Media Bias and Reputation

Matthew Gentzkow
Graduate School of Business
University of Chicago

Jesse M. Shapiro*
Department of Economics
Harvard University

January 2, 2005

Abstract

A Bayesian consumer who is uncertain about the quality of an information source will infer that the source is of higher quality when its reports conform to the consumer's prior expectations. We use this fact to build a model of media bias in which firms slant their reports toward the prior beliefs of their customers in order to build a reputation for quality. Bias emerges in our model even though it can make all market participants worse off. The model predicts that bias will be less severe when consumers receive independent evidence on the true state, and that competition between independently owned news outlets will reduce bias. We present a variety of empirical evidence consistent with these predictions.

JEL classification: L82, L10, D83

Keywords: bias, news, media, reputation

*We are grateful to Alberto Alesina, Attila Ambrus, Nigel Ashford, Chris Avery, Heski Bar-Isaac, Gary Becker, Tyler Cowen, Jeremy Fox, Drew Fudenberg, Ed Glaeser, Jerry Green, James Heckman, Tom Hubbard, Steve Levitt, Larry Katz, Kevin M. Murphy, Roger Myerson, Canice Prendergast, Matthew Rabin, Andrei Shleifer, Lars Stole, Richard Thaler, Richard Zeckhauser and seminar participants at Harvard University, the University of Chicago, and the Institute for Humane Studies for helpful comments. We thank Christopher Avery, Judith Chevalier, Matthew Hale, Martin Kaplan, Bryan Boulier and H. O. Stekler for generously providing access to their data. Fuhito Kojima, Jennifer Paniza and Tina Yang provided excellent research assistance. Shapiro acknowledges financial assistance from the Institute for Humane Studies, the Center for Basic Research in the Social Sciences, the Chiles Foundation and the National Science Foundation. E-mail: gentzkow@ChicagoGSB.edu, jmshapir@fas.harvard.edu.

1 Introduction

On December 2, 2003, American troops fought a battle in the Iraqi city of Samarra. Fox News began its story on the event with the following paragraph:

> In one of the deadliest reported firefights in Iraq since the fall of Saddam Hussein's regime, US forces killed at least 54 Iraqis and captured eight others while fending off simultaneous convoy ambushes Sunday in the northern city of Samarra. (Fox News 2003).

The *New York Times* article on the same event began:

> American commanders vowed Monday that the killing of as many as 54 insurgents in this central Iraqi town would serve as a lesson to those fighting the United States, but Iraqis disputed the death toll and said anger against America would only rise. (*New York Times* 2003).

And the English-language website of the satellite network Al Jazeera began:

> The US military has vowed to continue aggressive tactics after saying it killed 54 Iraqis following an ambush, but commanders admitted they had no proof to back up their claims. The only corpses at Samarra's hospital were those of civilians, including two elderly Iranian visitors and a child (AlJazeera.net 2003).

All of the accounts are based on the same set of underlying facts. Yet by selective omission, choice of words, and varying credibility ascribed to the primary source, each conveys a radically different impression of what actually happened. The choice to slant information in this way is what we will mean in this paper by media bias.

Such bias has been widely documented, both internationally and within the United States.[1] Concern about bias has played a prominent role in many policy debates, ranging from public diplomacy in the Middle East (Satloff, 2003; Peterson et al, 2003) to ownership regulation by the FCC (Cooper, Kimmelman and Leanza, 2001). Moreover, survey evidence revealing rising

[1] The differences between the slant of Arab and American news sources in covering the Middle East are documented at length by Ajami (2001). A sampling of recent works documenting bias in US national media includes books by Franken (2003), Coulter (2003), Goldberg (2003), and Alterman (2003). Underhill and Pepper (2003) discuss accusations of prejudicial reporting at the BBC.

polarization and falling trust in the news media has prompted concerns about the market's ability to deliver credible information to the public (Kohut, 2004).

In this paper, we develop a new model of media bias. Existing models of bias all take as given that some agents in the economy—consumers (Mullainathan and Shleifer, 2003), reporters (Baron, 2003), or governments (Besley and Prat, 2004)—*prefer* for news suppliers to distort the information they provide.[2] In contrast, our model shows that bias can arise even when news consumers care only about learning the truth, news sellers care only about maximizing profits, and eliminating bias could make all agents in the economy better off. The framework generates testable predictions about the way consumer beliefs, competition, ownership structure, and the characteristics of specific issues influence the way news is slanted in equilibrium. We present a range of empirical evidence consistent with these predictions.

We start from a simple assumption: A media firm wants to build a reputation as a provider of accurate information. If the quality of the information a given firm provides is difficult to observe directly, consumer beliefs about quality will be based largely on observations of past reports. Firms will then have an incentive to shape these reports in whatever way will be most likely to improve their reputations and thus increase their future profits by expanding the demand for their product.

We explore this intuition in a reputation model where long-lived media firms sell information to overlapping generations of short-lived consumers. We show that the reports of these firms are biased in equilibrium, and that this may strictly reduce the welfare of all agents in the economy. While distinguishing our model from existing theories of media slant, this feature relates to several recent papers that show how reputational incentives can lead to inefficient outcomes in information-exchange settings (see, for example, Ely and Välimäki, 2003; Ely, Fudenberg and Levine, 2002; Prendergast, 1993; and Morris, 2001).[3]

[2] An earlier version of Mullainathan and Shleifer's (2002) paper does not assume that consumers have a taste for confirmatory information but generates similar behavior through a mechanism in which consumers think categorically.

[3] More generally, this paper is related to the study of sender-receiver games. See, for example, Glazer and Rubinstein (1998, 2001, and 2004), Banerjee and Somanathan (2001) , Battaglini (2002), Gilligan and Krehbiel (1989), and Krishna and Morgan (2001).

Our first set of results shows that firms will tend to distort information to make it conform with consumers' prior beliefs. To see why, consider that a noisy or inaccurate signal is more likely to produce reports that contradict the truth. An agent who has a strong prior belief about the true state of the world will therefore expect inaccurate information sources to contradict that belief more often than accurate ones. Suppose, for example, that a newspaper reports that scientists have successfully produced cold fusion. If a consumer believes this to be highly unlikely *a priori*, she will rationally infer that the paper probably has poor information or exercised poor judgment in interpreting the available evidence. A media firm concerned about its reputation for accuracy will therefore be reluctant to report evidence at odds with consumers' priors, even if they believe the evidence to be true. The more priors favor a given position, the less likely the firm becomes to print a story contradicting that position.

Our second main result is that when consumers have access to a source that can provide *ex-post* verification of the true state of the world, firms' incentives to distort information are weakened. If a firm misreports its signal so as to move closer to consumers' priors, it runs the risk that the truth will come out and its report will be falsified, damaging its reputation. As the likelihood of *ex-post* feedback about the state of the world improves, the amount of bias occurring in equilibrium decreases. Our model therefore predicts relatively less bias in contexts where predictions are concrete and outcomes are immediately observable—forecasting weather, sports outcomes, or stock returns, for example. It predicts more bias in coverage of a foreign war, discussion of the impact of alternative tax policies, or summary of scientific evidence about global warming, contexts where outcomes are difficult to observe and are often not realized until long after the report is made.

The analysis of feedback yields as a corollary our third result: Competition in the news market leads to lower bias. A firm competing with another news outlet runs the risk that, if it distorts its signal, the competitor's report will expose the inaccuracy and thus reduce consumers' assessments of the distorting firm's quality. Our prediction that increased competition lowers the incentive to bias reports toward consumer priors contrasts sharply with that of Mullainathan and Shleifer (2003), who argue that increased competition will *tighten* the connection between priors and reports.[4] We also show that it is competition per se, and not simply increasing the number of firms, that drives this result: if all firms in a market are jointly owned, bias may remain unchanged even as the number of firms goes to infinity.

At the end of the paper, we present empirical evidence on each of these three key predictions. First, we confirm the intuitive prediction that prior beliefs affect media slant. We show, for example, that newspapers in poorer states were more likely to endorse Bush in 2000 than those in rich states, consistent with the fact that voters in poorer states were more pro-Bush. Next, we show that a wide range of existing evidence on financial reporting suggests that feedback can limit bias, and that in high-feedback settings, such as weather reporting, bias tends to be relatively minor. We also highlight the fact that local sports columnists do not excessively favor their local teams in forecasting game outcomes, which is consistent with an important role for rapid feedback in limiting the incentive to slant. Finally, we argue that in many cases markets with greater media competition are less likely to suppress facts that challenge pre-existing beliefs, and show quantitatively that television news reports leading up to the 2000 election were more equitable in their treatment of Bush and Gore in more competitive markets.

As a final note, we are not arguing that the mechanisms identified in previous models of media bias are unimportant. It is indisputable, for example, that pressure from governments sometimes shapes media coverage. There is also psychological evidence suggesting that consumers seek out evidence that conforms to their prior beliefs. The observation that bias may exist even in the absence of such preferences is important because it implies different predictions and has different implications for policy than models that assume a taste for bias. We present empirical evidence

[4] The intuition for their result is that if consumers have heterogeneous priors about the true state and have a psychological taste for bias in the same direction as their priors, increasing the number of products may cause the market to become segmented with products catering directly to each consumer type. This segmentation may increase bias. For example, if consumers are evenly spread from extreme left to extreme right, a monopoly firm might prefer to locate in the middle of the distribution and report with no bias; duopolists, on the other hand, might split the market with one biasing far right and the other biasing far left. We rule out this kind of segmentation effect by assumption in order to highlight a new effect. In a richer model in which both effects were allowed to operate, the net effect of competition on bias would be ambiguous.

below which is consistent with these predictions, suggesting that the mechanisms we identify play a significant role. However, the relative importance of the reputational effects we identify and the preferences discussed in past work remains an open empirical question.

In the next section we discuss the role of reputational incentives in media markets. In section 3 we introduce the building blocks of the model, and in section 4 we characterize its unique equilibrium and the welfare effect of bias. The key comparative statics are then proved in section 5: that equilibrium bias is correlated with consumer priors, decreasing in the amount of ex post feedback, and decreasing in the extent of market competition. In section 6, we extend the model to allow consumers with heterogeneous prior beliefs to coexist in the same market. We show that it is possible to have segmented equilibria where each firm provides information to only one type of consumer and slants its reports accordingly, and that the key comparative statics remain valid in this setting. Section 7 presents empirical evidence, and section 8 concludes. Appendix A contains extensions of the basic model.

2 Credibility and Quality in the Media

In this section, we present evidence supporting two key building blocks of our model: Media firms try to build a reputation for truthful reporting, and consumers' assessments of the quality of news sources depend on prior beliefs.

2.1 The Importance of Reputation in Media Markets

At the heart of our model will be media firms' desire to maintain a reputation for accuracy in reporting. The high costs firms are willing to incur to gather information provide strong evidence of such an incentive,[5] as does the response of media firms whose reports are revealed to have

[5] To take one example, Andrew Lack, President of NBC News, estimated at the beginning of the war in Afghanistan that covering it would cost each network approximately 1 million dollars per week—10 percent of their total weekly expenditures (Auletta 2001). One would not expect to see this level of expense if consumers were not significantly concerned with the factual content of news.

been inaccurate. For example, on September 8, 2004, CBS News anchor Dan Rather reported the emergence of new evidence indicating that President Bush's family had pulled strings in order to get him into the Texas Air National Guard and avoid his having to serve in Vietnam. When later information indicated that the documents on which the report was based may have been fabricated, both Rather and CBS President Andrew Heyward issued apologies emphasizing the importance of a reputation for truth-telling in journalism.

Heyward wrote that "nothing is more important to [CBS] than our credibility and keeping faith with the millions of people who count on us for fair, accurate, reliable, and independent reporting. We will continue to work tirelessly to be worthy of that trust" (Heyward, 2004). Rather's statement echoed Heyward's, explaining that "nothing is more important to [CBS] than people's trust in our ability and our commitment to report fairly and truthfully" (Rather, 2004).

Similarly, the exposure of Jayson Blair's fraudulent reporting at the *New York Times* prompted the resignation of top-ranking editors Howell Raines and Gerald M. Boyd. Former Tupperware chief executive Warren L. Batts remarked, "They, of course, had to resign... Any company has to sell the credibility of its product, but a media company has nothing else to sell" (Kirkpatrick and Fabrikant, 2003).[6]

2.2 The Influence of Priors on Quality Assessments

How do consumers determine whether a news source is trustworthy? Consider a weather forecaster who can either predict *sun* or *rain*. Imagine that forecasters come in two types: one who simply flips a fair coin to determine her report, and another who always perfectly predicts tomorrow's weather. If a consumer is living in Los Angeles, where sun is by far the more common state, she will believe that the coin-flipper is much more likely to report rain than the accurate forecaster. Upon seeing a forecast of rain, she will therefore update her beliefs about the type of the forecaster

[6] An investigation by the *Times* discovered that Blair had "fabricated comments," "concocted scenes," and "selected details from photographs to create the impression he had been somewhere or seen someone, when he had not" (Barry et al, 2003). *New York Times* publisher Arthur Sulzberger Jr. called Blair's deceptions "an abrogation of the trust between the newspaper and its readers" (Barry et al, 2003).

toward the view that the forecaster is a coin-flipper.

This link from priors to inferences about quality reflects a general property of Bayesian updating. It is closely related to the dynamic modeled in Prendergast (1993), which leads employees to reinforce the beliefs of their superiors even when they have evidence that those beliefs are wrong.[7] We show in Appendix A that the intuition applies much more generally than the simple model considered below.

Updating beliefs in this way may be normatively correct, but does it reflect actual consumer behavior? A large body of psychological research documents a strong connection between subjects' prior views and their assessments of information sources. In perhaps the best known paper on this subject, Lord, Ross, and Lepper (1979) show that experimental subjects evaluating studies of the deterrence effect of the death penalty rate studies supporting their prior beliefs as both more "convincing" and "better done."[8] This basic finding is replicated and expanded by Lord, Lepper, and Preston (1984), Miller et al. (1993), and Munro and Ditto (1997). Along the same lines, Koehler (1993) shows that scientists rate experiments as higher quality when the experimental results conform to the scientists' belief about a controversial issue.

Evidence on consumer assessments of media quality in the real world show a similar pattern. To take one example, Figure 1 shows that in a recent survey nearly 30 percent of respondents who described themselves as "conservative" indicated that they thought they could believe all or most of what the Fox Cable News Network says. In contrast, less than 15 percent of self-described liberals said that they could believe all or most of what the network reports. Ratings of National Public Radio, show the opposite pattern: more than 35 percent of liberals believe all or most of what NPR says, as opposed to less than 20 percent of conservatives.

[7] See also Brandenburger and Polak (1996).

[8] This paper is often cited as evidence that consumers have confirmatory bias—i.e. a taste for information that confirms their prior beliefs. We simply note that the evidence on evaluating the quality of information sources is equally consistent with a Bayesian model. A second finding in this paper that lends support to the confirmatory bias hypothesis is that seeing the same information led subjects with different prior beliefs to *diverge* in their opinions rather than converge toward the truth as Bayesian updating would suggest. Subsequent work (Miller et al. 1993; Munro and Ditto 1997), however, argues that this "attitude polarization" phenomenon is less robust than the paper's other results.

Other evidence comes from the 2002 Gallup Poll of the Islamic World (The Gallup Organization, 2002). Respondents in nine Islamic countries were asked to report whether each of the following five descriptions applies to CNN: has comprehensive news coverage; has good analyses; is always on the site of events; has daring, unedited news; has unique access to information. In Appendix C we show that an index of these quality assessments is strongly correlated with respondents' reported favorability toward the US. To deal with the possibility of reverse-causality, we also construct a proxy for favorability based on respondents' reported religiosity and show that this also has a strong correlation with quality assessments. Taken together these pieces of evidence strongly suggest that prior beliefs influence consumers' judgements of quality in the way a Bayesian model would predict.

3 A Model with Reputation

3.1 Overview

In this section, we introduce the building blocks of a model in which media firms endogenously choose to distort the information they report. We follow much of the previous literature in modeling the information provided by media firms as an informative signal about some unknown state of the world, and assuming that consumers value this information because they face some decision whose payoffs are connected with the true state. This could represent actual decisions that depend on the news, either with large instrumental consequences (whether to join a terrorist group opposing the United States) or with minor consequences (what position to support in an argument with friends). It could also represent consumers who value information intrinsically as in Grant, Kajii and Polak (1998).[9]

In the analysis that follows, we assume all consumers in the market share the same prior about

[9] Of course some information provided by the media would fit none of these descriptions. Much of the content of news programs or newspapers does not affect any strongly held beliefs and could properly be classified as entertainment—coverage of crime, auto accidents, or fires would be obvious examples. We will assume, however, that decisions about how to slant coverage of those issues that do impact beliefs are independent of choices about entertainment content.

the true state. Comparative statics with respect to this prior can be thought of as predictions about how bias would vary across markets with different consumers—why Republican and Democratic towns have different bias in their newspapers, for example, or why the slant of Al Jazeera differs from that of CNN. In section 6 below we consider a richer model with heterogeneous consumer priors in a single market, and show that the key comparative statics results continue to hold.

In the model several firms each print a newspaper. One firm in each period gets a "scoop" on the true state of the world. That is, it receives a noisy signal about the truth and then prints a report which is seen by readers of the paper. The newspaper can freely choose whether to report truthfully or with bias, and firms that do not get a scoop do not make a report.[10] We introduce reputational incentives by assuming that with a small probability each firm is "high quality." When a high-quality firm has a scoop it perfectly observes the true state of the world and always reports truthfully.[11] Firms will therefore have an incentive to build a reputation for quality.

The intuition we wish to capture about the effect of feedback and competition depends on the notion that information about the true state is gradually revealed over time. Firms that distort their reports when they have a scoop face the possibility that consumers will later learn the truth, either independently or from a competing firm. To capture this in the simplest possible way, we suppose that after consumers have chosen which newspaper to read and taken their actions whose payoff depends on the true state, there is a "feedback stage." Each consumer may learn the true state at this point with some probability.[12] All firms also learn the true state and print a report about it. Again, we assume that high types report honestly but that normal firms have discretion

[10] Allowing firms to freely distort information seems the simplest way to model the much richer ways in which slant is introduced in the real world. In reality, news sources may slant information by reporting some facts and omitting others, by changing the order in which facts are presented, by presenting sources as more or less credible, or by using language with positive or negative connotations. In Appendix A.2 we extend the model to allow for richer bias technologies, including the case where firms receive several signals each period and choose which to print and which to omit.

[11] The model could easily be extended to allow high-quality firms to have a noisy signal, provided it is more accurate than the signal of a low-quality firm. We show in Appendix A.3 that the model can also be extended to allow the high type to have discretion in reporting. This complicates the analysis because it potentially allows multiple equilibria, but we show that the equilibrium we focus on (where the high type always reports honestly) is the unique *stable* equilibrium of the richer game.

