

NBER WORKING PAPER SERIES

MEDIA BIAS AND REPUTATION

Matthew Gentzkow
Jesse M. Shapiro

Working Paper 11664
<http://www.nber.org/papers/w11664>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2005

We are grateful to Alberto Alesina, Attila Ambrus, Nigel Ashford, Chris Avery, Heski Bar-Isaac, Gary Becker, Tyler Cowen, Jonathan Feinstein, 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, the Institute for Humane Studies, and the University of British Columbia 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. Karen Bernhardt, Fuhito Kojima, Jennifer Paniza and Tina Yang provided excellent research assistance. Gentzkow acknowledges financial assistance from the Social Science Research Council and the Centel Foundation/Robert P. Reuss Faculty Research Fund. 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@uchicago.edu. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Matthew Gentzkow and Jesse M. Shapiro. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Media Bias and Reputation
Matthew Gentzkow and Jesse Shapiro
NBER Working Paper No. 11664
September 2005
JEL No. L82, L10, D83

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 of the world, and that competition between independently owned news outlets can reduce bias. We present a variety of empirical evidence consistent with these predictions.

Matthew Gentzkow
University of Chicago
5807 S. Woodlawn Avenue
Chicago, IL 60637
gentzkow@chicagogsb.edu

Jesse M. Shapiro
University of Chicago
5807 S. Woodlawn Avenue
Chicago, IL 60637
and NBER
jmshapir@uchicago.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 (Groseclose and Milyo, forthcoming).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 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. 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

¹The 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.

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.

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 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 foreshadows our third result: Competition in the news market can lead 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. We also show that if all firms in a market are jointly owned, bias can

remain unchanged even as the number of firms gets large.

At the end of the paper, we present empirical evidence on the determinants of bias. We review a range of existing evidence suggesting 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 discuss anecdotal evidence suggesting that media firms in more competitive markets have stronger incentives to reveal important information, 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.

Formally, our work is most closely related to the literature on “herding on the priors,” which considers the way agents’ incentives to act on or reveal information depend on the prior beliefs of those who will ultimately determine their rewards. In this vein, Brandenburger and Polak (1996) show that a firm manager concerned about his firm’s current stock price may choose the action favored by shareholders’ priors even when he has private information showing this is inefficient. The manager’s desire to maximize current stock prices plays a role similar to the reputational incentives in our framework: it gives the decision-maker an incentive to slant its action toward the prior beliefs of another agent. Prendergast (1993) shows that similar concerns will lead a worker to skew her reports to match the data that her manager has received, again leading valuable information to be lost in equilibrium.² In contrast to much of this literature, however, the reputational concerns we model provide an additional incentive for honesty, which in turn guarantees the presence of an informative equilibrium. In this sense, our model is also related to work on reputational effects in sender-receiver games (see, e.g., Morris, 2001; Effinger and Polborn, 2001; Avery and Meyer, 2003; Olszewski, 2004; and Ottaviani and Sørensen, 2005a and 2005b). We deviate from these papers in highlighting the importance of the receiver’s prior beliefs for the equilibrium reporting strategy, and in showing the effects of ex-post revelation on equilibrium reporting.³

²See also Heidues and Lagerlöf, 2000 for an application to political competition.

³See Ely and Välimäki (2003) and Ely, Fudenberg and Levine (2002) for other models in which reputational concerns lead to distortions in equilibrium.

Topically, our work is related to the growing body of economic research on media bias. Existing economic models of bias all take as given that some agents in the economy—consumers (Mullainathan and Shleifer, forthcoming), reporters (Baron, 2003), or governments (Besley and Prat, 2004)—*prefer* for news suppliers to distort the information they provide.⁴ 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. While we do not deny that some agents may prefer for the news media to distort their reports, our findings suggest that caution is warranted in interpreting media slant as evidence for such tastes. Additionally, our framework generates novel, testable predictions that distinguish it from these existing theories. Most notably, our prediction that increased competition lowers the incentive to bias reports toward consumer priors contrasts sharply with that of Mullainathan and Shleifer (forthcoming), who argue that increased competition will *tighten* the connection between priors and reports.

