Measure of Media Bias," *Quarterly Journal of Economics*, 120(4), 1191–1237.

Koszegi, B. and Rabin, M., 2009. "Reference-Dependent Consumption Plans," *American Economic Review*, 99(3), 909–936.

Kreps, D.M. and Porteus, E.L., 1978. "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica,* 185–200.

Mullainathan, S. and Shleifer, A., 2005. "The Market for News," *American Economic Review*, 1031–1053.

Nebeker, F., 1995. *Calculating the Weather: Meteorology in the 20th Century*, Academic Press.

Oster, E., Shoulson, I., and Dorsey, E., 2011. "Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease," National Bureau of Economic Research.

Proquest. 1851–2007. *New York Times.*

Proquest. 1889–1993. *Wall Street Journal.*

Sports Reference LLC, 2013. *Baseball Reference.*

Suen, W., 2004. "The Self-Perpetuation of Biased Beliefs," *Economic Journal*, 114(495), 377–396.