[12] This could either be independent across consumers or correlated.

over what to report. What the consumer learns in the feedback stage is important because if the consumer learns that the truth contradicts a firm's report, their estimate of that firm's quality will suffer.

3.2 Information Structure

We consider an infinite-horizon model where there is an unknown state of the world in each period which we denote $\omega \in \{L, R\}$. The true state is independent across periods, so that learning ω in one period has no effect on future beliefs. In each period, all consumers place prior probability θ on the event $\omega = R$. Nothing in the analysis would change if this prior were allowed to differ across periods. To rule out trivial cases and simplify the exposition, we assume $1 > \theta > .5$. We assume that all firms place prior probability $\frac{1}{2}$ on the event $\omega = R$. Other assumptions, for example that the firms' prior is identical to the consumers', would not change any of the results below.[13]

There are J firms indexed by j. Information revelation in each period happens in two stages. In the first stage, the *reporting stage*, one firm is randomly chosen to receive a signal $d \in \{L, R\}$ which is informative about ω.[14] We will say that this firm has a "scoop" on the story, and will assume that the probability of having a scoop is independent over time. In the second stage, the *feedback stage*, all firms learn the true value of ω and a fraction μ of consumers learn ω independently.[15]

With a small probability λ each firm is "high quality." In this case the firm observes the true state perfectly in the reporting stage ($d = \omega$). Otherwise, the firm is "normal" and d is a noisy signal of the truth, where $d = \omega$ with probability π. To focus attention on the most relevant cases, we will assume that normal newspapers are sufficiently informative that their underlying information in the reporting stage would be valuable to consumers. This will be guaranteed by assuming that

[13] See Morris (1995) for a discussion of the role of heterogeneous priors in economic theory, and Morris (1994) for an application with heterogeneous priors.

[14] In section A.2 we generalize the model to allow for a continuous signal space and show that none of our results are sensitive to the assumption of a binary signal.

[15] This might be because they observe the truth directly, as they would if the issue at stake is whether or not it is going to rain. Or, it could be that they learn the state by talking to friends or reading sources from outside of the market we model. If a local newspaper suppresses information about a national political scandal, for example, consumers might still learn about it by watching a national news report.

$\pi > \theta$.[16]

In the reporting stage of period t, the firm with a scoop (say firm j) prints a report whose value is either L or R. In the feedback stage, all firms print a report. We denote the reporting stage choice by n_j and feedback stage choice by $\tilde{n}_j$. A high-quality firm always reports truthfully—that is, sets $\tilde{n}_j = \omega$ always and $n_j = \omega$ whenever it has a scoop. A normal firm can freely choose to report L or R in both the reporting and feedback stages.

3.3 Consumer Problem

We assume there are overlapping generations of consumers, each of which lives for two periods. There are a large number of consumers in each generation. We index consumers by i and say that a consumer is *young* in the first period she is alive and *old* in the second.

Each consumer chooses an action $a \in \{L, R\}$ at the end of the reporting stage and receives utility 1 if $a = \omega$ and utility 0 otherwise. So as not to render irrelevant the feedback process we have modeled, we suppose that consumers receive utility and thus learn the true states of the world only at the end of their participation in the game.

A consumer's expected reporting-stage utility from an optimal decision if she receives no new information about ω will be θ. Expected utility from following the report of a firm known to be high quality will be 1. And expected utility from following the report of a normal firm will depend on the firm's reporting-stage strategy. We represent such a strategy by s and write the expected utility as $v(s)$.[17] Because $\pi < 1$ it must be the case that $v(s) < 1$.

[16] Relaxing this assumption would not change our results substantially in the current setup where the firm's prior is fixed at $\frac{1}{2}$, except that as π gets very low the model might have no equilibria where consumers value firms' information. In the alternative case where the firm's prior is θ, however, allowing $\pi < \theta$ would change the implications more substantially. The key difference in this case is that the firm's posterior belief about the true state is no longer equal to its signal d—i.e. after seeing $d = L$ the firm would still believe that R is more likely to be the true state. We would want to define "bias" in this case to be the probability that the firm's report deviates from its posterior (the proper definition with respect to welfare). Bias would therefore push the firm toward reporting L relatively more often rather than R relatively more often as in the case we analyze. However, the key comparative statics, that firm reports are correlated with consumer priors and that bias decreases with competition and feedback, would remain unchanged.

[17] Formally, a reporting-period strategy is a mapping from d to a probability distribution over $\{L, R\}$. The value

The added value of reading a paper with a scoop whose quality is uncertain can be written as:

$$V\left(\hat{\lambda}; s\right) = \max\left\{\hat{\lambda} + \left(1 - \hat{\lambda}\right) v\left(s\right) - \theta, 0\right\}. \tag{1}$$

where $\hat{\lambda}$ is the probability the consumer places on the firm being high quality. The first term in the maximum is the added value conditional on following the firm's report. If this term is negative, the consumer would prefer to ignore the firm's report in which case the added value would be 0.

In addition to the information value, we allow heterogeneous consumer-specific tastes for each product in each period. These could represent, for example, consumers' tastes for the non-news components of the product, such as the TV listings and comic strips. We denote the idiosyncratic benefit to consumer i from reading newspaper j in period t by δ_{ij}. We treat δ_{ij} as stochastic from the perspective of the firm, and assume it is uniformly distributed on the interval $[-1, 1]$, with draws independent across i and j.[18]

The gain to consumer i from reading newspaper j as a function of the consumer's belief about the firm's quality and the firm's strategy is:

$$u_{ij}\left(\hat{\lambda}; s\right) = \begin{cases} V\left(\hat{\lambda}; s\right) + \delta_{ij} & \text{if } j \text{ has a scoop} \\ \delta_{ij} & \text{otherwise} \end{cases} \tag{2}$$

Consumer i reads newspaper j if and only if $u_{ij}\left(\hat{\lambda}; s\right) \geq 0$.[19] We assume for simplicity that

of following such a report will be:

$$v\left(s\right) = \max \Pr\left(n = R | \omega = R, s\right)\theta + \Pr\left(n = L | \omega = L, s\right)\left(1 - \theta\right)$$

[18] Assuming a uniform distribution simplifies the analysis because it will make firm profits linear in consumer beliefs about quality, $\hat{\lambda}$. The results require only that profits are *increasing* in $\hat{\lambda}$, however, and so would be robust to assuming a more general distribution.

[19] It is important to emphasize two assumptions built into this specification. First, we assume that consumers are myopic and do not incorporate the gain to learning about quality in their first period decision making. If they were forward-looking, the probability of consuming each newspaper in the first period would increase, but neither the second-period choice probabilities nor, therefore, firms' strategies would change.

Second, we assume that the value of one newspaper does not depend on what other newspapers the consumer reads as well. In terms of the information component $V()$, this is an implication of the model—since only one newspaper consumers make a single consumption decision in each period—a consumer either reads product j in both the reporting and feedback stages of a given period, or does not read it at all.[20]

3.4 Firm Problem

We assume the firm receives some fixed advertising revenue for each consumer that buys its newspaper. To focus on the reporting decision, we abstract away from pricing, fixing all newspapers' prices at zero. Firms choose n_j and $\tilde{n}_j$ to maximize total discounted profits. (Since only normal firms have discretion about what to report, the analysis will focus on the equilibrium strategies of normal firms).

This problem is greatly simplified by the overlapping generations structure assumed above. First, the demand of young consumers is independent of firm decisions, because these consumers make their purchase decisions before seeing any firm reports. Second, the demand of old consumers in period t does not depend on any firm decisions in period t, again because consumers have made their consumption decisions before seeing the firms' period-t reports. The only link between firm decisions and demand (and thus profits) is that old consumers' demand depends on the reports they saw when they were young. Firms thus have an incentive to build a reputation for high quality among young consumers.

In order for this reputational dynamic to have bite, firm profits must in fact depend on consumers' beliefs about quality. $V\left(\hat{\lambda}; s\right)$ in Equation 1 is strictly increasing in $\hat{\lambda}$ if and only if it is strictly positive—i.e. if and only if $\hat{\lambda} + \left(1 - \hat{\lambda}\right) v\left(s\right) > \theta$ (remember that $v\left(s\right) < 1$). This will be guaranteed for any $\hat{\lambda}$ if $v\left(s\right) > \theta$. If a particular reporting-stage strategy s satisfies this condition, we will say strategy s is *informative.*

can have a scoop in a given reporting period, the expected information value is independent of what other papers have been read. The more substantive restriction is that the idiosyncratic value (δ_{ij}) is also independent of the bundle consumed. As will become clear below, this assumption greatly simplifies the analysis because it will mean that a given firm's demand only depends on its own perceived quality and not on the perceived quality of its competitors.

[20] Assuming that the same set of consumers reads in the reporting and feedback stages simplifies the feedback-stage analysis. What we would require in a more general model is a condition that guarantees that readership is not too unstable over time, so that a firm which has a scoop in the reporting stage faces a large number of the same readers in the feedback stage, and would therefore prefer not to contradict its earlier report.

To represent the firm's problem formally, let $\tilde{s}$ be a feedback-stage strategy,[21] and define $\mathbf{s} = \{s, \tilde{s}\}$. Let $\hat{\lambda}(n, \tilde{n}; \mathbf{s}, \theta)$ denote the posterior on quality of a young consumer at the end of the period, given that the consumer's prior on the true state was θ, the firm was playing strategy $\mathbf{s}$, the consumer read a paper that reported n and $\tilde{n}$ and the consumer received no other relevant information. Now suppose that $\hat{\theta}_{i-j}$ is young consumer i's posterior on the true state after seeing all other information received in the period except n and $\tilde{n}$ (i.e. after seeing reports from firms other than j that i read and any exogenous feedback that i received). Then the fact that firms' types are assumed to be independent means we can write the consumer's posterior on j's quality after seeing *all* information revealed in the period as $\hat{\lambda}\left(n, \tilde{n}; \mathbf{s}, \hat{\theta}_{i-j}\right)$.

Observe that the probability that any given consumer reads a product with a scoop in a given period is:

$$\Pr\left(\delta_{ij} \geq -V\left(\hat{\lambda}; s\right)\right) = \frac{1}{2} + \frac{1}{2}V\left(\hat{\lambda}; s\right), \tag{3}$$

where $\hat{\lambda}$ is the consumer's belief about the firm's quality and s is that firm's strategy. The probability of reading a paper without a scoop is just 1/2. Clearly, if s is informative, the probability the consumer reads is strictly increasing in $\hat{\lambda}$.

Suppose consumers expect a firm to play an informative strategy in the next period. Equations 1 and 3 imply that the firm's expected profit from a consumer who reads its paper in the current period is linearly increasing in $\hat{\lambda}\left(n, \tilde{n}; \mathbf{s}, \hat{\theta}_{i-j}\right)$. In both the reporting and the feedback stages, the firm will therefore choose its reports to maximize the expected value of $\hat{\lambda}$.

4 Reporting in Equilibrium

As shown in the previous section, the reputational incentives we are interested in disappear if consumers in period t do not expect firms to play informative strategies in period $t+1$. Since consumer posteriors $\hat{\lambda}()$ do not affect firm profits in this case, firms are indifferent about their

[21] $\tilde{s}$ is a map from the firm's first-stage report and the true state to its feedback stage report. That is, $\tilde{s} : \{L, R, 0\} \times \{L, R\} \to \{L, R\}$ where $n = 0$ represents the case where the firm did not have a scoop in the first period.

reporting decisions. The model thus permits (in some cases) babbling equilibria in which consumers ignore firms' information and reporting strategies are not uniquely determined by the model. To focus attention on the interesting cases, we will define the class of *informative Bayesian equilibria*, to be Bayesian equilibria of the game in which strategies in every period are informative.

The first decision variables of interest are reports in the feedback stage. Observe that a high quality firm will always report the same thing in the reporting and feedback stages. A normal firm that has a scoop in the reporting stage will therefore never contradict its initial report, because doing so would reveal it to be a normal type for sure. A firm that did not have a scoop, on the other hand, will prefer to report truthfully, since if it does not, then any consumers that see exogenous feedback will know that it is a normal type. Lemma 1 shows formally that these strategies are the unique equilibrium of the reporting stage.

Lemma 1 *In any informative Bayesian equilibrium, a firm j that had the scoop in the reporting stage will report $\tilde{n}_j = n_j$ in the feedback stage; a firm j that did not have the scoop in the reporting stage will report $\tilde{n}_j = \omega$ in the feedback stage.*

Proof. *See Appendix B.* ■

We now consider the equilibrium strategy of the firm j that has a scoop in the reporting stage. The feedback which determines $\hat{\theta}_{i-j}$ for a consumer $i \in I_j$ can potentially come from two places: reports of firms other than j and exogenous feedback. Since all other firms report truthfully in the feedback stage, any one of these sources will show the consumer the true state ω with certainty. The probability that any consumer sees such feedback is:

$$\tilde{\mu} = 1 - (1 - \mu)(1/2)^{J-1}. \tag{4}$$

(Recall that the probability a consumer reads a paper without a scoop is 1/2.) Let $\hat{\lambda}^L(n; s)$ and $\hat{\lambda}^R(n; s)$ represent a consumer's posterior on a firm's quality conditional on learning that the true state is L or R respectively. Let $\hat{\lambda}^0(n; s, \theta)$ represent the posterior conditional on receiving no feedback. (These depend only on the reporting-stage strategy s and not the strategy profile $\mathbf{s}$; $\hat{\lambda}^{L}(n;s)$ and $\hat{\lambda}^{R}(n;s)$ do not depend on θ because the prior has no effect once consumers have learned the true state).

The firm's reporting stage problem is thus to choose n to maximize:

$$\hat{\theta}(d)\,\tilde{\mu}\hat{\lambda}^{R}(n;s)+\left(1-\hat{\theta}(d)\right)\tilde{\mu}\hat{\lambda}^{L}(n;s)+(1-\tilde{\mu})\,\hat{\lambda}^{0}(n;s,\theta)\,. \tag{5}$$

where $\hat{\theta}(d)$ is the firm's own posterior on the true state after seeing its information d. Note that since firms have a neutral prior we have that $\hat{\theta}(L)=(1-\pi)$ and $\hat{\theta}(R)=\pi$ independently of θ.

To characterize the equilibrium solution to Equation 5, we require two intermediate steps. The first is to define more formally the firm's reporting-stage strategy, s. A possibly mixed strategy for reporting $n\in\{L,R\}$ conditional on having seen signal d is, in principle, defined by two probabilities of distortion: the probability the paper reports $n=R$ after seeing $d=L$ and the probability it reports $n=L$ after seeing $d=R$. We show in the proof of Proposition 1 below, however, that in an equilibrium where it strictly prefers to build a reputation for quality, the firm will only distort in one direction or the other with positive probability.

We will therefore represent the firm's strategy by a single number $\rho\in[-1,1]$, where positive values indicate bias toward R and negative values indicate bias toward L:

$$\begin{aligned}\rho &= \Pr(n=R|d=L)\ \text{ if } \rho\geq 0 \\ &= \Pr(n=L|d=R)\ \text{ if } \rho<0.\end{aligned} \tag{6}$$

This provides a single-dimensional characterization of firm strategies, and a convenient metric of bias.

Completing the characterization of the equilibrium requires showing that an equilibrium strategy ρ^{*} exists, is unique, and is informative. We also show that $\rho^{*}\in[0,1)$.

Proposition 1 *The model has a unique informative Bayesian equilibrium. In every period:*

- *The firm that has the scoop has equilibrium bias* $\rho^* \in [0, 1)$ *in the reporting stage and repeats its report in the feedback stage;*
- *All other firms report truthfully in the feedback stage.*

Proof. *See Appendix B.* ■

As emphasized in the introduction, existing theories of media bias assume that bias arises because some agents in the economy—consumers, reporters, or governments—are better off when media firms distort the truth of their reports. In our model, by contrast, bias arises despite the fact that consumers care only about obtaining accurate information and firms care only about profits. While these actors may in fact have a taste for bias, our model demonstrates that the mere existence of bias cannot be taken as evidence for such a preference.

To highlight this fact, we show in the next proposition that bias can arise even when it makes all agents in the economy strictly worse off, and firms would prefer to commit ex ante to report their signals truthfully.

Proposition 2 *Suppose that* ρ^* *is bounded away from zero in a neighborhood of* $\lambda = 0$*. Then for small* λ *both consumers and firms would be strictly better off if all firms were required to report their signals truthfully.*

Proof. *See Appendix B.* ■

That bias reduces the welfare of consumers is immediately apparent. As we have modeled it, bias is pure distortion and so adds noise to the firm's signal—this reduces the information value of the signal. It also means that consumers learn less about the firm's true type, and so make less informed decisions when they are old.

The effect of eliminating bias on the firm is somewhat more subtle. High-quality firms are clearly better off, since consumers are more likely to learn the firm's quality. Normal firms face a trade-off: their signal is more informative and so they attract more young consumers, but these consumers are also more likely to learn their true quality. When λ is small, however, the learning effect is small and these firms are also strictly better off.

5 Determinants of Bias

In this section, we analyze the determinants of bias in the equilibrium of Proposition 1. Three key intuitions emerge from this discussion. First, the firm's desire to maintain a reputation for quality leads equilibrium slant to be correlated with consumers' prior beliefs. Second, bias will be greater the weaker is the ex post feedback consumers receive about the true state of the world. Finally, bias will be smaller the more competitive is the market.

To frame this discussion, imagine a war between a foreign army and a domestic insurgency somewhere overseas. Different parties make official statements giving their own points of view—the army presents itself as valiantly fighting off terrorists while domestic sources claim it is recklessly targeting civilians. The situation on the ground is dangerous for journalists, and so independent information is fragmentary and difficult to verify.

Suppose that in this world, as in the model, there are two types of news sources. A few are high quality, with an exceptional ability to tell the difference between a reliable source and an unreliable one, to parse the language of official reports, and to piece together an accurate account of the events that actually transpired. The rest have fewer resources, are less skilled, and are left uncertain about what actually took place.

5.1 Consumer Priors

Proposition 3 characterizes the relationship between consumer priors and equilibrium bias. In particular, it shows that the firm's report is more right-biased the greater is the consumer's prior θ. Proposition 3 also shows that there is a region of θ close to .5 where the firm reports with no bias.

Proposition 3 *The equilibrium bias ρ^* is continuous and weakly increasing in θ. In particular, there exists $\theta^* \in [.5, \pi)$ such that:*

- *If $\theta \in [.5, \theta^*]$, $\rho^* = 0$;*
- *If $\theta \in (\theta^*, \pi)$, $\rho^* > 0$, and ρ^* is strictly increasing in θ.*

Proof. *See Appendix B.* ■

To see the intuition for this result, consider a consumer who has a strong prior belief that the army does not target civilians. Assuming for a moment that all news sources do their best to truthfully report what took place, her prior will lead her to expect high quality news sources to consistently produce reports showing low civilian casualties and emphasizing efforts the army is taking to avoid them, and ordinary news sources to be more likely to print stories casting the army in an unfavorable light. When she sees reports that contradict her prior beliefs, she will infer that the news source is less likely to be high quality.