In the next section we discuss the role of reputational incentives in media markets. In section 3 we present the model for a simple monopoly case, and show that equilibrium bias is correlated with consumer priors and decreasing in the amount of ex post feedback. In section 4, we extend the model to allow for multiple firms, and develop the intuition that competition can reduce bias by increasing the likelihood that erroneous reports are exposed. In section 5, 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 6 presents empirical evidence supporting our key findings, and section 7 concludes.

2 Credibility, Quality, and Bias 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

⁴An 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.

sources depend on prior beliefs. We also present evidence confirming the intuition suggested in the first paragraph of the introduction that firms' reporting strategies are highly related to the prior beliefs of their consumers.

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 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

⁵To 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.

⁶An 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).

2.2 The Influence of Priors on Quality Assessments

Our model will draw upon a general property of Bayesian updating about source quality; namely, that a source is judged to be of higher quality when its reports are more consistent with the agent’s prior beliefs.⁷ 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.⁹ Taken together these pieces of evidence strongly suggest that prior beliefs influence consumers’ judgements of quality in the way a Bayesian model would predict.

⁷Our supplemental appendix, available online, provides a proof of this property under fairly general conditions. See also Prendergast (1993) and Brandenburger and Polak (1996).

⁸This 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.

⁹Other 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 our supplemental appendix, 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.

2.3 The Influence of Consumer Priors on Media Reports

A large body of anecdotal evidence suggests a connection 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.

3 Media Bias in a Monopoly

In this section we present a simple model in which a single media firm reports to homogeneous consumers. There is a binary state of the world $S \in \{L, R\}$ and a large population of consumers who must each choose a binary action $A \in \{L, R\}$ which gives payoff 1 if $A = S$ and payoff 0 otherwise.¹⁰ We assume that consumers and the firm place prior probability θ and θ^F , respectively, on the true state being R .¹¹

At the beginning of the game, the firm receives a signal $s \in \{l, r\}$ about the true state, whose distribution depends on the firm’s quality. With probability λ , the firm is “high quality” and has a signal that perfectly reveals the true state. With probability $(1 - \lambda)$, the firm is “normal,” and has an imperfect but informative signal distributed according to

$$\Pr(l | L) = \Pr(r | R) = \pi,$$

where $\pi \in (\frac{1}{2}, 1)$. We assume that the firm knows its own quality with certainty. We also assume that $\theta \in [\frac{1}{2}, \pi)$ and $\theta^F \in (1 - \pi, \pi)$. These restrictions ensure that the firm’s best guess as to the true state will depend on its signal and also that its signal will be valuable to consumers.¹² The

¹⁰We 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).

¹¹Although we will allow for the possibility that $\theta = \theta^F$, the choice to specify firm and consumer beliefs separately rather than assuming a common prior does affect the results slightly because it means the comparative statics we present on θ will hold the firm prior constant. As we point out below, however, the comparative statics would only be strengthened if we moved firm and consumer priors together.

¹²In the case where $\pi < \theta^F$, the firm’s best guess is that the true state is R regardless of its signal. We would want to define “bias” in this case to be the probability that the firm’s report deviates from its best guess of the true state

assumption that $\theta \geq \frac{1}{2}$ is without loss of generality and serves to limit the set of cases we need to consider.

After seeing its signal, the firm publishes a report $\hat{s} \in \{\hat{l}, \hat{r}\}$. Normal firms are free to report either \hat{l} or \hat{r} , and we write a normal firm's strategy conditional on its signal by $\sigma_s(\hat{s}) = \Pr(\hat{s} | s)$. Without loss of generality, we will restrict attention to strategies with $\sigma_r(\hat{r}) \geq \sigma_l(\hat{r})$ (cases where this does not hold are equivalent to a relabeling of \hat{l} and \hat{r}). We assume that high quality firms always report their signal (and thus the true state) honestly.¹³