A newspaper editor whose readers think that the army probably does not target civilians will think twice before printing reports of a high civilian casualty toll or of malicious behavior by soldiers, even if the editor knows these reports to be true. An editor whose readers believe the army *does* recklessly target civilians will play up the civilian deaths and de-emphasize steps the soldiers took to avoid them.

The intuition for the proposition can be seen more formally by returning to Equation 5. The only term in this expression that depends directly on consumers' prior belief is $\hat{\lambda}^0 (n; \rho, \theta)$. We show that $\hat{\lambda}^0 (R; \rho, \theta)$ is increasing in θ and $\hat{\lambda}^0 (L; \rho, \theta)$ is decreasing in θ. Increasing θ thus makes reporting R relatively more attractive.

As the probability of feedback $\tilde{\mu}$ goes to zero (i.e. when $J = 1$ and $\mu \to 0$), we obtain an even starker result. In this case, the firm reports whatever value of n will maximize $\hat{\lambda}^0 (n; \rho, \theta)$. Note that a consumer expects a high quality firm to report R with probability θ and L with probability $1 - \theta$. Unbiased reporting by normal firms thus cannot be an equilibrium, since consumers would expect a normal firm to report R with probability $\pi\theta + (1-\pi)(1-\theta) < \theta$ and so consumers would judge the firm's quality to be strictly higher when it reported R. The only equilibrium is thus the value $\rho_m^* > 0$ where $\hat{\lambda}^0(L; \rho^*, \theta) = \hat{\lambda}^0(R; \rho^*, \theta)$. This ρ_m^*, which we will refer to as the *minimum-feedback equilibrium* is strictly increasing in θ. We note that ρ_m^* is independent of λ, and in particular that for $\theta > \frac{1}{2}$, $\lim_{\lambda \to 0} \rho_m^* > 0$. Therefore, as is common in reputation-based models, even an arbitrarily small chance that the firm is high-type is sufficient to pin down its reporting incentives (Kreps and Wilson, 1982; Milgrom and Roberts, 1982).

5.2 Feedback

Recall that $\tilde{\mu}$ was defined to be the probability that consumers receive feedback either exogenously or from other firms. Proposition 4 shows that bias is decreasing in $\tilde{\mu}$.

Proposition 4 *The equilibrium bias ρ^* becomes weakly smaller as the probability of feedback $\tilde{\mu}$ increases. If $\rho^* > 0$ at a given θ and $\tilde{\mu}$, ρ^* becomes strictly smaller as $\tilde{\mu}$ increases. Bias disappears entirely if $\tilde{\mu}$ is sufficiently close to* 1.

Proof. *See Appendix B.* ■

Return once more to the position of a newspaper editor covering our hypothetical war. In addition to the daily reports of casualty tolls, there are reports of widespread torture of captured insurgents in army prisons. The editor knows that an independent commission is investigating these allegations, and that most of the paper's readers are likely to hear about the commission's report once it is published. How does this change the editor's incentives? *Ex ante*, consumers might think it highly unlikely that the army's soldiers would torture prisoners and thus downgrade their assessment of a newspaper's quality if it gave credence to such allegations. On the other hand, suppressing the information is risky because consumers may learn the truth *ex post* and thus realize that the paper was inaccurate after all. The temptation to distort the information should thus be disciplined by the fact that consumers will receive feedback after the fact on the true state of the world.

In terms of the model, imagine a firm that has observed an underlying signal $d = L$ and is considering whether to report truthfully. Suppose that in the absence of feedback, the consumer's posterior on quality would be the same regardless of what the firm prints. How would feedback change this? Because the firm has seen $d = L$, its own prior on the true state has shifted toward L. It therefore thinks it is relatively more likely that the consumer will learn that L is the true state. Since the consumer's ex post estimate of quality will be higher if the feedback matches the firm's report, this will make the firm prefer to truthfully report $n = L$. The more likely consumers are to learn the true state, the stronger this incentive will be.

Figure 3 illustrates the effects of priors and feedback on bias graphically. With no feedback, there is always bias in equilibrium, with the amount increasing in θ. With some level of feedback, there is a range of priors close to $\frac{1}{2}$ such that no bias occurs. For priors that exceed this range, bias is strictly increasing in θ. Moreover, there is always more bias in the minimum-feedback equilibrium than in the equilibrium with some feedback.

5.3 Competition

Proposition 4 also determines the effect of competition on bias. Since firms that did not have a scoop in the reporting stage always report the true state truthfully in the feedback stage, and since increasing the number of firms strictly increases the likelihood that a given consumer will read such a report, increasing the number of competitors simply increases $\tilde{\mu}$. Proposition 4 thus implies directly that increased competition reduces bias. This proposition and the definition of $\tilde{\mu}$ also imply that bias disappears entirely in the limit as the number of firms grows large. We state these results as Corollary 1.

Corollary 1 *The equilibrium bias ρ^* becomes weakly smaller as the number of firms increases, and becomes strictly smaller if $\rho^* > 0$ at a given θ and number of firms. Furthermore, for any prior belief θ, there exists some number of firms $J^*(\theta)$ such that $\rho^* = 0$ for all $J \geq J^*(\theta)$.*

In terms of our example, suppose that the hypothetical newspaper editor again faces the choice of whether or not to print the allegations of torture. In this scenario, there is no independent commission report that consumers will see on their own, but there are a large number of competing news outlets. Since competitors interested in beefing up their market share would be quick to point out any inconsistencies or omissions in the editor's coverage, the gain to distorting information in the direction of consumer beliefs would be severely limited. Thus, the more competitive the market, the more likely the editor would be to report the truth.

One point we wish to stress is that it is competitive market forces and not a mechanical increase in the ability of firms to reach consumers that is at the heart of Corollary 1. One way to make this point clear is to ask how increasing the number of newspapers would affect bias if all the newspapers were owned by a single firm. We assume that all jointly owned newspapers have the same quality, and that $\lambda \in (0, 1)$ is the consumer's prior belief that a firm is high quality. We also focus on the case where the probability of exogenous feedback is small, and so a monopoly firm's strategy is close to the minimum-feedback equilibrium defined in section 5.1. The result is stated in Proposition 5.

Proposition 5 *Suppose all newspapers are jointly owned. Then for μ sufficiently close to zero, there is an equilibrium where increasing the number of papers J has no effect on the level of equilibrium bias ρ^*.*

Proof. *See Appendix B.* ■

To see the intuition for this result, suppose for example that there are two jointly owned newspapers and that paper 1 has the scoop in period t and reports L. Consider paper 2's decision in the feedback stage after learning that $\omega = R$, assuming we are in a candidate equilibrium in which consumers expect paper 2 to always report truthfully (as in Lemma 1). Since there is a very low probability of exogenous feedback, the beliefs of consumers who did not also read paper 1 will be essentially unchanged by paper 2's report. For consumers who did read paper 1, their posteriors will be $\hat{\lambda} = 0$ if paper 2 reports R and $\hat{\lambda} > 0$ if paper 2 reports L. Truthful reporting by firm 2 thus cannot be an equilibrium.

On the other hand, it *is* an equilibrium for paper 2 to simply repeat the report of paper 1. To see this, consider first the set of consumers who read paper 2 but not paper 1. Since feedback is minimal, the only way their assessment of the firm's quality would change based on paper 2's report is if paper 2 reports either R or L relatively more often than a high-quality firm on average. But we know paper 1's reports have exactly the same distribution as a high-quality firm's signals because there is essentially no exogenous feedback. So paper 2's reports cannot affect the posteriors of these consumers. The other group of consumers are those that read both papers. We assume that in the zero probability event where paper 2 contradicts the reports of paper 1, they believe the paper to be normal with probability 1—i.e. $\hat{\lambda} = 0$. We know $\hat{\lambda} > 0$ otherwise. So paper 2 prefers to follow the equilibrium strategy.

6 Heterogeneous Priors

In the model presented above, all consumers have identical beliefs about the state of the world, and consequently any two consumers who see the same report and feedback will make identical inferences about newspaper quality. In addition, all firms in a given market will report with the same bias in equilibrium. We argued above that this is a reasonable starting point for thinking about differences in bias across markets—for example, why Al Jazeera and CNN differ, and why markets with competing newspapers might have lower bias overall. But in many key settings of interest, we see firms with different biases competing in the same market. And as Figure 1 suggests, consumers with different pre-existing intuitions form different assessments of the quality of these alternative media outlets.

We show in this section that the basic intuitions developed above extend to the case of heterogeneous priors. In markets where consumer beliefs are polarized and feedback is weak, there exist segmented equilibria where firms effectively serve only one side of the market or the other.

We show that in these equilibria, the comparative statics developed above with respect to priors, feedback, and competition continue to hold. We also show that there is an additional effect of competition in markets with heterogeneous priors—when competition is high enough, segmented equilibria disappear and all firms serve both types of consumers, reporting with no bias.

The key conceptual difference between the cases of heterogeneous and homogenous priors is that when priors are heterogeneous, consumers' beliefs about a firm's strategy determine the composition of its readers, and hence its incentives to bias. In particular, a firm that is right-biased will attract relatively more right-leaning readers and a firm that is left-biased will attract relatively more left-leaning viewers.

Note that this is similar in some respects to a Hotelling model of horizontal differentiation where firms are differentiated according to their expected bias. There are, however, fundamental differences. First, *all* consumers in the model would prefer a firm that reported with less bias, and the locations firms can "choose" are limited by their inability to commit. Second, consumers are allowed to choose multiple newspapers, and under our simplified utility structure choosing one paper does not reduce a consumer's utility from reading a second paper. The latter simplification rules out the possibility that increasing the number of firms causes the market to become more fragmented as in Mullainathan and Shleifer (2003).

For ease of exposition we consider a simple model of heterogeneous priors, with two equal-sized groups of consumers, denoted L and R. Consumers in group R have prior belief θ, and consumers in group L have prior belief $(1-\theta)$, where $\theta \in \left[\frac{1}{2}, 1\right)$.[22] To define a segmented equilibrium, let $\rho^*(\theta, \tilde{\mu})$ represent the equilibrium strategy defined in Proposition 1 in a market where all consumers have prior beliefs θ and the probability of feedback is $\tilde{\mu}$. It follows from symmetry that a market with only consumers of type $1-\theta$ would have equilibrium bias $-\rho^*(z, \tilde{\mu})$. Define s_j^t as before to be firm j's strategy in period t.

[22] The model could be extended to allow more consumers of one type or to relax the restriction that the beliefs are symmetric around 1/2. Doing so would change the details of the equilibrium (roughly, as the fraction of consumers in group R increases or their beliefs become more extreme, equilibrium strategies will tend to become more right-biased), but none of the key comparative statics results would change.

Definition 1 *A strategy profile* $\mathbf{s}$ *is a* ***segmented equilibrium*** *if: for each* j *either* $s_j^t = \rho^*(\theta, \tilde{\mu})$ *for all* t *or* $s_j^t = -\rho^*(\theta, \tilde{\mu})$ *for all* t*; consumers in group* R *base their actions only on the reports of firms playing* $\rho^*(\theta, \tilde{\mu})$ *while consumers in group* L *base their actions only on the reports of firms playing* $-\rho^*(\theta, \tilde{\mu})$*.*

Our first proposition shows that such equilibria can exist in markets where beliefs are polarized (i.e. θ is far from 1/2) and feedback and competition are weak. We define $\rho_0(\theta, \tilde{\mu}) = \lim_{\lambda \to 0} \rho^*(\theta, \tilde{\mu})$ which must exist by the continuity of $\Delta(L; \rho^*, \theta, \tilde{\mu})$ in λ.

Proposition 6 *Pick* $(\theta^*, \tilde{\mu}^*)$ *such that the consumers in group* L *strictly prefer not to follow the report of a low-quality firm playing strategy* $\rho_0(\theta^*, \tilde{\mu}^*) > 0$*. Then for any* $(\theta, \tilde{\mu})$ *such that* $\theta \geq \theta^*$ *and* $\tilde{\mu} \leq \tilde{\mu}^*$*, there exists* $\lambda^* > 0$ *such that a segmented equilibrium exists for any* $\lambda < \lambda^*$*.*

Proof. *See Appendix B.* ■

The proposition requires that the ex-ante fraction of firms that are high quality, λ, is small. To see the intuition for this, note that when λ is large both type-L and type-R consumers would choose to base their actions on a firm's report regardless of ρ^* (i.e. the strategies they expect low-quality firms to play) since they believe that the firm is probably high quality. This would mean immediately that segmentation would be impossible.

Because equilibrium bias in a segmented equilibrium is the same as in a homogeneous equilibrium with either all θ or all $1 - \theta$ consumers, the following corollary is an immediate implication of Propositions 3 and 4 and Corollary 1. This result uses the fact that although there may be multiple segmented equilibria that exist for a given market (i.e. depending on the identity of the firms serving each type of consumer), the magnitude of equilibrium bias is uniquely defined.

Corollary 2 *Consider* θ*,* μ*, and* λ *strictly in the interior of the parameter space such that segmented equilibria exist. Then the magnitude of bias* $|\rho^*|$ *in any segmented equilibrium is:*

1. *strictly increasing in consumer beliefs* θ*;*

2. *strictly decreasing in the probability of feedback μ;*

3. *and strictly decreasing in the number of firms J.*

Our final result shows that competition and feedback not only reduce bias within the set of segmented equilibria, but, when they are strong enough, make truthful reporting the unique equilibrium.

Proposition 7 *There exists μ^* and J^* such that if either $\mu > \mu^*$ or $J > J^*$, the unique informative Bayesian equilibrium is for all firms to report with no bias.*

Proof. *See Appendix B.* ■

7 Evidence on the Determinants of Bias

In the model above, we showed that three factors play a key role in determining the direction and strength of bias: consumer priors, *ex post* feedback, and competition. In this section, we review existing evidence and present new evidence on each of these implications.

7.1 Consumer Priors

A large body of anecdotal evidence supports the connection highlighted by Proposition 3 between consumers' prior beliefs and media firms' slant. Consider, for example, Ames' (1938) description of the problem faced by southern newspaper editors in their coverage of lynching: "As individuals, they are unanimously opposed to mob violence but, as editors who are caught in the general atmosphere of a given trade territory, they do not reflect their own ideas but those of the people upon whose goodwill their papers depend for revenue." The result of this pressure was that southern editorials in the period almost universally condoned lynchings.

A more recent example is the reported difference in coverage of the war in Iraq between U.S. networks and Arabic-language news channels such as Al Jazeera. As Lieutenant Josh Rushing, an

American press officer, explains in the documentary *Control Room*, "It benefits Al Jazeera to play to Arab nationalism because that's their audience, just like Fox plays to American patriotism for the exact same reason" (Turan, 2004).

Even *within* a given firm slant can vary depending on the audience. For example, CNN's domestic cable channel broadcasts quite different content from CNN International, which is broadcast worldwide. Chris Cramer, president of CNN International, writes that their audience "expects us to have a non-U.S. viewpoint." The difference is also illustrated by coverage in the aftermath of September 11: the domestic channel prominently displayed an American flag during its broadcasts while the international broadcasts quickly dropped the flag (Kempner, 2001).

Turning to more systematic evidence, newspaper endorsements of presidential candidates display a pattern of conformity to local political opinion. As Figure 2 shows, in the 2000 U.S. presidential election, Bush's share of the two-party vote was considerably lower in richer states (Glaeser, Ponzetto and Shapiro, 2004). Bush received almost 60 percent of the two-party vote among states in the lowest income quartile, as against just over 40 percent in the states in the highest quartile. As the figure illustrates, newspaper endorsements displayed a similar pattern, with almost 90 percent endorsing Bush in the bottom quartile and less than 55 percent in the top quartile. Although this graph is by no means conclusive, it certainly suggests a significant connection between consumer beliefs and media slant.

Existing work in political science also suggests a correlation between the editorial position of newspapers and the views of their readers. For example, Dalton, Beck and Huckfeldt (1998) show survey evidence from the 1992 presidential election suggesting that the editorial stance of local newspapers is correlated with local perceptions of candidates. Erikson (1976) documents a similar relationship using aggregate data on voting patterns in the 1960s. In Gunther's (1992) analysis of national survey data, he finds that only two percent of respondents had political views categorized as "very distant" from those of their primary newspaper.

7.2 Feedback

As Proposition 4 shows, our model predicts that *ex-post* feedback about the true state of the world will tend to reduce media bias. An extreme example of an issue where feedback is immediate and unambiguous is weather reporting. Although the notion of bias in weather reporting seems strange, consumers certainly have strong prior beliefs about the next day's weather—a forecaster who predicts snow in New Mexico in July would be viewed with suspicion. But since feedback is immediate, this should not deter her from making such a prediction if she truly believes it to be the most likely scenario.

In fact, studies of weather forecasters' predictions reveal excellent reliability. Probability forecasts match up well with observed relative frequencies; i.e., a forecast of a 20 percent chance of precipitation tends to be followed by precipitation roughly 20 percent of the time. Additionally, reported confidence intervals for temperatures show nearly exact coverage (Murphy and Winkler, 1977 and 1984). Given the fact that the weather is known with certainty soon after forecasters make their predictions, it is not surprising that we find little evidence of bias in predicting the weather.

Of course, weather differs from politics not only because of the strength of feedback but because people take actions with concrete and immediate consequences in response to forecasted conditions. Some authors (such as Glaeser, 2004) have argued that psychological biases will have less influence on decisions with larger stakes. The presence of such a force could allow theories based on confirmatory bias to accommodate the observation that weather reporting is relatively unbiased.

We therefore turn next to a forecasting environment with rapid feedback in which emotions run high and concrete stakes tend to be low: sports picking by local newspaper columnists. We draw on data collected by Boulier and Stekler (2003) on *New York Times* sports editors' predictions from 1994-2000. For each game, the dataset contains the opening betting line (as published in *USA Today*) and the editors' picks. If bias is driven primarily by consumers' desire to hear felicitous reports, a natural hypothesis is that the *Times* would favor the New York teams—the Jets and the

Giants—to win (relative to the betting market's expectation). In contrast, because outcomes are observed soon after reports are made, our model predicts little such bias in this context.

To investigate this issue we calculate a measure $\hat{\delta}_i$ of the experts' slant towards team i by estimating a regression model of the form

$$win_j = \alpha + \delta_i \left[(home_j = i) - (away_j = i) \right] + \gamma \left(line_j \right) + \varepsilon_j$$

where win_j denotes whether the editor picked the home team to win game j, $home_j$ indexes the home team in game j, $away_j$ indexes the visiting team in game j, and $line_j$ is a vector of dummy variables representing deciles of the betting line.

Figure 4 presents graphically our estimates of δ_i for each team, measured relative to the Seattle Seahawks (the experts' least favored team over this time period). As the figure shows, the Giants are picked more often than average and the Jets less often, but there is no evidence of overall favoritism toward the New York teams. This fact is surprising if we start from taste-based theories of bias, but is consistent with the implications of Proposition 4 above.