These assumptions treat bias as pure distortion of the facts. In reality, of course, 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. Note, however, that in all of these cases, different slants convey information to consumers about what the firm believes is the true state of the world. For example, Al Jazeera's opening paragraph about the Samarra incident uses evidence selectively to imply that the American attack harmed civilians. What is important for modelling is the information conveyed in such a report, not the form the communication takes. In this sense, our model could be a useful approximation to many technologies for bias, provided that these different strategies involve conveying different impressions about the facts underlying the news story.¹⁴

Consumers choose whether or not to purchase the firm's product and thus learn the value of \hat{s} . Rather than specify the pricing and profit structure in detail, we make two reduced-form assumptions: (i) all consumers purchase the product in equilibrium; (ii) a change in the distribution of \hat{s} (conditional on the true state) that increases consumers' willingness to pay for the firm's report does not reduce consumer welfare and strictly increases firm profits. These assumptions would hold, for example, in a model where the firm sets a monopoly price and extracts all of the consumer surplus.¹⁵

(the proper definition with respect to welfare). Bias would therefore involve the firm reporting \hat{l} "too often" rather than \hat{r} too often as in the case we will analyze. However, the key comparative statics, that firm reports are correlated with consumer priors and that bias decreases with feedback, would remain unchanged.

¹³Allowing both types to have discretion would complicate the exposition and lead the model to have multiple equilibria in general. As we discuss in more detail in the supplemental appendix, however, the equilibrium we analyze in which high types report truthfully is in fact the unique equilibrium of the more general model under an intuitive stability criterion.

¹⁴In the supplemental appendix, we extend the model to allow for a continuous underlying signal, which approximates the case where firms receive several signals each period and choose which to print and which to omit.

¹⁵Alternatively, we could relax assumption (i) and allow the number of consumers who buy the firm's product to

After making the purchase decision and possibly learning \hat{s} , each consumer chooses an action A . Note that if consumers received utility from this action immediately, they would also know the true state. Since we wish to model feedback about the truth explicitly, we assume consumers do not receive utility until the end of the game, but all consumers immediately learn the true state with probability $\mu \in [0, 1]$. Letting $X \in \{L, R, 0\}$ denote the feedback received, with $X = 0$ indicating the case of no feedback, we will abuse notation slightly and denote by $\lambda(\hat{s}, X)$ a consumer's posterior probability that the firm is high quality given a report of \hat{s} and feedback of X .

Finally, we assume that the firm receives a continuation payoff $\sum_i f(\lambda_i)$, where f is a continuous and strictly increasing function of its argument and λ_i is the posterior belief of consumer i . This could easily be derived by assuming a second period with exactly the same structure as the game just described. Second period profits would then be increasing in λ_i because each consumers' expected gain from learning \hat{s} will be higher the greater the probability the firm is high quality. It can also be generated by a repeated structure in which overlapping generations of consumers interact with a long-lived firm (Gentzkow and Shapiro, 2005).

The timing of the game is summarized in figure 3. Note that the structure of the game implies that the firm's choice of a report \hat{s} only affects its objective function via the continuation payoffs $f(\lambda_i)$. Therefore, any \hat{s} played in equilibrium must maximize the expectation of the continuation payoff over possible realizations of the consumer feedback.

3.1 Inferring quality from content

As a first step to analyzing the outcome of this game, we consider how a consumer's posterior belief about firm quality depends on the firm's strategy, the accuracy of its signal, π , and the consumer's prior, θ . We focus for now on the case when consumers do not receive feedback.

Consider the posterior after a report of \hat{r} and no feedback, $\lambda(\hat{r}, 0)$. Setting aside the degenerate case where the firm always reports \hat{l} , this is a strictly increasing function of the likelihood ratio:

$$\frac{\Pr(\hat{r} \mid \text{high})}{\Pr(\hat{r} \mid \text{normal})} = \frac{\theta}{\theta [\sigma_r(\hat{r}) \pi + \sigma_l(\hat{r}) (1 - \pi)] + (1 - \theta) [\sigma_r(\hat{r}) (1 - \pi) + \sigma_l(\hat{r}) \pi]}.$$

depend on the expected value of the information in the firm's report. A firm that set a zero price but extracted revenues from advertising would then be better off the more useful its reports were expected to be.