A similar test can be conducted using data collected by Avery and Chevalier (1999) on picks of six experts published in the *Boston Globe* between 1983 and 1994. Results are in Figure 5, and details of the exercise are provided in notes to the figure. While the writers are more favorable to home team (the Patriots) than average, the Patriots are not the experts' most favored team, and are treated comparably to many other teams in the league. Overall, this data shows limited evidence for a taste for confirmation as a driving factor in sports reporting.

Evidence from financial reporting also suggests a role for feedback in limiting media bias. Lim (2001) presents evidence indicating that analysts' forecasts of corporate earnings are more optimistically biased for smaller firms and for firms with more volatile historical earnings and stock returns (see also Das, Levine and Sivaramakrishnan, 1998). Although Lim interprets these findings as evidence that analysts make optimistic forecasts so as to win favor with firms and obtain access to non-public information, we propose our model as an alternative explanation. When earnings are less volatile, inaccurate reporting is more easily detected, and so analysts concerned with their reputation for high-quality reports will be less likely to bias their forecasts.

Relatedly, several authors have argued that biases in earnings forecasts become less severe as the length of time between the publication of the forecast and the earnings announcement decreases (see Kang, O'Brien and Sivaramakrishnan, 1994 and Raedy, Shane and Yang, 2003). Again, a model with reputational concerns offers a possible interpretation of this fact: When the announcement comes quickly on the heels of the forecast, bias is more likely to be detected and to influence consumer decisions about which forecast to purchase in the future.

7.3 Competition

Because competition increases the likelihood that erroneous reports will be exposed *ex-post*, our model predicts that added competition will tend to reduce bias. The reaction to Dan Rather's report on President Bush's service in the National Guard, discussed in more detail in section 2.1, illustrates the role that competing media outlets can play in exposing flaws in journalism.

Anecdotes about the impact of Al Jazeera's relatively independent reporting on media in the Arabic-speaking world provide another example of the competitive discipline we model. As Otis (2003) reports, "Many experts contend that Egyptian newspapers have improved dramatically in recent years. During the Six-Day War against Israel in 1967, the heavily censored press largely ignored battlefield defeats. Today, Al-Jazeera and other television stations beam raw images of military conflicts into people's homes, preventing newspapers from straying too far from the truth."

To take another example, when an American civilian was beheaded by militants in Iraq, reporting of the story was more common in countries with competitive media environments. In Syria, where all local press and television are state-owned (Djankov, McLiesh, Nenova and Shleifer, 2003), newspapers completely ignored this event. By contrast, in Lebanon, which has a relatively competitive press, most newspapers did report on the beheading (Associated Press, 2004). This fact seems to support the view that suppression and distortion of information are less attractive when competition makes the truth likely to come out.[23]

For a more quantitative investigation of the effects of competition on media bias, we have obtained data from the 2000 Local News Archive (Kaplan and Hale, 2001). The dataset encodes the characteristics of local election news coverage broadcast between 5:00pm and 11:35pm during the 30 days prior to the general election on November 7, 2000. It covers 74 stations in 58 of the top 60 media markets in the US. Most importantly for our purposes, it contains a coding of the number of seconds of speaking time given to George W. Bush and Al Gore. By calculating the total number of seconds given to each candidate by each station i, we can then construct the following measure of biased treatment:

$$bias_i = \left(\frac{bush_i}{bush_i + gore_i} - \frac{1}{2} \right)^2$$

where x_i denotes the number of seconds given to candidate x by station i. This measure takes on a value of 0 when Bush and Gore received equal time in local news coverage of the election, and a value of .25 when only one candidate is given coverage. The average of $bias_i$ across the stations in a market will serve as our measure of bias.

Given this measure, we can investigate whether the degree of bias is lower in markets with a greater number of local news broadcasts, as predicted by the model. As Figure 6 shows, this is indeed the case. In markets with 4 stations or fewer broadcasting the local news, the above-described measure of bias is more than six times larger than in markets with 6 news broadcasts. Table 1 presents this fact in a regression framework. As column (1) of the table illustrates, one additional television station is associated with a statistically significant reduction in bias of about .006, which is equivalent to about one-third of a standard deviation. As column (2) shows, this effect is robust to

[23] Another example of the role of competition is the coverage of the allegations of torture in the Abu Ghraib prison in Iraq. The CBS program *60 Minutes* was the first to obtain photos from the prison, but it delayed broadcasting them for two weeks. The incentive to suppress the photos in this case was not consumer beliefs but direct pressure from the government—Chairman of the Joint Chiefs of Staff Richard Myers had personally asked Dan Rather not to broadcast the photos. But what led them to finally be aired was competition: once CBS learned that Seymour Hersh was working on the same story for the *New Yorker*, they decided to put the report on the air (Folkenflik, 2004).

the inclusion of controls for Census region, so it is not merely driven by geographic differences in the thickness of markets. Column (3) highlights that including population, an important determinant of the number of local broadcasts (the correlation between log(population) and the number of news broadcasts is about .5), does not substantially reduce the size of the competition effect or eliminate its statistical significance. Finally, column (4) shows that the effect is robust to an additional control for income per capita, which could also drive differences across locations in the competitiveness of the news market.[24]

Several existing studies of bias in reporting also show effects of competition consistent with our model's predictions. Dyck and Zingales (2003) argue that newspapers put less "spin" on their reports of company earnings when many alternative sources of information are available. Lim (2001) presents evidence from analysts' earnings forecasts suggesting that bias is lower the more analysts are providing reports on a given company.[25] Gentzkow, Glaeser and Goldin (2004) document that the emergence of independent (i.e. non-party-affiliated) newspapers in the United States was faster in larger cities, suggesting a role for competition in encouraging the growth of more informative news outlets.

8 Conclusions

The model in this paper presents a new way to understand media bias. Bias in our model does not arise from consumer preferences for confirmatory information, reporters' incentives to promote their own views, or politicians' ability to capture the media. Instead, it arises as a natural consequence of firms' desire to build a reputation for accuracy, and in spite of the fact that eliminating bias could make all agents in the economy better off.

An advantage of our model is that it generates sharp predictions about where bias will arise and when it will be most severe. We wish to highlight two policy implications of these results. The first

[24]The estimated effect of competition is also robust to controlling for the average total amount of candidate speaking time aired by stations in each market.

[25]See also Firth and Gift (1999).

concerns the regulation of media ownership. In the current debate over FCC ownership regulation in the U.S., the main argument in favor of limits on consolidation has been the importance of "independent voices" in news markets. Proposition 5 offers a new way to understand the potential costs of consolidation: independently owned outlets can provide a check on each others' coverage and thereby limit equilibrium bias, an effect that may be absent if the outlets are jointly owned.

As a second implication, the effect of competition described by Corollary 1 has important implications for the conduct of public diplomacy. The U.S. government is currently engaged in a debate about the most effective way to counter what it sees as anti-American bias in the Arab media, especially Al Jazeera. Efforts along these lines have included condemnation of Al Jazeera by top U.S. officials (Rumsfeld, 2001), appeals to the Emir of Qatar (who sponsors the network) to change the tone of Al Jazeera's coverage (Campagna, 2001), and the recent closing of Al Jazeera's Baghdad office by the U.S.-backed Alawi government in Iraq. Our model suggests a different approach: supporting the growing competitiveness of the Middle Eastern media market and in particular increasing the availability of alternative news sources in local languages. Aside from the direct effect on the beliefs of those who watch these sources (Gentzkow and Shapiro, 2004), introducing more news outlets could have the effect of disciplining existing stations and reducing the overall amount of bias in the region.

References

[1] Fouad Ajami. What the Muslim world is watching. *New York Times Magazine*, page 48, November 18 2001.

[2] Eric Alterman. *What Liberal Media? The Truth About Bias and the News.* Basic Books, February 2003.

[3] Jessie Daniel Ames. Editorial treatment of lynchings. *The Public Opinion Quarterly*, 2(1):77–84, January 1938.

[4] Ken Auletta. Battle stations: How long will the networks stick with the news? *New Yorker*, December 10 2001.

[5] Christopher Avery and Judith Chevalier. Identifying investor sentiment from price paths: The case of football betting. *Journal of Business*, 72(4):493–521, 1999.

[6] Abhijit Banerjee and Rohini Somanathan. A simple model of voice. *Quarterly Journal of Economics*, 116(1):189–227, February 2001.

[7] David P. Baron. Persistent media bias. *Stanford Graduate School of Business Research Paper*, 1845, February 2004.

[8] Dan Barry, David Barstow, Jonathan D. Glater, Adam Liptak, and Jacques Steinberg. Correcting the record: Times reporter who resigned leaves long trail of deception. *New York Times*, page 1, May 11 2003.

[9] Marco Battaglini. Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401, July 2002.

[10] Timothy Besley and Andrea Prat. Handcuffs for the grabbing hand? media capture and government accountability. *London School of Economics Mimeograph*, April 2004.

[11] Bryan L. Boulier and H. O. Stekler. Predicting the outcomes of national football league games. *International Journal of Forecasting*, 19:257–270, 2003.

[12] Adam Brandenburger and Ben Polak. When managers cover their posteriors: Making the decisions the market wants to see. *RAND Journal of Economics*, 27(3):523–541, Autumn 1996.

[13] Joel Campagna. Between two worls: Qatar's Al-Jazeera satellite channel faces conflicting expectations. *CPJ Press Freedom Reports*, October 2001. http://www.cpj.org.

[14] Mark Cooper, Gene Kimmelman, and Cheryl Leanza. Comments of Consumers Union, Consumer Federation of America, Civil Rights Forum, Center for Digital Democracy, Leadership Conference on Civil Rights and Media Access Project. December 2001.

[15] Ann Coulter. *Slander : Liberal Lies About the American Right*. Three Rivers Press, September 2003.

[16] Russell J. Dalton, Paul A. Beck, and Robert Huckfeldt. Partisan cues and the media: Information flows in the 1992 presidential election. *The American Political Science Review*, 92(1):111–126, March 1998.

[17] Somnath Das, Carolyn B. Levine, and K. Sivaramakrishnan. Earnings predictability and bias in analysts' earnings forecasts. *Accounting Review*, 73(2):277–294, April 1998.

[18] Simeon Djankov, Caralee McLiesh, Tatiana Nenova, and Andrei Shleifer. Who owns the media? *Journal of Law and Economics*, 46(2), October 2003.

[19] Alexander Dyck and Luigi Zingales. The media and asset prices. *Harvard Business School mimeograph*, August 2003.

[20] Jeffrey C. Ely, Drew Fudenberg, and David K. Levine. When is reputation bad? October 2002. Harvard University mimeo.

[21] Jeffrey C. Ely and Juuso Välimäki. Bad reputation. *Quarterly Journal of Economics*, 118(3):785–814, August 2003.

[22] Robert S. Erikson. The influence of newspaper endorsements in presidential elections: The case of 1964. *American Journal of Political Science*, 20(2):207–233, May 1976.

[23] Michael Firth and Michael Gift. An international comparison of analysts' earnings forecast accuracy. *International Advances in Economic Research*, 5(1):56–64, February 1999.

[24] David Folkenflik. Iraq prison story tough to hold off on, CBS says. *Baltimore Sun*, page 1D, May 5 2004.

[25] Al Franken. *Lies and the Lying Liars Who Tell Them: A Fair and Balanced Look at the Right.* E P Dutton, August 2003.

[26] Matthew A. Gentzkow, Edward L. Glaeser, and Claudia D. Goldin. The rise of the fourth estate: How newspapers became informative and why it mattered. July 2004. Harvard University mimeograph.

[27] Matthew A. Gentzkow and Jesse M. Shapiro. Media, education, and anti-americanism in the muslim world. *Journal of Economic Perspectives*, Summer 2004. Forthcoming.

[28] Thomas W. Gilligan and Keith Krehbiel. Asymmetric information and legislative rules with a heterogeneous committee. *American Journal of Political Science*, 33(2):459–490, May 1989.

[29] Edward L. Glaeser. Psychology and the market. *American Economic Review*, 94(2):408–413, May 2004.

[30] Edward L. Glaeser, Giacomo A. M. Ponzetto, and Jesse M. Shapiro. Strategic extremism: Why republicans and democrats divide on religious values. *Harvard University Mimeograph*, August 2004.

[31] Jacob Glazer and Ariel Rubinstein. Motives and implementation: On the design of mechanisms to elicit opinions. *Journal of Economic Theory*, 79:157–173, 1998.

[32] Jacob Glazer and Ariel Rubinstein. Debates and decisions: On a rationale of argumentation rules. *Games and Economic Behavior*, 36:158–173, 2001.

[33] Jacob Glazer and Ariel Rubinstein. On optimal rules of persuasion. *Econometrica*, 2004. Forthcoming.

[34] Bernard Goldberg. *Bias : A CBS Insider Exposes How the Media Distort the News.* Perennial, January 2003.

[35] Simon Grant, Atsushi Kajii, and Ben Polak. Intrinsic preference for information. *Journal of Economic Theory*, 83:233–259, 1998.

[36] Albert C. Gunther. Biased press or biased public? attitudes toward media coverage of social groups. *The Public Opinion Quarterly*, 56(2):147–167, Summer 1992.

[37] Andrew Heyward. Cbs statement on bush memos. *CBSNews.com*, September 20 2004.

[38] Sok-Hyon Kang, John O'Brien, and K. Sivaramakrishnan. Analysts' interim earnings forecasts: Evidence on the forecasting process. *Journal of Accounting Research*, 32(1):103–112, Spring 1994.

[39] Martin Kaplan and Matthew Hale. *Local TV coverage of the 2000 general election*. Norman Lear Center, USC Annenberg School of Communication, February 2001.

[40] Matt Kempner. CNN on top of the world; CNN International crafts coverage with a global emphasis that appeals to viewers overseas. *Atlanta Journal-Constitution*, page 1G, October 7 2001.

[41] David D. Kirkpatrick and Geraldine Fabrikant. Changes at the times: Reactions; advertisers and Wall St. see an end to turmoil. *New York Times*, page B9, June 6 2003.

[42] Jonathan Koehler. The influence of prior beliefs on scientific judgments of evidence quality. *Organizational Behavior and Human Decision*, 56(1):28–55, 1993.

[43] Andrew Kohut. Media myopia: More news is not necessarily good news. *New York Times*, July 11 2004.

[44] David Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27:253–279, 1982.

[45] Vijay Krishna and John Morgan. A model of expertise. *Quarterly Journal of Economics*, 116(2):747–775, May 2001.

[46] Terence Lim. Rationality and analysts' forecast bias. *Journal of Finance*, 56(1):369–385, February 2001.

[47] Charles G. Lord, Mark R. Lepper, and Elizabeth Preston. Considering the opposite: A corrective strategy for social judgment. *Journal of Personality and Social Psychology*, 47(6):1231–1243, 1984.

[48] Charles G. Lord, Lee Ross, and Mark R. Lepper. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11):2098–2109, 1979.

[49] Paul Milgrom and John Roberts. Predation, reputation and entry deterrence. *Journal of Economic Theory*, 27:280–312, 1982.

[50] Arthur G. Miller, John W. McHoskey, Cynthia M. Bane, and Timothy G. Dowd. The attitude polarization phenomenon: Role of response measure, attitude extremity, and behavioral consequences of reported attitude change. *Journal of Personality and Social Psychology*, 64(4):561–574, 1993.

[51] Stephen Morris. Trade with heterogeneous prior beliefs and asymmetric information. *Econometrica*, 62(6):1327–1347, November 1994.

[52] Stephen Morris. The common prior assumption in economic theory. *Economics and Philosophy*, 11:227–253, 1995.

[53] Stephen Morris. Political correctness. *Journal of Political Economy*, 109(2):231–265, 2001.

[54] Sendhil Mullainathan and Andrei Shleifer. Media bias. *Harvard Institute of Economic Research Discussion Paper*, (1981), October 2002.

[55] Sendhil Mullainathan and Andrei Shleifer. The market for news. December 2003. Harvard University mimeo.

[56] Geoffrey D. Munro and Peter H. Ditto. Biased assimilation, attitude polarization, and affect in reactions to stereotyped-relevant scientific information. *Personality and Social Psychology Bulletin*, 23(6):636–653, 1997.

[57] Allan H. Murphy and Robert L. Winkler. Reliability of subjective probability forecasts of precipitation. *Applied Statistics*, 26(1):41–47, 1977.

[58] Allan H. Murphy and Robert L. Winkler. Probability forecasting in meteorology. *Journal of the American Statistical Association*, 79(387):489–500, September 1984.

[59] The Gallup Organization. *The 2002 Gallup Poll of the Islamic World: Executive Report*. The Gallup Poll Tuesday Briefing. Princeton, NJ, February 2002.

[60] John Otis. Arab media accused of bias; experts say most reports are accurate, but with anti-US slant. *Houston Chronicle*, page A25, April 5 2003.

[61] Peter G. Peterson, Kathy Bloomgarden, Henry Grunwald, David E. Morey, and Shibley Telhami. *Finding America's Voice: A Strategy for Reinvigorating U.S. Public Diplomacy*. Council on Foreign Relations, New York, 2003.

[62] Canice Prendergast. A theory of "yes men". *American Economic Review*, 83(4):757–770, September 1993.

[63] Associated Press. Arab media muted in coverage of beheading; many papers play down story, others ignore it. May 12 2004.

[64] Jana Smith Raedy, Philip Shane, and Yanhua Yang. Horizon-dependent underreaction in financial analysts' earnings forecasts. *University of Colorado Mimeograph*, September 2003.

[65] Dan Rather. Dan rather statement on memos. *CBSNews.com*, September 20 2004.

[66] Donald H. Rumsfeld. Secretary rumsfeld remarks at stakeout outside ABC TV. *Department of Defense News Transcript*, October 28 2001.

[67] Robert Satloff. Rethinking public diplomacy in the Middle East. *Weekly Standard*, 8(46), August 2003.

[68] Wing Suen. The self-perpetuation of biased beliefs. *Economic Journal*, 2004. Forthcoming.

[69] Kenneth Turan. Through the looking glass with Al Jazeera: Jehane Noujaim's documentary 'Control Room' challenges what some might expect of the news channel. *Los Angeles Times*, page E6, June 18 2004.

[70] William Underhill and Tara Pepper. The media: Losing trust in Auntie Beeb. *Newsweek*, September 29 2003.

Figure 1 *Political views and assessments of news media believability*

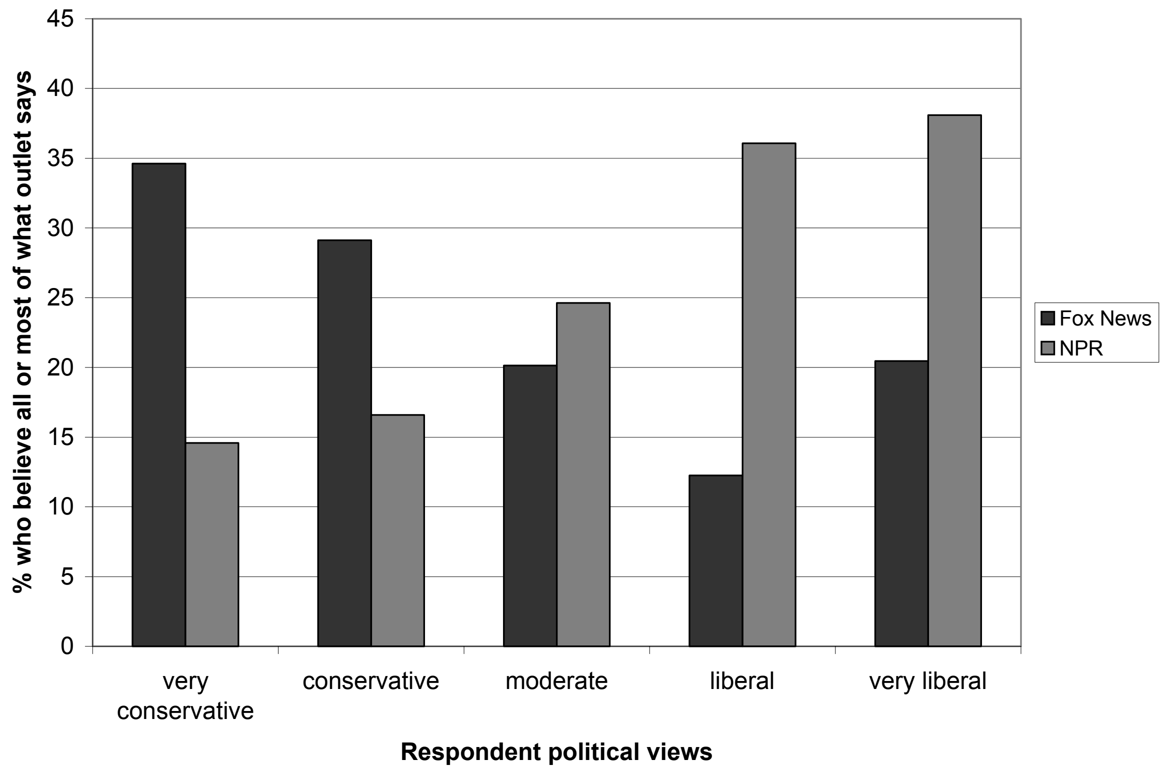

Notes: Data from the Pew Research Center for the People and the Press's 2002 News Media Believability Survey. Exact wording for survey question on respondent political views:

> In general, would you describe your political views as very conservative, conservative, moderate, liberal, or very liberal?

Exact wording for survey question on media believability:

> Now, I'm going to read a list. Please rate how much you think you can BELIEVE each organization I name on a scale of 4 to 1. On this four point scale, "4" means you can believe all or most of what the organization says. "1" means you believe almost nothing of what they say. How would you rate the believability of {The Fox News CABLE Channel / National Public Radio} on this scale of 4 to 1?

Figure 2 *Newspaper endorsements and ideology across U.S. states in the 2000 election*

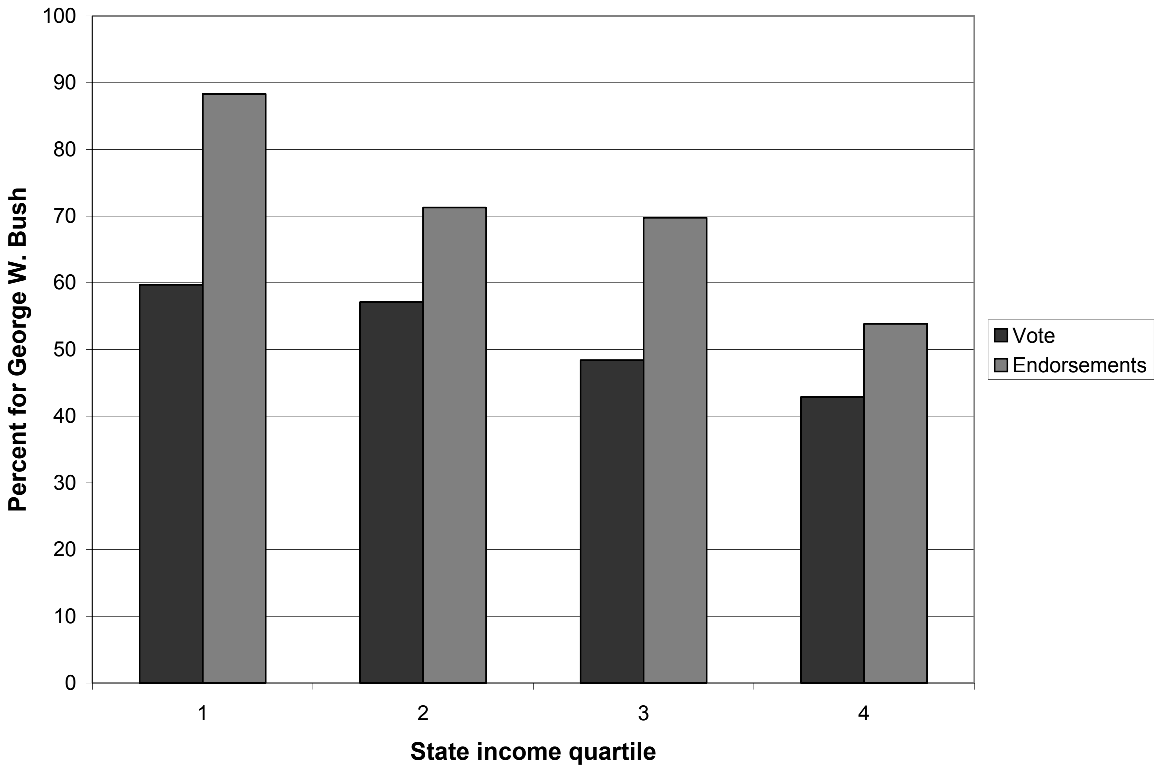

Notes: Data on voting behavior taken from Federal Election Commission (http://www.fec.gov/elections.html). Percent for Bush reflects percent of two-party vote. Data on state per capita income from USA Counties 1998 CD-ROM. Data on newspaper endorsements from Bush (www.georgewbush.com) and Gore (www.algore.com) official campaign sites; data posted at http://www.wheretodoresearch.com/Political.htm#Endorsements. Percent for Bush reflects percent among papers endorsing either Bush or Gore.

Figure 3 *A numerical example*

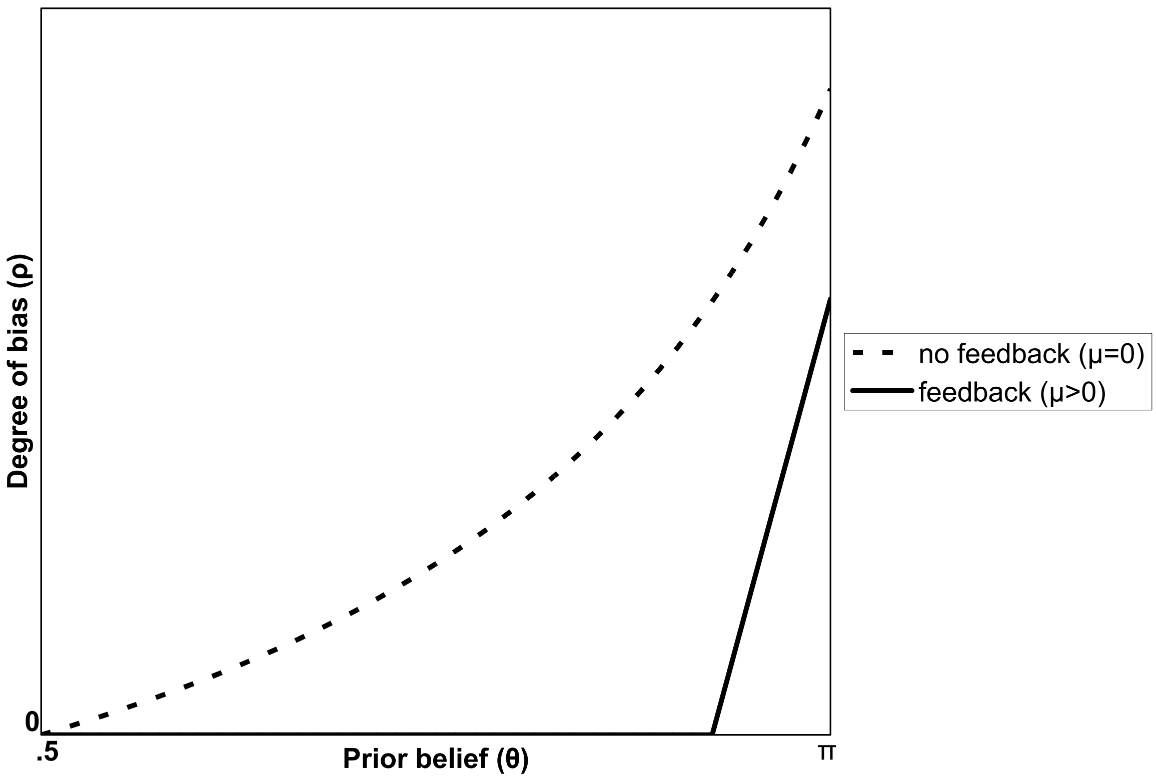

Notes: Vertical axis shows equilibrium probability of reporting R given that the newspaper with the scoop receives a signal of L. Horizontal axis shows prior probability θ. Dashed line shows equilibria for the limit case of no feedback ($\mu = 0$). Solid line shows equilibria for the case of a positive probability of feedback ($\mu > 0$).

Figure 4 *Sports picking by* New York Times *sports editor, 1994-2000*

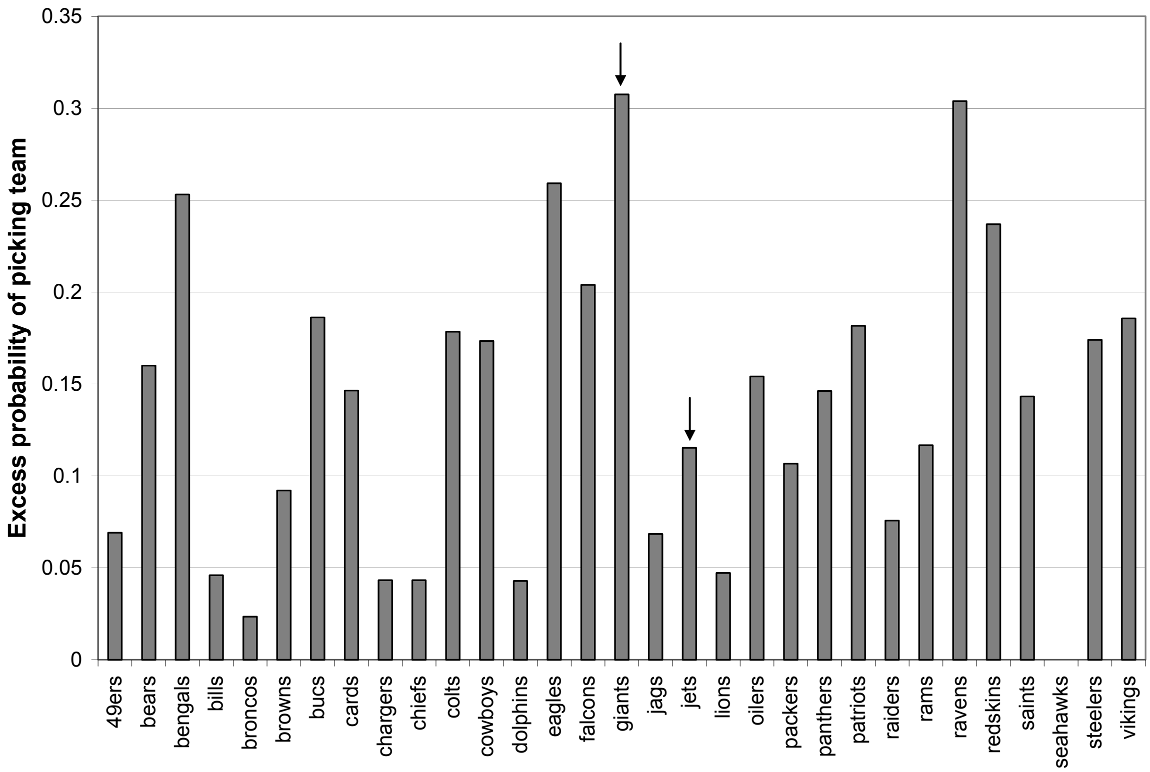

Notes: Data from Boulier and Stekler (2003). Dataset contains information on the picks of the New York Times sports editor for NFL games in the 1994-2000 seasons, as well as the outcome of the game and the betting line. The bar for team i represents the estimated coefficient $\hat{\delta}_i$ in a regression of the form

$$win_j = \alpha + \delta_i \left[(home_j = i) - (away_j = i) \right] + \gamma \left(line_j \right) + \varepsilon_j$$

where win_j denotes whether the editor picked the home team to win game j, $home_j$ indexes the home team in game j, $away_j$ indexes the visiting team in game j, and $line_j$ is a vector of dummy variables representing deciles of the betting line.

Figure 5 *Sports picking by* Boston Globe *columnists, 1983-1994*

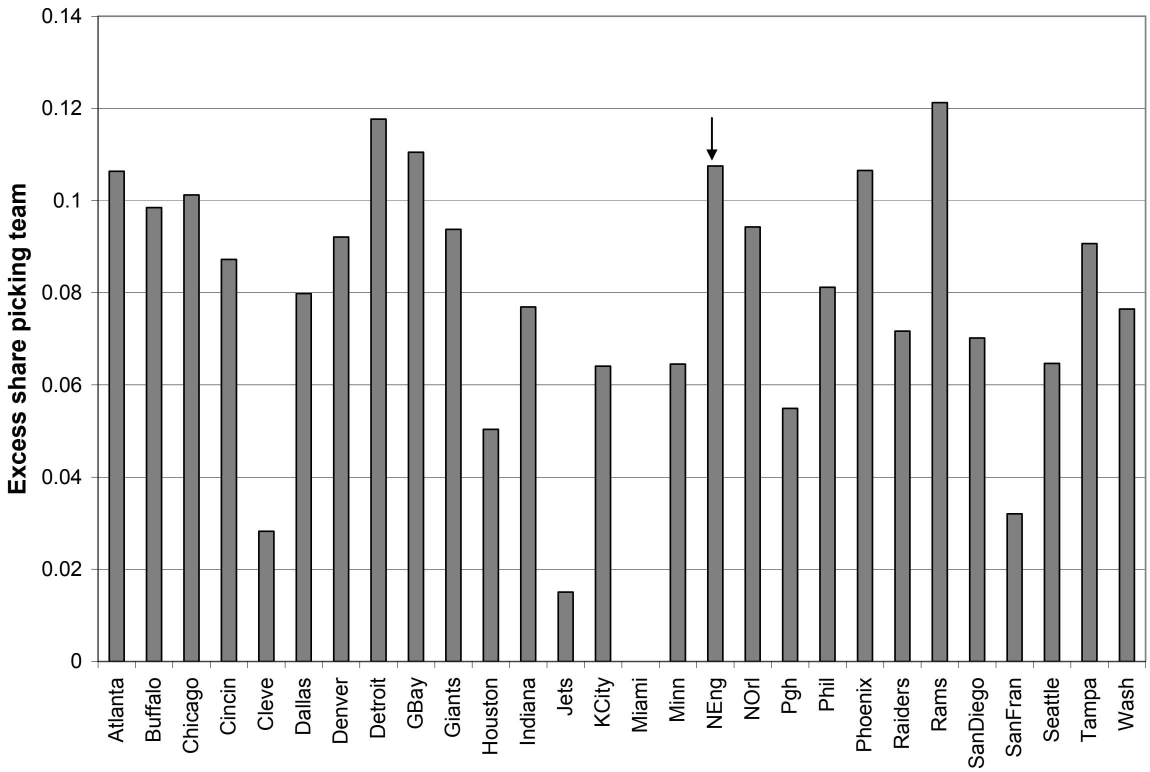

Notes: Data from Avery and Chevalier (1999). Dataset contains information on the picks of *Boston Globe* sports columnists for NFL games in the 1984-1994 seasons, as well as the outcome of the game and the opening betting line. The bar for team i represents the estimated coefficient $\hat{\delta}_i$ in a regression of the form

$$share_j = \alpha + \delta_i \left[(home_j = i) - (away_j = i)\right] + \gamma \left(line_j\right) + \varepsilon_j$$

where $share_j$ denotes the share of local columnists picking the home team to win game j, $home_j$ indexes the home team in game j, $away_j$ indexes the visiting team in game j, and $line_j$ is a vector of dummy variables representing deciles of the opening betting line.

Figure 6 *Competition and bias in local news coverage of the 2000 election*

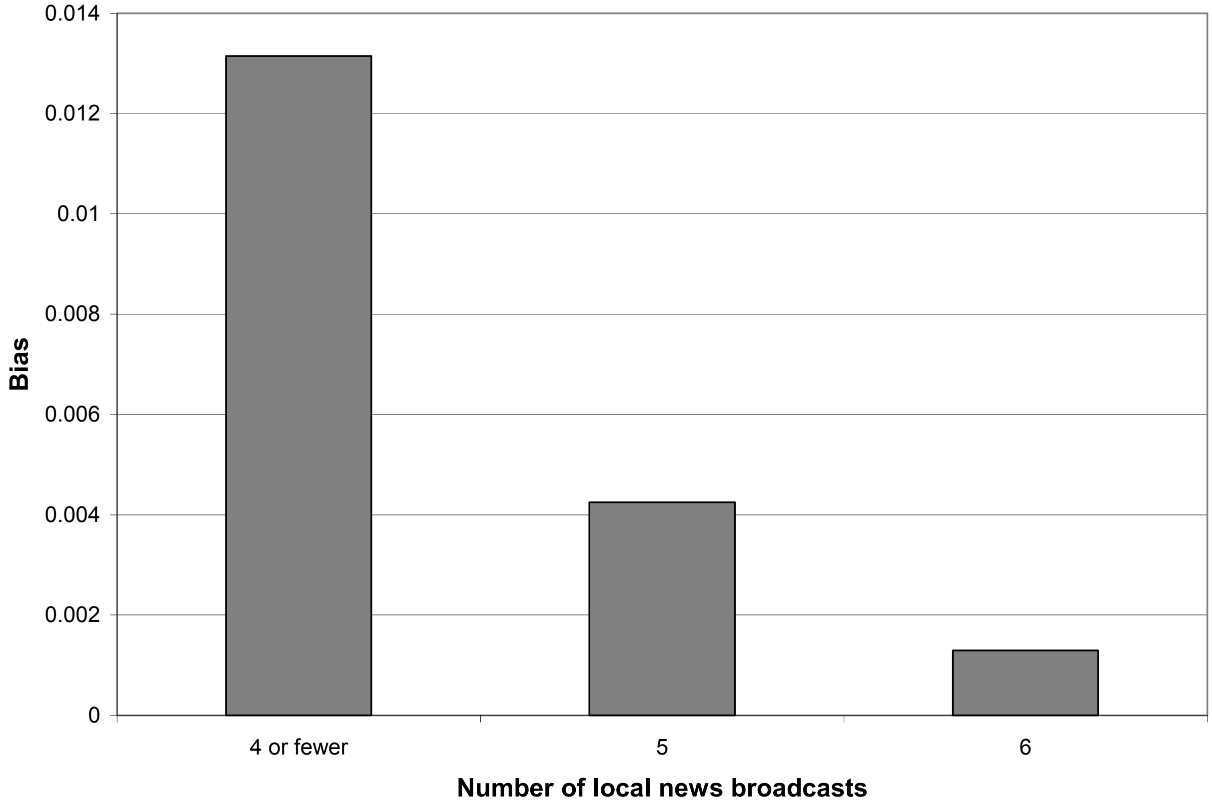

Notes: Bias is measured as the average across all stations in a market of

$$bias_i = \left(\frac{bush_i}{bush_i + gore_i} - \frac{1}{2} \right)^2$$

where x_i denotes the number of seconds given to candidate x by station i, with data taken from the Local News Archive (Kaplan and Hale, 2001). Number of local news broadcasts reflects the number of stations showing local news coverage at some point during the day as of July 2002, compiled from www.tvguide.com.