The key observation that drives many of the results below is that the more consumers' prior beliefs favor the state R , the more highly they will rate the quality of a firm that reports \hat{r} . To see this, note that the derivative of the likelihood ratio with respect to θ has the same sign as $\sigma_r(\hat{r})(1-\pi) + \sigma_l(\hat{r})\pi$, which is strictly positive under our maintained assumption that $\pi < 1$. Intuitively, as θ increases, the probability that a high type reports \hat{r} increases faster than the probability that a normal type reports \hat{r} , since the latter's report is more weakly related to the true state. This makes \hat{r} a better indicator of quality.

We can also see immediately that the likelihood ratio is strictly decreasing in both $\sigma_r(\hat{r})$ and $\sigma_l(\hat{r})$ (since increasing these makes \hat{r} a better indicator that the firm is normal). It is also weakly decreasing in π (strictly if $\theta > \frac{1}{2}$ and $\sigma_r(\hat{r}) > \sigma_l(\hat{r})$) since increasing π makes normal firms more likely to receive a signal of r and thus more likely to report \hat{r} , but does not affect the distribution of the high type's reports.

Applying the same logic to beliefs after a report of \hat{l} , we arrive at lemma 1.

Lemma 1 *Suppose that normal firms report both \hat{r} and \hat{l} with positive probability and that $\theta > \frac{1}{2}$. Then $\lambda(\hat{r}, 0)$ is strictly increasing in θ and strictly decreasing in $\sigma_r(\hat{r})$ and $\sigma_l(\hat{r})$. Similarly, $\lambda(\hat{l}, 0)$ is strictly decreasing in θ , and strictly increasing in $\sigma_r(\hat{r})$ and $\sigma_l(\hat{r})$. Finally, $\lambda(\hat{r}, 0)$ is decreasing in π and $\lambda(\hat{l}, 0)$ is increasing in π , strictly whenever $\sigma_r(\hat{r}) > \sigma_l(\hat{r})$.*

3.2 Two special cases

To fix ideas, we treat first the special cases of no feedback ($\mu = 0$) and perfect feedback ($\mu = 1$). In the no feedback case, the consumer's posterior on the firm's quality will be $\lambda(\hat{s}, 0)$, which depends only on the firm's report. In equilibrium we must have $\lambda(\hat{r}, 0) = \lambda(\hat{l}, 0)$. To see this, note that if $\lambda(\hat{r}, 0) > \lambda(\hat{l}, 0)$, normal firms would prefer to always report \hat{r} which in turn would imply $\lambda(\hat{l}, 0) = 1$, a contradiction.

In order for the posterior to be independent of the firm's report, it must be the case that, from the consumer's perspective, the equilibrium distributions of the high and normal types' signals are the same. This requires that the normal firm report \hat{r} with probability θ and \hat{l} with probability $(1-\theta)$. Note that this immediately rules out the normal firm truthfully reporting its signal, since a truthful normal firm reports \hat{r} with probability $(1-\theta)(1-\pi) + \theta\pi < \theta$.

There are many possible strategies for the normal firm which would cause its signal to have the necessary distribution, and thus many possible equilibria.¹⁶ Since the firm must distort its signal in a way that increases the overall probability of reporting \hat{r} , all of these will involve some garbling of the signal l (that is $\sigma_l(\hat{r}) > 0$). For now, we will focus on the most informative equilibrium—the one in which the firm does not simultaneously garble r signals.¹⁷ We show below that with even a small probability of feedback, μ , the model has a unique equilibrium, and that the one considered here is the limit of these unique equilibria as $\mu \rightarrow 0$.

In equilibrium, then, the firm biases its signal in the direction of consumers' prior.¹⁸ Using lemma 1, we can see immediately that the intensity of this bias, $\sigma_l^*(\hat{r})$, will be increasing in the strength of the prior θ . Increasing θ increases $\lambda(\hat{r}, 0)$ and decreases $\lambda(\hat{l}, 0)$, thus making it more attractive to report \hat{r} ; $\sigma_l(\hat{r})$ must then increase to restore equilibrium. The same logic shows that increasing π reduces the intensity of bias. Intuitively, since R is the most likely state, a normal firm with a more accurate signal is more likely to report \hat{r} , and so requires less distortion to match the high type's signal. Finally, note that since the prior on quality λ cancels out in the equilibrium condition $\lambda(\hat{r}, 0) = \lambda(\hat{l}, 0)$, bias will be independent of λ . 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).