Table 1 *Competition and bias in local news coverage of the 2000 election*

	(1)	(2)	(3)	(4)
		$\left(\frac{bush_i}{bush_i+gore_i} - \frac{1}{2}\right)^2$		
Number of local news broadcasts, 2002	-0.0057 (0.0025)	-0.0064 (0.0025)	-0.0062 (0.0030)	-0.0062 (0.0031)
Census region controls?	NO	YES	YES	YES
log(population), 2000			-0.0006 (0.0043)	-0.0004 (0.0053)
log(income per capita), 1999				-0.0013 (0.0188)
N	58	58	58	58
R^2	0.0834	0.1747	0.1751	0.1752

Notes: Bias is measured as the average across all stations in a market of

$$bias_i = \left(\frac{bush_i}{bush_i + gore_i} - \frac{1}{2}\right)^2$$

where x_i denotes the number of seconds given to candidate x by station i, with data taken from the Local News Archive (Kaplan and Hale, 2001). Number of local news broadcasts reflects the number of stations showing local news coverage at some time during the day as of July 2002, compiled from www.tvguide.com. Data on population, income per capita, and Census region are taken from the U.S. Census, 2000.

A Appendix: Extensions and Generalizations

A.1 Consumer inferences about quality

In this extension, we show that the basic link between consumer priors and inferences about quality holds in a much larger class of information structures than the simple model considered in the paper. That is, it is a robust property of Bayesian belief formation.

Suppose the true state of the world is $\omega \in \{L, R\}$. Information sources, which may be high or low quality, report a signal $x \in X$. The density of a high-quality signal conditional on the state ω is $\bar{\pi}^{\omega}(x)$ and the density of a low-quality signal is $\pi^{\omega}(x)$. Here $\bar{\pi}$ and π may be either PMFs or PDFs.

We say that a value x *supports* R if $\bar{\pi}^{R}(x) > \bar{\pi}^{L}(x)$—i.e. if seeing x from a high-quality source provides information that R is the true state. We assume that the high-quality source is uniformly more informative than the low-quality source in the sense that:

$$
\begin{aligned}
\frac{\bar{\pi}^{R}(x)}{\bar{\pi}^{L}(x)} &> \frac{\pi^{R}(x)}{\pi^{L}(x)} \text{ if } x \text{ supports } R; \\
\frac{\bar{\pi}^{L}(x)}{\bar{\pi}^{R}(x)} &> \frac{\pi^{L}(x)}{\pi^{R}(x)} \text{ if } x \text{ supports } L.
\end{aligned}
\tag{7}
$$

This is the intuitive definition of informativeness in the case where both $\bar{\pi}^{R}(x) - \bar{\pi}^{L}(x)$ and $\pi^{R}(x) - \pi^{L}(x)$ have the same sign. There may also be situations where a given value of x provides evidence for R under the high-quality signal and evidence for L under the low-quality signal (or vice versa), in which case equation 7 holds trivially.

Suppose that a consumer has prior probability θ that the true state is R, and prior probability λ that the source is high quality. The following proposition characterizes how the signal x influences the consumer's posterior estimate of quality: $\hat{\lambda}(x, \theta)$. It shows that consumers who believe R is likely to be the true state will judge quality to be higher than those who believe L is likely to be the true state if and only if x supports their prior beliefs.[26]

Proposition 8 *$\hat{\lambda}(x, \theta)$ is strictly increasing in θ if x supports R and strictly decreasing in θ if x supports L.*

Proof. See Appendix B. ■

[26] Note that Proposition 8 only concerns the comparison of $\hat{\lambda}(x, \theta)$ across consumers with different prior beliefs. One could ask whether seemingly analogous statement holds about the comparison across different x—i.e. is $\hat{\lambda}(x, \theta)$ increasing in x (where x is ordered so that higher x provide more support for R) if and only if $\theta > .5$? The answer is clearly no, because we have imposed no restrictions on the way low and high quality signals weight different x. It may be, for example, that there is some value x^* that supports L but that is only reported by high quality sources. In this case, $\hat{\lambda}(x^*, \theta)$ could be greater than $\hat{\lambda}(x', \theta)$, even if x' supports R and θ is close to 1. To continue the previous example, it may be that a news source would be judged higher quality for making a moderate liberal claim (the death penalty does not deter crime) than an outlandish conservative claim (all Democratic senators are evil aliens from Mars), no matter how conservative the views of its customers.

A.2 More general signal space

The model presented in the body of the paper assumes that firms receive a binary signal of the state of the world. In this section, we show that this assumption can be relaxed without affecting any of the results in the paper. With a more general signal space, firms seeking to emulate the behavior of the high type will still have a temptation to lean towards the prior beliefs of their customers. And, as before, the presence of competition or ex-post feedback will tend to discipline this incentive and therefore to reduce the amount of equilibrium bias. In particular, we model firms as receiving signals drawn from a continuous distribution and then choosing a binary report, either L or R. This can be thought of as a continuous approximation to a game in which firms receive a finite set of binary signals, and must choose one to report.[27]

To see this formally, imagine now that if a normal firm j has a scoop in period t then it receives a signal $x_{jt} \in (-b, b)$ with $b \in (0, \infty]$ whose distribution function $F(\cdot)$ depends on the state of the world.[28] After observing this signal, the firm has the option of reporting either R or L. (We continue to assume that if a high-type firm has the scoop it always reports the true state.) We assume that $F(\cdot)$ has full support on $(-b, b)$, and that higher values of x indicate a greater likelihood that the true state is R. More precisely, we assume that

$$\frac{f(x \mid R)}{f(x \mid L)} \tag{R1}$$

is strictly increasing in x, where $f(\cdot)$ is the (continuous and differentiable) probability density function associated with $F(\cdot)$.

We also impose the following restrictions:

$$\begin{aligned}
\lim_{x \to -b} \frac{f(x \mid R)}{f(x \mid L)} &= 0 && \text{(R2)} \\
\lim_{x \to b} \frac{f(x \mid R)}{f(x \mid L)} &= \infty && \text{(R3)} \\
\frac{f(0 \mid R)}{f(0 \mid L)} &= 1 && \text{(R4)} \\
1 - F(0|R) &= F(0|L) > \theta && \text{(R5)}
\end{aligned}$$

Restrictions (R2) and (R3) imply that as the value of x approaches the boundaries, it is strong enough to overwhelm any non-doctrinaire prior. Restriction (R4) normalizes the signal space so that a signal of 0 provides no information about the true state. The first part of (R5) is a symmetry condition that requires that the probability of a positive signal if the true state is R is equal to the probability of a negative signal if the true state is L. The second part of (R5) puts a lower bound

[27] See also Suen (2004) for a model in which an information provider provides a binary report based on a continuous signal. Suen's model highlights the fact that if a provider wishes to maximize the value of its report to consumers, it will simply state the correct posterior based on its consumers' prior belief. Thus there is a sense in which optimal reporting involves a dependency between firms' reports and consumers' priors, although this arises from an intuition separate from the one we capture here.

[28] Here we use $b = \infty$ to denote the case in which $(-b, b) = \mathbb{R}$.

on the informativeness of the firm's signal by guaranteeing that consumers in either group would rather take action R when $x > 0$ and L when $x < 0$ than the action that is optimal given their priors. (This is analogous to our assumption that $\pi > \theta$ in the two-signal model.)

Given these conditions, we have the following characterization of equilibrium behavior:

Proposition 9 *There exists a number $k^* \in (-b, b)$ such that in any informative equilibrium all firms report R if and only if $x \geq k^*$. Moreover,*

$$\begin{aligned} k^* &\lesseqgtr 0 \Leftrightarrow \left(\theta \gtreqless \frac{1}{2}\right) \\ \frac{d\,|k^*|}{d\mu} &\leq 0 \\ \frac{d\,|k^*|}{dJ} &\leq 0 \end{aligned}$$

with these inequalities strict whenever $|k^*| > 0$.

Proof. *See Appendix B.* ■

Note that given the firm's prior belief of $\frac{1}{2}$, the most informative possible strategy is for the firm to report R if and only if $x \geq 0$. The proposition therefore shows that in equilibrium firms are too likely to report R when $\theta > \frac{1}{2}$ relative to an efficient benchmark. Additionally, it confirms that bias is decreasing in the probability of feedback μ and decreasing in the number of firms J, all of which replicate the findings in the two-signal model.

A.3 Allowing for a dishonest high type

In the model presented in the body of the paper, we assume that high-type firms both know the true state of the world *and* always report their signals honestly in the reporting stage. In this subsection, we relax the latter assumption and permit the high type to choose its reporting-stage action so as to maximize profits. (We continue to assume for simplicity that the high type reports honestly in the feedback stage.) While there are multiple equilibria in this case, we show that the strategy profile studied in the body of the paper is unique with respect to an intuitive stability criterion.

It is easy to verify that there exists an equilibrium in which high-type firms report honestly and normal-type firms play the strategy defined as ρ^* in Proposition 1. That is, the equilibrium studied in the body of the paper survives when we permit high-type firms to choose their actions optimally.[29] However, other equilibria are also possible. Whenever consumers believe that in

[29] To see why, suppose that the high type is honest and the normal type plays as in Proposition 1. Given that consumers believe that firms play this strategy, both types wish to appear to be the high type and thus to match the feedback. Since high-type firms know that the feedback will always reveal their own signal, they will always receive a strictly greater expected return to honesty than normal-type firms. But if normal-type firms play $\rho^* > 0$, then we know that they are exactly indifferent between playing R and playing L given that they see a signal of L, and strictly prefer to report R given that they see a signal of R. As a consequence, high-type firms will strictly prefer to report honestly, and the equilibrium is sustained. If $\rho^* > 0$ normal firms strictly prefer to report honestly and so high-type firms do as well

equilibrium high types are more informative than normal types in period $t+1$, normal types will try to emulate high types in period t. If the high type is not being perfectly honest in period t, in general the strategy ρ^* defined in Proposition 1 will not be an equilibrium in that period; rather the normal type's equilibrium play will involve additional bias in the direction of matching the high type's behavior.

Such equilibria are unstable in an intuitive sense, however. In any equilibrium in which the high type reports dishonestly with positive probability, a small perturbation to the high type's behavior would lead the proposed equilibrium to "unravel." To see why, consider that if the high type sometimes reports R when its signal is L, then high-type firms must be indifferent between reporting R and reporting L given consumers' correct beliefs about the firms' strategies.[30] But then, in any equilibrium in which firms strictly prefer to be thought of as high-type, a small increase in the probability of the high type reporting R will increase the incentives for the high type to do so. This in turn will lead high-type firms to move towards reporting R more frequently, and so on until the process reaches a boundary.

By contrast, the equilibrium characterized by Proposition 1 is stable in the sense that high-type firms strictly prefer to play their equilibrium strategies, and when normal firms become more likely to report R, this reduces the incentive for them to say R, so that behavior has a tendency to return to the equilibrium point.

Formally, let $q \in \{0,1\}$ index whether a firm is high-type (with $q=1$ denoting a high-type firm), and let $\alpha_d^q \in [0,1]$ be the probability that type q reports R given a signal of d. Analogously, let $\Delta_d^q(\alpha)$ be type q's net return to reporting R given a signal of d and equilibrium play $\alpha = \left\{\alpha_d^q\right\}_{d\in\{L,R\}}^{q\in\{0,1\}}$.[31] We will say that an equilibrium is *stable* if for all q and d, either $\left|\Delta_d^q(\alpha)\right| > 0$ or $\Delta_d^q(\alpha) = 0$ and $\frac{\partial \Delta_d^q(\alpha)}{\partial \alpha_d^q} < 0$. That is, an equilibrium is stable if for each signal d and type q, either the type strictly prefers playing R or L given a signal of d, or it is indifferent between reports and an increase in its probability of its playing R *decreases* its return to doing so—i.e. the function $\Delta_d^q(\alpha)$ crosses the x-axis from above. This definition captures an idea that when a type's behavior is perturbed, it ought to have an incentive to move back to the equilibrium point.

We also restrict attention to equilibria in which the high type is indeed perceived to be higher quality, conditional on its strategy. This is an intuitive requirement for the reputational incentives we have introduced to be meaningful. In cases where it fails to hold, the equilibria effectively reverse the role of the high-quality and normal firms. Let $V(q,\alpha)$ represent consumers' information value from a firm of type q playing strategy α^q. We now re-define an *informative equilibrium* to be an equilibrium which satisfies two conditions: (i) $V(0,\alpha) > 0$ in every period as before; and (ii) $V(1,\alpha) > V(0,\alpha)$ in every period. This allows us to state the following result.

Proposition 10 *There exists a unique stable informative equilibrium α^*. In this equilibrium, high-quality firms report honestly, and normal firms play the equilibrium strategy defined in Proposition 1.*

[30] Note that our definition of informativeness will rule out equilibria in which the high type plays a pure "lying" strategy—i.e. always reports L when its signal is R and vice-versa. If such equilibria were allowed, they would have all the properties of the equilibria we discuss except that the signals would be re-labeled.

[31] Since our definition of stability will apply separately to each firm and time period, we suppress the subscripts j and t in defining equilibrium strategies and incentives.

Proof. *See Appendix B.* ■

B Appendix: Proofs

In this section, we prove the lemmas and propositions that were stated without proof in the text. As a preliminary step, we will introduce some additional notation for consumers' posteriors on the firm's quality. Define the probability that a normal firm reports R in some stage conditional on the true state being R to be π_R. Similarly, define the probability that the firm reports L conditional on this being the state by π_L. We can then define a firm's strategy s to be the pair $\{\pi_R, \pi_L\}$.

We will write the posteriors of a consumer that sees this one report and no additional information about the true state as:

$$\hat{\lambda}^0(R; s, \theta) = \frac{\theta\lambda}{\theta\lambda + [\pi_R\theta + (1-\pi_L)(1-\theta)](1-\lambda)} \tag{8}$$

$$\hat{\lambda}^0(L; s, \theta) = \frac{(1-\theta)\lambda}{(1-\theta)\lambda + [\pi_L(1-\theta) + (1-\pi_R)\theta](1-\lambda)}.$$

The posteriors of a consumer who learns for certain that the true state is R are:

$$\hat{\lambda}^R(R; s) = \frac{\lambda}{\lambda + \pi_R(1-\lambda)}$$

$$\hat{\lambda}^R(L; s) = 0.$$

And the posteriors of a consumer who learns for certain that the true state is L are:

$$\hat{\lambda}^L(L; s) = \frac{\lambda}{\lambda + \pi_L(1-\lambda)}.$$

$$\hat{\lambda}^L(R; s) = 0.$$

Note that in the special case when strategies can be represented by a single number ρ as defined in Equation 6, we have:

$$\begin{aligned} \pi_R(\rho) &= \pi + (1-\pi)\rho \\ \pi_L(\rho) &= \pi(1-\rho). \end{aligned} \tag{9}$$

when $\rho > 0$ and:

$$\begin{aligned} \pi_R(\rho) &= \pi(1+\rho) \\ \pi_L(\rho) &= \pi - (1-\pi)\rho. \end{aligned}$$

when $\rho < 0$.

When it is understood that strategies take this form, we will abuse notation slightly and write expressions directly as a function of ρ. For example:

$$\hat{\lambda}^0(n;\rho,\theta) = \hat{\lambda}^0(n;s,\theta)$$
$$\hat{\lambda}^R(n;\rho) = \hat{\lambda}^0(n;s,\theta)$$

where s is the pair $\{\pi_R(\rho), \pi_L(\rho)\}$.

PROOF OF LEMMA 1

Observe that because attention is restricted to informative Bayesian equilibria, expected future profits at the end of any period will be linearly increasing in young consumers' posterior belief $\hat{\lambda}$.

Consider first the decision of a firm j that has the scoop in a given period. All consumers that read the paper in the feedback stage also read in the reporting stage. If $\tilde{n}_j \neq n_j$, their posterior will be $\hat{\lambda} = 0$ since high-quality firms always report the same thing in both periods.[32] If $\tilde{n}_j = n_j$, their posterior will depend on the feedback they receive (it will be $\hat{\lambda} = 0$ if the consumer learns for certain that $n_j \neq \omega$ and $\hat{\lambda} > 0$ otherwise), but will be greater than zero with positive probability.[33] We thus have for all i and ω:

$$E_{\hat{\theta}_{i-j}|\omega}\hat{\lambda}\left(n, n; \mathbf{s}, \hat{\theta}_{i-j}\right) > E_{\hat{\theta}_{i-j}|\omega}\hat{\lambda}\left(n, \tilde{n}; \mathbf{s}, \hat{\theta}_{i-j}\right) \text{ if } \tilde{n} \neq n.$$

Since future profits are linearly increasing in $\hat{\lambda}\left(n, \tilde{n}; \mathbf{s}, \hat{\theta}_{i-j}\right)$, the firm prefers to report $\tilde{n}_j = n_j$.

Next, consider the decision of some firm j that did not have the scoop in period t. Suppose for a moment that consumers expect a firm in this position to always report truthfully. Then consumers who do not see exogenous feedback about the true state will not change their estimate of the firm's quality based on its report. Consumers who do see exogenous feedback will increase their posterior on quality to $\hat{\lambda} > \lambda$ if $\tilde{n}_j = \omega$ and have $\hat{\lambda} = 0$ if $\tilde{n}_j \neq \omega$. This means that if s^+ is informative:

$$E_{\hat{\theta}_{i-j}|\omega}\hat{\lambda}\left(0, \omega; \mathbf{s}, \hat{\theta}_{i-j}\right) > E_{\hat{\theta}_{i-j}|\omega}\hat{\lambda}\left(0, \tilde{n}; \mathbf{s}, \hat{\theta}_{i-j}\right) \text{ if } \tilde{n} \neq \omega,$$

where we let $n = 0$ indicate that a firm did not have the scoop in the first period. So truthful reporting is indeed an equilibrium.

We show now that this is the unique equilibrium. Let χ be an indicator for the event that the consumer receives exogenous feedback about the true state. The expected future profit that j receives from a young consumer i, given that it did not have a scoop in the reporting stage (i.e. $n = 0$) and it reports $\tilde{n}$ in the feedback stage, is:

$$\mu\hat{\lambda}^\omega(\tilde{n};\pi_R,\pi_L) + (1-\mu)E_{\hat{\theta}_{i-j}|\chi=0,\omega}\hat{\lambda}^0\left(\tilde{n};\pi_R,\pi_L,\hat{\theta}_{i-j}\right)$$

[32] There is a minor technical issue in that conditional on normal firms reporting $\tilde{n}_j = n_j$ with probability 1, the observation $\tilde{n}_j \neq n_j$ is a zero probability event and consumer beliefs in this case can be freely specified in a Bayesian equilibrium. If they held beliefs such that the firm would actually prefer to deviate and report $\tilde{n}_j \neq n_j$ (i.e. believing in this event that firms are high-quality for sure), the model would have *no* equilibrium. In any equilibrium, therefore, consumers must have beliefs that make reporting $\tilde{n}_j = n_j$ a dominant strategy.

[33] By assumption, there is positive probability that any given consumer sees neither exogenous feedback nor the report of any firm other than j. In this case $\hat{\theta}_{i-j} = \theta$ and $\hat{\lambda}\left(n, \tilde{n}; s, \hat{\theta}_{i-j}\right) > 0$.

where π_R and π_L are the expected distribution of $\tilde{n}$ given the firm's strategy. (The expectation in this expression is over consumer beliefs $\hat{\theta}_{i-j}$, conditional on the consumer receiving no exogenous feedback and the true state being ω).

Suppose that $\pi_L < 1$, so that the firm sometimes reports R when the true state is L. Because $\hat{\lambda}^L(L;s) > \hat{\lambda}^L(R;s)$ and $\mu > 0$, it must then be the case that

$$E_{\hat{\theta}_{i-j}|\chi=0,\omega=L}\hat{\lambda}^0\left(R;s,\hat{\theta}_{i-j}\right) > E_{\hat{\theta}_{i-j}|\chi=0,\omega=L}\hat{\lambda}^0\left(L;s,\hat{\theta}_{i-j}\right). \tag{10}$$

If the true state is L, there cannot be information in equilibrium such that $\hat{\theta}_{i-j} = 1$. Furthermore Equation 8 implies that $\hat{\lambda}^0(R;s,0) < \hat{\lambda}^0(L;s,0)$. We must therefore have $\hat{\lambda}^0\left(R;s,\hat{\theta}_{i-j}\right) > \hat{\lambda}^0\left(L;s,\hat{\theta}_{i-j}\right)$ for at least some $\hat{\theta}_{i-j}$ strictly between zero and one. From Equation 8 this requires:

$$\pi_L + (1-\pi_R)\frac{\theta}{(1-\theta)} > \pi_R + (1-\pi_L)\frac{(1-\theta)}{\theta}.$$

Observe that the fact that $\hat{\lambda}^R(R;s) > \hat{\lambda}^R(L;s)$ and Equation 10 together mean we must have $\pi_R = 1$ (the firm always reports R when the state is R). But then we would need

$$\pi_L > 1 + (1-\pi_L)\frac{(1-\theta)}{\theta}.$$

which cannot be true since $\pi_L < 1$.

A similar contradiction obtains if we begin by assuming the firm sometimes reports L when the true state is R. ■

PROOF OF PROPOSITION 1

Because feedback-stage strategies will be as in Lemma 1, the firm will choose its reporting-stage strategies to maximize Equation 5. Using the fact that $\hat{\lambda}^R(L;s) = \hat{\lambda}^L(R;s) = 0$, define the difference between the value of this expression when $n = R$ and $n = L$ to be:

$$\begin{aligned}\Delta(d;s,\theta) &= \tilde{\mu}\left[\hat{\theta}(d)\hat{\lambda}^R(R;s) - \left(1-\hat{\theta}(d)\right)\hat{\lambda}^L(L;s)\right] \\ &\quad + (1-\tilde{\mu})\left[\hat{\lambda}^0(R;s,\theta) - \hat{\lambda}^0(L;s,\theta)\right].\end{aligned} \tag{11}$$

The firm will report R after seeing information d only if $\Delta(d;s,\theta) \geq 0$.

We first show that if a firm distorts L signals with positive probability in the reporting stage (sometimes reports R after seeing L), it will never distort R signals, and vice versa. This allows us to characterize strategies in terms of ρ as defined in Equation 6.

Suppose the firm sometimes reports R after seeing $d = L$. This implies $\Delta(L;s,\theta) \geq 0$ (where $\Delta()$ is defined as in Equation 11). The fact that $\hat{\theta}(R)$ is strictly greater than $\hat{\theta}(L)$ means that $\Delta(R;s,\theta) > 0$, so the firm never reports L after seeing R. Analogous reasoning shows that if the firm sometimes reports L after seeing R, it can never report R after seeing L.

Next, we show that $\rho^* = 1$ or $\rho^* = -1$ could never be an equilibrium. If $\rho = 1$, normal firms always report R, and so a consumer seeing $n = L$ knows that the firm is high quality while a consumer seeing $n = R$ will believe the firm to be normal with positive probability. This means all firms strictly prefer to report L which contradicts the assumption that $\rho = 1$. An analogous contradiction obtains if $\rho = -1$.

This implies that in any informative equilibrium of the stage game, each firm's reporting-stage strategy ρ_j^* must satisfy the following conditions:

$$\begin{array}{lll} \text{If } \rho_j^* > 0: & \Delta\left(L;\rho_j^*,\theta\right) = 0 \text{ and } \Delta\left(R;\rho_j^*,\theta\right) \geq 0; \\ \text{If } \rho_j^* = 0: & \Delta\left(L;\rho_j^*,\theta\right) \leq 0 \text{ and } \Delta\left(R;\rho_j^*,\theta\right) \geq 0; \\ \text{If } \rho_j^* < 0: & \Delta\left(L;\rho_j^*,\theta\right) \leq 0 \text{ and } \Delta\left(R;\rho_j^*,\theta\right) = 0. \end{array} \tag{12}$$

(Where again we abuse notation and write $\Delta(d;s,\theta)$ directly as a function of ρ.) Note that the equalities in the first and third lines follow from the fact that ρ_j^* cannot be either 1 or -1 and so if $\rho_j^* \neq 0$, the firm must be playing a mixed strategy.

We now show that the equilibrium ρ^* exists and is unique, and that $\rho^* \geq 0$.

Observe, first, that by Equations 9 and 8, $\hat{\lambda}^R(R;\rho)$ and $\hat{\lambda}^0(R;\rho,\theta)$ are strictly decreasing in ρ and $\hat{\lambda}^L(L;\rho)$ and $\hat{\lambda}^0(L;\rho,\theta)$ are strictly increasing in ρ for all $\theta \in (.5,1)$. Therefore $\Delta(d;\rho,\theta)$ is strictly decreasing in ρ for all d and $\theta \in (.5,1)$.

Second, because $\hat{\theta}(R) > \hat{\theta}(L)$, $\Delta(R;\rho,\theta) > \Delta(L;\rho,\theta)$ for all ρ and $\theta \in (.5,1)$.

Third, by Equations 9 and 8, it is clear that $\hat{\lambda}^L(L;0) = \hat{\lambda}^R(R;0)$. Thus, the fact that $\hat{\theta}(R) > .5$ implies:

$$\hat{\theta}(R)\,\hat{\lambda}^R(R;0) - \left(1 - \hat{\theta}(R)\right)\hat{\lambda}^L(L;0) > 0.$$

Inspection of Equation 8 also shows that $\hat{\lambda}^0(R;0,\theta) > \hat{\lambda}^0(L;0,\theta)$ for $\theta > .5$. Combining these facts gives us $\Delta(R;0,\theta) > 0$ for $\theta \in (.5,1)$.

Finally, again from Equation 8, we know that as $\rho \to 1$: $\hat{\lambda}^0(R;\rho,\theta) < \hat{\lambda}^0(L;\rho,\theta)$ and $\hat{\lambda}^R(R;\rho) < \hat{\lambda}^L(L;\rho)$. Since $\hat{\theta}(L) < .5$, this implies $\lim_{\rho\to1}\Delta(L;\rho,\theta) < 0$.

Define ρ_L to be the unique value of ρ such that $\Delta(L;\rho,\theta) = 0$. Suppose that $\rho_L > 0$. Then $\rho > 0$ whenever $\Delta(L;\rho,\theta) \leq 0$, meaning that $\rho^* = \rho_L \in (0,1)$ is the unique strategy that satisfies the conditions of Equation 12. Suppose, on the other hand, that $\rho_L \leq 0$. Then $\rho^* = 0$ clearly satisfies these conditions, and no $\rho^* \neq 0$ could satisfy them because this would require either $\rho^* > 0$ and $\Delta(L;\rho^*,\theta) = 0$ (which is impossible because $\rho_L \leq 0$ is the unique value such that $\Delta(L;\rho,\theta) = 0$) or $\rho^* < 0$ and $\Delta(R;\rho^*,\theta) = 0$ (which is impossible because $\Delta(R;\rho,\theta)$ is strictly decreasing in ρ and $\Delta(R;0,\theta) > 0$). This shows that ρ^* exists, is unique, and is greater than or equal to zero.

The final step in completing the proof is to show that ρ^* is informative. If a consumer followed the report of firm it knew to be normal and reporting with bias $\rho \geq 0$, her expected gain would be:

$$(\pi + (1-\pi)\rho)\,\theta + \pi(1-\rho)(1-\theta) - \theta$$

which is weakly positive by the maintained assumption that $\pi > \theta$,. and strictly positive if $\rho < 1$.■

PROOF OF PROPOSITION 2

Note that ρ^* is continuous in λ and suppose that

$$\lim_{\lambda \to 0} \rho^* > 0$$

so that $\rho^* > 0$ for small λ. Begin with consumers. If a consumer expects the firm to report with bias ρ, her surplus from the firm's information in the period when she is young is:

$$V(\lambda;\rho) = \lambda + (1-\lambda)\left[(\pi + (1-\pi)\rho)\theta + \pi(1-\rho)(1-\theta)\right] - \theta$$

which is strictly decreasing in ρ. (The idiosyncratic portion of consumer utility, δ_{ij}, is unaffected by bias.)

To see that surplus in the period when she is old cannot be lower, note that if consumers are more likely to follow the report of the firm when there is no bias then they must be better off. But as λ gets small, consumers' posteriors on quality depending on what they observe must all approach 0, so that the effect on the value of information dominates any effects on the revelation of information about firm type. Therefore old consumers will also be made better off.

Next consider firms' welfare. We showed above that the information value of the firm's signal $V(\lambda;\rho)$ is strictly higher when bias is eliminated and that the difference is decreasing in λ. This means profits from young consumers are strictly higher without bias and the gain is bounded away from 0 as λ becomes small.

The lowest profits from old consumers could be when bias is eliminated is $V(0;0)$—this would be the case if the consumer would learn the firm was normal with certainty. Let $\hat{\lambda}^*$ denote the highest possible belief among old consumers that the firm is high quality given that the firm is playing ρ^*. Note that $\hat{\lambda}^* < 1$ since there is no outcome that with certainty identifies the firm as high-type. The gain in profits from old consumers to eliminating bias is thus bounded below $V(0;0) - V\left(\hat{\lambda}^*;\rho^*\right)$ which becomes strictly positive as $\lambda \to 0$ since $\hat{\lambda}^*$ goes to 0 as λ goes to 0.

For small λ, therefore, when $\rho^* > 0$ the gain from providing superior information must outweigh the loss from the quality assessments of older consumers, making overall profits higher when bias is eliminated.■

PROOF OF PROPOSITION 3

Define $\Delta(d;\rho,\theta)$ as in Equation 11 in the proof of Proposition 1.

Define ρ_L as before to be:

$$\rho_L(\theta) = \{\rho : \Delta(L;\rho,\theta) = 0\}.$$

In the proof of Proposition 1 we showed that $\Delta(L;\rho,\theta)$ is decreasing in ρ.

The only terms in $\Delta(L;\rho,\theta)$ that depend on θ are $\hat{\lambda}^0(R;\rho,\theta) - \hat{\lambda}^0(L;\rho,\theta)$ and $\hat{\theta}(L)$. By Proposition 8, $\hat{\lambda}^0(R;\rho,\theta)$ is increasing in θ and $\hat{\lambda}^0(L;\rho,\theta)$ is decreasing in θ. Also, $\Delta(L;\rho,\theta)$ is clearly increasing in $\hat{\theta}(L)$ which in turn is increasing in θ. So $\Delta(L;\rho,\theta)$ is increasing in θ.

Thus, $\rho_d(\theta)$ is increasing in θ for all d. Because $\Gamma()$ is continuous in both θ and ρ, $\rho_d(\theta)$ is also continuous.

Suppose that $\theta \approx .5$ and $\rho = 0$. Because by inspection of Equation 8 $\hat{\lambda}^0(L;0,.5) = \hat{\lambda}^0(R;0,.5)$

and $\hat{\lambda}^{L}(L;0) = \hat{\lambda}^{R}(R;0) > 0$, and because $\hat{\theta}(L) < .5$, we know that $\Delta(L;0,.5) < 0$. This shows $\lim_{\theta \to .5} \rho_L(\theta) < 0$.

Inspection of Equation 8 also shows that $\lim_{\theta \to \pi} \left[\hat{\lambda}^{0}(R;0,\theta) - \hat{\lambda}^{0}(L;0,\theta) \right] > 0$. At $\theta = \pi$, $\hat{\theta}(L) = .5$. Therefore, $\Delta(L;0,\pi) > 0$. This shows $\rho_L(\pi) > 0$.

These properties together imply there is a unique $\theta^* \in (.5,\pi)$ such that $\rho_L(\theta^*) = 0$. From the proof of Proposition 1 we know that $\rho^* = 0$ if $\theta \leq \theta^*$ and $\rho^* = \rho_L(\theta)$ which is strictly increasing in θ if $\theta > \theta^*$. Continuity follows from the continuity of $\rho_L(\theta)$. ■

PROOF OF PROPOSITION 4

Let $\Delta(L;\rho,\theta,\tilde{\mu})$ be defined as in Equation 11 but with the dependence on $\tilde{\mu}$ written explicitly. Define $\rho_L(\theta;\tilde{\mu})$ to be:

$$\rho_L(\theta;\tilde{\mu}) = \{\rho : \Delta(L;\rho,\theta,\tilde{\mu}) = 0\}.$$

>From the construction of the equilibrium in the proof of Proposition 1, it will suffice to show that $\rho_L(\theta;\tilde{\mu})$ is strictly decreasing in $\tilde{\mu}$ wherever $\rho_L(\theta;\tilde{\mu}) \geq 0$—this will imply both that ρ^* is strictly decreasing in $\tilde{\mu}$ in the range where $\rho^* = \rho_L(\theta;\tilde{\mu}) > 0$, and that θ^* (the point where $\rho_L(\theta;\tilde{\mu}) = 0$) is strictly increasing in $\tilde{\mu}$.

To show this, suppose ρ^* is the equilibrium bias given some θ and $\tilde{\mu}$: $\rho_L(\theta;\tilde{\mu})$. Holding ρ^* fixed, we establish that $\Delta(L;\rho^*,\theta,\tilde{\mu})$ is strictly decreasing in $\tilde{\mu}$.

Consider the term multiplied by $\tilde{\mu}$ in Equation 11 evaluated at ρ^*:

$$\hat{\theta}(L)\,\hat{\lambda}^{R}(R;\rho^*) - \left(1 - \hat{\theta}(L)\right)\hat{\lambda}^{L}(L;\rho^*). \tag{13}$$

By Equation 8, $\rho \geq 0$ implies $\hat{\lambda}^{L}(L;\rho) \geq \hat{\lambda}^{R}(R;\rho) > 0$. Since $\hat{\theta}(L) = (1-\pi) < .5$, this term is strictly negative.

Since $\Delta(L;\rho_L(\theta;\tilde{\mu}),\theta,\tilde{\mu}) = 0$ by definition of $\rho_L(\theta;\tilde{\mu})$, we know that the term multiplied by $(1-\tilde{\mu})$,

$$\hat{\lambda}^{0}(R;\rho,\theta) - \hat{\lambda}^{0}(L;\rho,\theta),$$

must be strictly positive.

This means that $\Delta(L;\rho^*,\theta,\tilde{\mu})$ is strictly decreasing in $\tilde{\mu}$.

To show the final part of the proposition, we need to verify that for any $\theta \in (.5,1)$ there exists $\tilde{\mu}^* < 1$ such that $\rho_L(\theta;\mu) < 0$ for all $\tilde{\mu} \geq \tilde{\mu}^*$. Observe that since $\hat{\lambda}^{R}(R;0) = \hat{\lambda}^{L}(L;0) > 0$, Equation 13 is strictly negative evaluated at $\rho^* = 0$. This implies that we can find $\tilde{\mu}^* < 1$ such that $\Delta(L;0,\theta,\tilde{\mu}^*) < 0$. This implies $\rho_L(\theta;\mu) < 0$ for all $\tilde{\mu} \geq \tilde{\mu}^*$. ■

PROOF OF PROPOSITION 5

Suppose there are J papers and assume without loss of generality that paper 1 has the scoop in period t. Suppose all $j > 1$ report $\tilde{n}_j = n_1$ and consider a deviation in which at least one paper j reports $\tilde{n}_j \neq n_1$.

Note that when consumers expect the firms to always repeat the report of firm 1, these firms' reports have no independent effect on the consumers' beliefs. Their posteriors will thus be the same as in a monopoly market where there is only exogenous feedback. Firm 1's relative gain from reporting R rather than L after seeing a signal L (defined above in Equation 11) will now be:

$$\begin{aligned}\Delta(L;\rho,\theta) &= \mu\left[\hat{\theta}(L)\,\hat{\lambda}^{R}(R;\rho)-\left(1-\hat{\theta}(L)\right)\hat{\lambda}^{L}(L;\rho)\right]\\ &\quad+(1-\mu)\left[\hat{\lambda}^{0}(R;\rho,\theta)-\hat{\lambda}^{0}(L;\rho,\theta)\right].\end{aligned}$$

As μ becomes small, the difference $\hat{\lambda}^{0}(R;\rho^{*},\theta)-\hat{\lambda}^{0}(L;\rho^{*},\theta)$ approaches zero and both $\hat{\lambda}^{0}(R;\rho^{*},\theta)$ and $\hat{\lambda}^{0}(L;\rho^{*},\theta)$ become arbitrarily close to λ.