Consider, next, the special case in which the consumer receives feedback with certainty ($\mu = 1$). Because any disagreement between feedback and report implies that the firm must be the normal type, the firm will always prefer to match the feedback if possible. Intuitively, this will mean the firm should report whichever state it believes is most likely after seeing its signal, s . Since we have assumed that the signal is sufficiently strong to outweigh the firm's prior (i.e. that $\pi > \theta^F > 1 - \pi$), this will lead to truthful reporting.

¹⁶For example, there is a pure babbling equilibrium in which normal firms randomly choose \hat{r} with probability θ and \hat{l} with probability $1 - \theta$ (independently of s).

¹⁷It is straightforward to show that this is the equilibrium that maximizes the probability that the consumer chooses the correct action, from the perspective of the consumer's prior θ .

¹⁸In this simple case, we can solve for the bias explicitly. Equilibrium requires that:

$$\Pr(\hat{r} \mid \text{normal}) = \sigma_r(\hat{r}) [\theta\pi + (1 - \theta)(1 - \pi)] + \sigma_l(\hat{r}) [\theta(1 - \pi) + (1 - \theta)\pi] = \theta$$

Setting $\sigma_r(\hat{r}) = 1$, we have:

$$\sigma_l(\hat{r}) = \frac{(2\theta - 1)(1 - \pi)}{\theta(1 - \pi) + (1 - \theta)\pi}.$$

It is straightforward to check that this is increasing in θ and decreasing in π .

This argument would be exactly correct if the gain in continuation payoff from correctly matching the R state were equal to the gain from correctly matching the L state. This is not *a priori* obvious, since the consumer's judgments will depend on her conjecture about the normal firm's strategy. Suppose, however, that the gain from correctly reporting R were greater—if $\lambda(\hat{r}, R) > \lambda(\hat{l}, L)$. This would mean the firm might have an incentive to deviate and report \hat{r} after seeing a signal of l but would always report truthfully after seeing r . But this would imply that normal firms are more likely to correctly match R states than L states and so we would have $\lambda(\hat{l}, L) > \lambda(\hat{r}, R)$, a contradiction. From this logic it follows that the only equilibrium is honest reporting. Note that this is also the most informative strategy, in the sense that it most efficiently conveys the firm's information to the consumer.

3.3 Equilibrium reporting

We turn next to a formal treatment of the more general case of $\mu \in (0, 1)$. The firm's expected payoff is essentially a weighted average of the two special cases described above. With probability $(1 - \mu)$, the consumer receives no additional information about the true state, and forms a posterior based on her prior beliefs and her conjectures about the firm's strategy. With probability μ , a feedback report arrives, and the consumer can base her judgment of the firm's quality on the true state. Because the equilibrium conditional on feedback involves truthful reporting, we would expect equilibrium bias in the general case to be decreasing in μ .

Formally, we can write the expected gain to reporting \hat{r} rather than \hat{l} after seeing signal s as

$$\Delta(s) = (1 - \mu) \Delta^{nf} + \mu \Delta^f(s),$$

where

$$\begin{aligned} \Delta^{nf} &= f(\lambda(\hat{r}, 0)) - f\left(\lambda\left(\hat{l}, 0\right)\right) \\ \Delta^f(s) &= \theta^F(s) [f(\lambda(\hat{r}, R)) - f(0)] - (1 - \theta^F(s)) \left[f\left(\lambda\left(\hat{l}, L\right)\right) - f(0) \right]. \end{aligned}$$

(We abuse notation slightly and let $\theta^F(s)$ represent the firm's posterior on the true state after seeing s .) Note that the two terms in square brackets are the gain to correctly matching the R

and L states respectively. These follow from the fact that a firm that fails to match the feedback report cannot be a high type, and so $\lambda(\hat{l}, R) = \lambda(\hat{r}, L) = 0$. Note that since $\theta^F(r) > \theta^F(l)$ and so $\Delta^f(r) > \Delta^f(l)$, the gain to reporting \hat{r} (\hat{l}) is strictly larger (smaller) after seeing r than after seeing l .