Any consumer who read a report from a paper other than j (that is they saw either n_1 or some $\tilde{n}_k$ for $k\neq j$) will place zero probability on the event $\tilde{n}_j\neq n_1$. We assume that their beliefs in this case are such that $\hat{\lambda}=0$. If j does not deviate, consumers will have $\hat{\lambda}\approx\lambda$. Profits from these consumers are thus strictly higher if the firm does not deviate, and the expectation of $\hat{\lambda}$ approaches λ as μ approaches 0.

The remaining consumers are those who only see the report of firm j. The only difference between the problem firm j faces with these consumers and that faced by firm 1 in the reporting stage is that j now knows the true state. It will therefore be indifferent between its reports if $\Delta(\omega;\rho^{*},\theta)=0$ where $\hat{\theta}(\omega)=1$ if $\omega=R$ and 0 otherwise. Since μ is close to zero and $\hat{\lambda}^{0}(R;\rho^{*},\theta)-\hat{\lambda}^{0}(L;\rho^{*},\theta)\approx 0$, $\Delta(\omega;\rho^{*},\theta)\approx 0$ and the increased profits the firm would get from these consumers if it deviates are arbitrarily small.

The firm therefore strictly prefers to follow its strategy, and the strategy $\tilde{n}_j=n_1$ for all j is indeed an equilibrium. ■

PROOF OF PROPOSITION 6

We will show that under the conditions of the proposition, it is an equilibrium for any firm j to play $s_j^t=\rho^{*}(\theta)$ for all t and that only type-R consumers will have positive information value from j's report. The fact that $s_j^t=-\rho^{*}(\theta)$ is also an equilibrium follows by symmetry.

Let $V\left(\rho,x,\hat{\lambda}\right)$ denote the expected value to a consumer with prior x and belief $\hat{\lambda}$ about the firm's quality of following the firm's report. Next define

$$\begin{aligned}D(\rho,x) &= V\left(\rho,x,\hat{\lambda}^{0}(R;x,\rho)\right)-V\left(\rho,x,\hat{\lambda}^{0}(L;x,\rho)\right)\\ D^{R}(\rho,x) &= V\left(\rho,x,\hat{\lambda}^{R}(R;x,\rho)\right)-V(\rho,x,0)\\ D^{L}(\rho,x) &= V(\rho,x,0)-V\left(\rho,x,\hat{\lambda}^{L}(L;x,\rho)\right)\end{aligned}$$

to be the net effect of reporting R on the expected valuations of old consumers with prior x conditional on no feedback, a feedback of R, and a feedback of L, respectively. Next, using Equation 3, let

$$\sigma(\rho,\theta)=\frac{.5+V(\rho,\theta,\lambda)}{1+V(\rho,\theta,\lambda)+V(\rho,1-\theta,\lambda)}$$

denote the fraction of a newspaper's young readers who have prior θ. Then we can write the net expected return to reporting R given a signal of L, denoted $\Delta(L;\rho,\theta,\tilde{\mu})$, as

$$\begin{aligned}\Delta(L;\rho,\theta,\tilde{\mu}) &= \sigma(\rho,\theta)\left[(1-\tilde{\mu})D(\rho,\theta)+\tilde{\mu}\pi D^{L}(\rho,\theta)+\tilde{\mu}(1-\pi)D^{R}(\rho,\theta)\right] \\ &\quad +(1-\sigma(\rho,\theta))\left[(1-\tilde{\mu})D(\rho,1-\theta)+\tilde{\mu}\pi D^{L}(\rho,1-\theta)+\tilde{\mu}(1-\pi)D^{R}(\rho,1-\theta)\right]\end{aligned}$$

Pick some $(\theta^{*},\tilde{\mu}^{*})$ such that

$$(1-\theta)(\pi+(1-\pi)\rho_{0}(\theta^{*},\tilde{\mu}^{*}))+\theta(\pi(1-\rho_{0}(\theta^{*},\tilde{\mu}^{*})))<\theta$$

and pick some $(\theta,\tilde{\mu})$ such that $\theta\geq\theta^{*}$ and $\tilde{\mu}\leq\tilde{\mu}^{*}$. It is immediate to verify from the definition of $\rho_{0}(\cdot)$ that $\rho_{0}(\theta,\tilde{\mu})\geq\rho_{0}(\theta^{*},\tilde{\mu}^{*})>0$, so that for small enough λ we must have that $\rho^{*}(\theta,\tilde{\mu})>0$ and hence that it is sufficient to show that

$$(1-\tilde{\mu})D(\rho^{*}(\theta,\tilde{\mu}),1-\theta)+\tilde{\mu}\pi D^{L}(\rho^{*}(\theta,\tilde{\mu}),1-\theta)+\tilde{\mu}(1-\pi)D^{R}(\rho^{*}(\theta,\tilde{\mu}),1-\theta)$$

is equal to 0 for small enough λ. Observe that

$$(1-\theta)(\pi+(1-\pi)\rho_{0}(\theta,\tilde{\mu}))+\theta(\pi(1-\rho_{0}(\theta,\tilde{\mu})))<\theta$$

Moreover, note that old consumers with prior $(1-\theta)$ have expected valuation (after seeing feedback $z\in\{0,L,R\}$ and report $n\in\{L,R\}$) given by

$$V\left(\rho^{*}(\theta,\tilde{\mu}),1-\theta,\hat{\lambda}_{n}^{\tilde{n}}\right)=\max\left\{(\left(1-\hat{\lambda}_{n}^{\tilde{n}}\right)((1-\theta)(\pi+(1-\pi)\rho^{*}(\theta,\tilde{\mu}))+\theta(\pi(1-\rho^{*}(\theta,\tilde{\mu}))))+\hat{\lambda}_{n}^{\tilde{n}}-\theta,0\right\}.$$

But note that for any z and n we have

$$\begin{aligned}&\lim_{\lambda\to0}\left[\left(1-\hat{\lambda}_{n}^{\tilde{n}}\right)((1-\theta)(\pi+(1-\pi)\rho^{*}(\theta,\tilde{\mu}))+\theta(\pi(1-\rho^{*}(\theta,\tilde{\mu}))))+\hat{\lambda}_{n}^{\tilde{n}}-\theta\right] \\ =\ &(1-\theta)(\pi+(1-\pi)\rho_{0}(\theta,\tilde{\mu}))+\theta(\pi(1-\rho_{0}(\theta,\tilde{\mu})))-\theta<0\end{aligned}$$

so that for λ sufficiently small it must be the case that old consumers with prior $(1-\theta)$ never adhere to the firm's report regardless of z and n, and therefore that $D(\rho,1-\theta)=D^{L}(\rho,1-\theta)=D^{R}(\rho,1-\theta)=0$. It then follows that $\Delta(L;\rho^{*}(\theta,\tilde{\mu}),\theta,\tilde{\mu})=0$, so that $\rho^{*}(\theta,\tilde{\mu})$ is an equilibrium for sufficiently small λ. ■

PROOF OF PROPOSITION 7

Pick $(\theta,\tilde{\mu})$ such that $\rho^{*}(\theta,\tilde{\mu})=0$. Choose some $\rho>0$. It must be the case by assumption that

$$\sigma(\rho,\theta)\left[(1-\tilde{\mu})D(\rho,\theta)+\tilde{\mu}\pi D^{L}(\rho,\theta)+\tilde{\mu}(1-\pi)D^{R}(\rho,\theta)\right]<0$$

so it is sufficient to show that

$$(1-\sigma(\rho,\theta))\left[(1-\tilde{\mu})D(\rho,1-\theta)+\tilde{\mu}\pi D^{L}(\rho,1-\theta)+\tilde{\mu}(1-\pi)D^{R}(\rho,1-\theta)\right]\leq0$$

in order to demonstrate that $\Delta(L;\rho,\theta,\tilde{\mu})<0$. First consider the case of no feedback. Old consumers with prior $(1-\theta)$ value the firm's report weakly more if it reported L when they were young than if it reported R. Thus we have three cases. First, if old consumers with prior $(1-\theta)$ do not follow the firm's report regardless of its report when they were young, then $D(\rho, 1-\theta)=0$. Second, if they follow the firm's report when old regardless of its report when they were young, then $D(\rho, 1-\theta)<0$ by our proposition 6. Third, if they follow if and only if the firm reports R, then $D(\rho, 1-\theta)<0$. Therefore $D(\rho, 1-\theta) \leq 0$. A similar logic shows that $\left(\pi D^{L}(\rho, 1-\theta)+(1-\pi) D^{R}(\rho, 1-\theta)\right) \leq 0$, so that $\Delta(L ; \rho, \theta, \tilde{\mu})<0$ and therefore the unique stationary informative equilibrium is honest reporting. ■

PROOF OF PROPOSITION 8

The posterior on quality will be an increasing function of the likelihood ratio:

$$\mathcal{L}=\frac{\bar{\pi}^{L}(x)(1-\theta)+\bar{\pi}^{R}(x) \theta}{\pi^{L}(x)(1-\theta)+\pi^{R}(x) \theta}$$

The sign of $d \mathcal{L} / d \theta$ is the same as the sign of:

$$\begin{aligned}
& \left[\pi^{L}(x)(1-\theta)+\pi^{R}(x) \theta\right]\left[\bar{\pi}^{R}(x)-\bar{\pi}^{L}(x)\right] \\
& \qquad -\left[\bar{\pi}^{L}(x)(1-\theta)+\bar{\pi}^{R}(x) \theta\right]\left[\pi^{R}(x)-\pi^{L}(x)\right] \\
= \; & \pi^{L}(x) \bar{\pi}^{R}(x)-\pi^{R}(x) \bar{\pi}^{L}(x)
\end{aligned}$$

The result then follows by equation 7.■

PROOF OF PROPOSITION 9

The first step is to show that the firm's equilibrium strategy must involve a cutoff k^{*} such that the firm reports R if and only if $x \geq k^{*}$. Let $C \subset(-b, b)$ be the set of signals such that the firm reports R. In any informative equilibrium C is nonempty. Pick some $x^{\prime} \in C$. It is sufficient to prove that the firm will want to report R if it receives a signal of $x^{\prime \prime} \in\left(x^{\prime}, b\right)$. Note that

$$\begin{aligned}
\Delta\left(x^{\prime} ; C, \theta\right) \;=\; & \tilde{\mu}\left[\hat{\theta}\left(x^{\prime}\right) \hat{\lambda}^{R}(R ; C)-\left(1-\hat{\theta}\left(x^{\prime}\right)\right) \hat{\lambda}^{L}(L ; C)\right] \\
& +(1-\tilde{\mu})\left[\hat{\lambda}^{0}(R ; C, \theta)-\hat{\lambda}^{0}(L ; C, \theta)\right]
\end{aligned}$$

must be weakly positive when $x^{\prime}$ since $x^{\prime} \in C$. (Here the firm's strategy s is denoted by C.) Since by (R1) $\hat{\theta}\left(x^{\prime \prime}\right)>\hat{\theta}\left(x^{\prime}\right)$, it follows that $\Delta\left(x^{\prime \prime} ; C, \theta\right)>\Delta\left(x^{\prime} ; C, \theta\right)$, so that $x^{\prime \prime} \in C$ as desired.

Let k^{*} be the cutoff that characterizes the firm's strategy. We must have $\Delta\left(k^{*} ; k^{*}, \theta\right)=0$. (Here we abuse notation slightly in representing the strategy s by the cutoff value k^{*}.) Note that by (R2) and (R3) we have

$$\begin{aligned}
\lim_{k^{*} \rightarrow-b} \Delta\left(k^{*} ; k^{*}, \theta\right) & <0 \\
\lim_{k^{*} \rightarrow b} \Delta\left(k^{*} ; k^{*}, \theta\right) & >0
\end{aligned}$$

and that (R1) combined with Bayesian inference about quality implies that $\Delta\left(k^{*} ; k^{*}, \theta\right)$ is strictly increasing in k^*. Therefore there exists a unique equilibrium cutoff value $k^* \in (-b, b)$. Next note that because of (R4)

$$\Delta(0; 0, \theta) = (1 - \tilde{\mu}) \left(\hat{\lambda}^0 (R; 0, \theta) - \hat{\lambda}^0 (L; 0, \theta) \right)$$

which is strictly positive if $\theta > \frac{1}{2}$, 0 if $\theta = \frac{1}{2}$, and strictly negative if $\theta < \frac{1}{2}$. Therefore $k^* \lesseqgtr 0 \Leftrightarrow \left(\theta \gtreqless \frac{1}{2}\right)$.

The final task is to prove the comparative statics on $\tilde{\mu}$. It is possible to show that (R1) and (R5) imply that

$$\hat{\theta}(k^*) \hat{\lambda}^R (R; k^*) < \left(1 - \hat{\theta}(k^*)\right) \hat{\lambda}^L (L; k^*)$$

for $k^* < 0$, and since $\Delta(k^*; k^*, \theta) = 0$ in equilibrium we therefore have that $\Delta(k^*; k^*, \theta)$ is strictly decreasing in $\tilde{\mu}$. Thus whenever $k^* < 0$, k^* increases as $\tilde{\mu}$ increases, which proves the desired comparative statics with respect to μ and J. ■

PROOF OF PROPOSITION 10

Note that in equilibrium α^* all types except possibly the normal type who has seen a signal of L strictly prefer their equilibrium strategy to any other strategy. Note also that $\frac{\partial \mathbf{\Delta}_L^0(\alpha^*)}{\partial \alpha_L^0} < 0$ whenever $\alpha_L^0 > 0$, so that α^* is stable.

Since consumers strictly prefer to follow the report of a high-type firm, each firm will try to maximize old consumers' assessments of its quality. It is then immediate that in any equilibrium in which high-type firms randomize given a signal d, an increase in the probability of high-type firms reporting R will lead to an increase in the incentive to do so, i.e. that $\frac{\partial \mathbf{\Delta}_d^1(\alpha^*)}{\partial \alpha_d^1} > 0$, so that such equilibria fail to meet the definition of stability. Therefore the only stable equilibria in this case are those in which the high type reports honestly, and Proposition 1 implies that α^* is the unique informative equilibrium in the class in which the high type is honest. ■

C Appendix: Evidence from the Gallup Poll of the Islamic World

In this appendix, we study the relationship between prior opinions and assessments of news media quality using survey evidence from the Muslim world on consumer evaluations of the satellite news network CNN International. This exercise has two limitations relative to the experimental approaches discussed in section 2.2. First, we cannot control exactly what information survey respondents receive. If two individuals give different evaluations of the quality of CNN, this could occur because the individuals reacted differently to the same content, or because they saw slightly different content (say, two different CNN news programs). Second, because the data are cross-sectional, we do not have a direct measure of the opinions respondents possessed before exposure to CNN. We will therefore need to seek proxies for pre-existing attitudes and ask whether these proxies are correlated with perceptions of CNN's quality.

The data come from the 2002 Gallup Poll of the Islamic World (The Gallup Organization, 2002). The sample consists of 10,004 respondents from nine predominantly Muslim countries.[34]

[34] Sample sizes by country are as follows: Pakistan (2,043), Iran (1,501), Indonesia (1,050), Turkey (1,019),

Respondents in all countries (except Iran) were asked to report whether each of the following five descriptions applies to CNN: has comprehensive news coverage; has good analyses; is always on the site of events; has daring, unedited news; has unique access to information. We have constructed an overall measure of perceived quality equal to the share of these characteristics the respondent feels CNN possesses. This measure has a correlation of over .7 with each individual component, and therefore seems like a good proxy for the respondent's overall attitude toward the quality of CNN's news coverage.

As we discuss in Gentzkow and Shapiro (2004), relative to the media environment in the sample countries, CNN is quite pro-United States in its coverage. In the context of the above model, then, we would expect respondents whose prior opinions are less pro-United States to rate CNN as being of lower quality. To execute this test, we will first need a measure of *prior* opinions–opinions formed before exposure to CNN content. We will use the respondent's ranking of the importance of religion in her life relative to four other concepts (own family/parents, extended family/local community, country, and own self). The rank varies from one to five, and we have re-scaled (by subtracting one and dividing by four) so that the measure varies from zero to one, with one implying that religion is the most important among the list of five. It seems likely that the importance of religion in the respondent's life is predetermined with respect to television news viewership.

We predict that respondents who rank religion as being of greater importance are likely to have more negative prior attitudes toward the United States. Columns (1) and (2) of Table 1 check this prediction by regressing a measure of the respondent's general attitude toward the United States on the importance of religion variable. The measure of the respondent's general attitude comes from a question of the form "In general, what opinion do you have of the following nations?...The United States." Responses range from one ("very unfavorable") to five ("very favorable"). We have re-scaled this measure to vary from zero to one, with one being the most favorable toward the United States.

As column (1) shows, respondents who indicate that religion plays an important role in their lives tend to report less favorable attitudes toward the United States. Column (2) shows that this relationship is robust to the inclusion of a wide set of demographic controls, indicating that it is not likely to be driven by demographic variation in the population. Similar results can be obtained using alternative measures of attitudes toward the United States, such as beliefs about the justifiability of the September 11 attacks (results not shown).

Now that we have established the relationship between the importance of religion and attitudes toward the United States, we can ask whether respondents who are likely to have a negative prior opinion toward the United States–that is, respondents for whom religion is more important–rate CNN as being of lower quality. Column (3) shows that this prediction of the above model is indeed correct. An increase in the importance of religion of one standard deviation is associated with a decrease in the perceived overall quality of CNN of about five percent of a standard deviation. As column (4) shows, this finding is robust to the inclusion of a large set of demographic controls.

Lebanon (1,050), Morocco (1,000), Kuwait (790), Jordan (797), and Saudi Arabia (754). Other than a slight oversampling of urban households, the samples are designed to be representative of the adult (18 and over) population in each country. Further details on sample selection and survey methodology are available at http://www.gallup.com/poll/summits/islam.asp.

Appendix Table: Prior opinions and assessments of media quality

	(1)	(2)	(3)	(4)
	General attitude toward US (Mean = .33, SD = .33)		Overall CNN quality rating (Mean = .10, SD = .24)	
Importance of religion	-0.1711	-0.1520	-0.0418	-0.0291
(Mean = .76, SD = .30)	(0.0132)	(0.0132)	(0.0101)	(0.0100)
Country fixed effects?	Yes	Yes	Yes	Yes
Demographic controls?	No	Yes	No	Yes
N	8566	8566	7451	7451
R^2	0.1432	0.1597	0.1575	0.1745

Notes: Respondents with missing data on dependent variable or importance of religion have been omitted from the regressions reported. Results are weighted as recommended by the data providers. Demographic controls include dummies for education, gender, age, urban/rural status, marital status. Missing data dummies are included for all demographic controls.