To gain some intuition for the comparative statics, suppose we begin from the case of perfect feedback ($\mu = 1$) in which normal firms report their signals truthfully. Because a truthful normal firm reports \hat{r} less often than a high quality firm, we know that $\Delta^{nf} > 0$ —if there were no feedback the firm would want to deviate and report \hat{r} after seeing l . Truthful reporting also implies that $\Delta^f(l) < 0$, since $\lambda(\hat{r}, R) = \lambda(\hat{l}, L)$ and $\theta^F(l) < \frac{1}{2}$ —conditional on feedback, the firm strictly prefers to report l signals truthfully.

Consider, now, what happens as we decrease μ . For a firm that saw a signal l , this shifts weight from the negative term $\Delta^f(l)$ to the positive term Δ^{nf} , decreasing the net gain to reporting truthfully. As long as firm strategies don't change, neither $\Delta^f(l)$ nor Δ^{nf} changes with μ , so we will eventually reach a value of μ strictly greater than 0 where $\Delta(l) = 0$. If μ falls beyond this point, the firm will strictly prefer to report \hat{r} even after seeing l . It cannot be an equilibrium for the firm to always report \hat{r} because in this case reporting \hat{l} would lead consumers to believe the firm was high quality for sure.¹⁹ The firm must therefore begin garbling l signals with a small probability—this increases the probability that a normal firm reports \hat{r} , causing the net reputational payoff to reporting \hat{r} to fall. There will thus be a range of high μ where the firm reports truthfully, and a range of lower μ where the firm plays a mixed strategy defined by the condition $\Delta(l) = 0$. In this range, the firm biases its signal toward consumer priors and the probability of bias is decreasing in μ .

It is easy to see that when the firm is playing a mixed strategy, the bias will also be increasing in θ . Increasing θ increases Δ^{nf} by lemma 1 and does not change $\Delta^f(s)$.²⁰ Since neither $\lambda(\hat{r}, R)$ nor $\lambda(\hat{l}, L)$ changes, this provides a strict incentive to report \hat{r} and in order to restore equilibrium the firm must garble l signals into \hat{r} more often.

¹⁹Making this argument precise requires us to handle the zero probability event that the firm deviates and reports \hat{l} but then the consumer receives feedback that the true state was R . We discuss the formalities in the proof of proposition 1.

²⁰Note that if we assumed that the firm and consumers had identical priors, then $\theta^F(s)$ and hence $\Delta^f(s)$ would be increasing in θ , thus strengthening the comparative static on θ .

Given the general form we have assumed for the continuation payoff, there is no clear comparative static with respect to π .²¹ However, in the limit as π approaches 1, a normal firm's signal becomes nearly as accurate as that of the high type, so that the normal firm's reputational incentives no longer favor distortion and honest reporting becomes the unique equilibrium strategy.

The complete characterization of equilibrium is given in proposition 1, which is proved in the appendix. Note that the equilibrium reporting strategy is now unique—with even arbitrarily little feedback, the firm will choose the most informative strategy to try to match the conditional distribution of the high type's reports. Note too that since the firm only garbles l signals, the limit of these equilibria as $\mu \rightarrow 0$ is the one discussed for the case of no feedback above.

Proposition 1 *The unique perfect Bayesian equilibrium reporting strategy is for a normal firm to report r signals truthfully and to misrepresent l signals with probability $\sigma_l^*(\hat{r}) \in [0, 1)$. This probability is weakly increasing in θ , strictly whenever $\sigma_l^*(\hat{r}) > 0$. For any $\theta \in (\frac{1}{2}, 1)$, there exists $\mu^* \in (0, 1)$ such that this probability is zero for $\mu \geq \mu^*$ and is positive and strictly decreasing in μ for $\mu < \mu^*$. The limit of $\sigma_l^*(\hat{r})$ as $\pi \rightarrow 1$ is zero.*

To make clearer how the equilibrium works, figure 4 shows the equilibrium bias as a function of θ for two different values of μ . When μ is very close to zero, there is bias for almost all θ and it is strictly increasing in θ . As μ increases the bias falls, and there is a range of low θ for which the equilibrium is truth-telling.

3.4 Welfare

Because we have modeled bias as a garbling of the firm's signal, it should be immediately clear that consumers' willingness to pay for an unbiased signal would be strictly higher than for a biased signal. Recall that we have assumed a change in the distribution of \hat{s} that increases consumers' willingness to pay for the firm's report does not reduce consumer welfare, and strictly increases firms' profits in the first period. This means eliminating bias cannot make consumers worse off (from the perspective of the consumer's prior), and that it strictly increases firm profits in the first

²¹Increasing π makes it more likely that a normal firm's report matches the feedback, and so decreases both $\lambda(\hat{r}, R)$ and $\lambda(\hat{l}, L)$. Depending on the shape of the function $f(\cdot)$, this could either increase or decrease $\Delta^f(l)$, and so the net effect on the incentive to report \hat{r} is ambiguous.

period.²²

The other element of welfare to consider is the effect of eliminating bias on the firm's expected continuation payoff, where the expectation is taken with respect to the firm's prior θ^F . This is more subtle because reducing bias would improve the accuracy of consumers' inferences about quality. This would tend to increase the continuation payoffs of high quality firms and decrease the continuation payoffs of normal firms. Importantly, however, the magnitude of this learning effect will be small if the change in consumer beliefs about quality is small. This will be true in particular if λ (the prior probability that the firm is high quality) is close to zero. To see this, observe that since the normal firm reports both \hat{r} and \hat{l} with positive probability in equilibrium, there is no event that reveals the firm to be high quality for sure. Therefore, as $\lambda \rightarrow 0$, the posteriors $\lambda(\hat{s}, X)$ become arbitrarily close to zero for all \hat{s} and X . This means that the change in continuation payoffs when bias is eliminated will approach zero.

We can thus state the following result.

Proposition 2 *Suppose that the equilibrium bias $\sigma_l^*(\hat{r})$ is bounded away from zero in a neighborhood of $\lambda = 0$. Then for small λ consumers would be weakly better off and the firm strictly better off if the firm were required to report its signal truthfully.*

It is easy to find parameters for which the equilibrium bias is indeed bounded away from zero around $\lambda = 0$. Recall that in the model with no feedback we showed that bias in the most informative equilibrium was independent of λ and also that this was the limit of the full model's equilibria as $\mu \rightarrow 0$. It is also easy to see that the bias will be continuous in λ . For any $\theta > \frac{1}{2}$ we can therefore find μ small enough that the limit of bias as $\lambda \rightarrow 0$ is strictly positive.

4 Competition and Endogenous Feedback

Thus far we have taken the probability of feedback, μ , to be exogenous. This was meant to capture a variety of ways that consumers might obtain additional information *ex post*—direct experience,

²²Note as an aside that we have not assigned consumers a continuation payoff that depends on their beliefs about the firm's quality. In a model where future interactions were defined explicitly, we might expect consumers to be better off the more they learned about firm quality. But eliminating bias must also at least weakly increase this payoff, since an unbiased signal provides strictly more information about the firm's type, and so we would still find a strict increase in overall willingness to pay.

contact with friends, and so forth. In this section, we explore a specific case in more detail: feedback that is provided by competing media firms. We discuss the kinds of incentives that would lead competing firms to provide accurate feedback in equilibrium and show that increasing competition will tend to decrease reporting bias.

We assume that there are J firms in the market. One of these firms has a “scoop” on the true state of the world—this means that the firm receives a possibly noisy signal of the true state s and reports it exactly as the monopoly firm did in the last section. That is, if the scoop firm is high quality (with probability λ) it receives a perfectly accurate signal and reports it honestly; if it is normal, it receives a noisy signal equal to the true state with probability π and is free to report either \hat{r} or \hat{l} .²³ At the end of the game, the firm receives a continuation payoff equal to $\sum_i f(\lambda_i)$. We extend the model by adding a “feedback stage” in which the remaining $J - 1$ firms observe the scoop firm’s report and also learn the true state with certainty. Each of these firms j then publishes a report which we denote $\tilde{s}_j \in \{\tilde{l}, \tilde{r}\}$. We continue to assume that there is an exogenous probability $\tilde{\mu} > 0$ that the true state is revealed after all reports have been made.

In specifying consumer behavior, we want to capture the idea that increasing the number of competing firms makes it more likely that consumers will read a feedback report *ex post*. We continue to assume for simplicity that all consumers read the scoop firm’s report and that they must choose their action A immediately after reading it. We then assume that each consumer can choose one of the $J - 1$ feedback reports or the outside option of reading none of them, which provides some exogenous constant level of utility. The utility of reading each of the feedback reports includes an i.i.d. random component with full support, so that the probability that each consumer reads at least one feedback report, denoted $\phi(J)$, is strictly increasing in J and approaches one as $J \rightarrow \infty$.

Observe that if all $J - 1$ firms report the true state accurately in the feedback stage, the assumptions we have made mean the scoop firm’s problem will be exactly equivalent to the monopoly problem analyzed above—it chooses its report to maximize its reputation knowing consumers will

²³Implicitly, we are assuming that the firm with the scoop always makes a report, but in principle we could endogenize this choice. Of course, a firm that knows it is the normal type might prefer not to make a report, so as to avoid making an error and revealing its type. This temptation would be especially strong if consumers do not know which firm has the scoop. On the other hand, if high types always make a report when they have the scoop, then there will be reputational incentives to make a report even if consumers do not know which firm has the scoop. Understanding how these trade-offs play out in equilibrium is an interesting issue beyond the scope of this paper.

either have no feedback about the truth or learn the true state with certainty. The only difference is that the probability of feedback will now be a strictly increasing function of J . Proposition 1 thus implies that given truthful reporting, bias is weakly decreasing in J , and strictly decreasing in J whenever the equilibrium calls for positive bias.

This observation treats competing firms as a mechanical device whose reports are exogenously fixed to perfectly reveal the true state. But reputational incentives such as those that drive our earlier results can also lead competing firms to uncover the true state and to report it truthfully. To see this, assume that each of the J firms may be high quality with probability λ , and that each firm j receives a continuation payoff $\sum_i f(\lambda_{ij})$, where λ_{ij} is consumer i 's posterior on j 's quality and $f(\cdot)$ is strictly increasing. Assume that in the feedback stage, high quality firms always report the true state honestly, but that normal firms can choose their reports freely. Each normal firm j will choose its feedback report \tilde{s}_j to maximize the *ex post* perception of its quality knowing that there is a probability μ that consumers receive exogenous feedback. Note that the game between a feedback firm and its readers is identical to our monopoly model above, with the restriction that the normal type knows the true state for sure rather than having a noisy signal ($\pi = 1$).

It then follows almost immediately from proposition 1 that the unique equilibrium is for all feedback firms to report the true state honestly. We showed that the limit of the unique equilibria as $\pi \rightarrow 1$ is honest reporting. The only complication at $\pi = 1$ is that if consumers expect a firm to report honestly the event that its report disagrees with the exogenous feedback becomes a zero probability event. Assuming that their belief in this case is that the firm is normal for sure easily supports the equilibrium. Any other candidate equilibrium would have positive equilibrium probabilities at all nodes and would be ruled out by the the same argument that established uniqueness for π close to 1.

Proposition 3 *Let $\bar{\mu} = \tilde{\mu} + (1 - \tilde{\mu}) \phi(J)$ denote the total probability that a given consumer receives feedback about the true state. Then in equilibrium all feedback firms report truthfully, and the scoop firm plays the equilibrium strategy of the monopoly game with feedback probability $\mu = \bar{\mu}$. Equilibrium bias is therefore weakly decreasing in the number of firms J (strictly whenever there is positive bias in equilibrium) and zero for J sufficiently large.*

This proposition is straightforward because the feedback firms are only motivated by their own reputations. Their only incentive is thus to ensure that their reports match any feedback consumers receive *ex post*. To highlight the importance of this independence, we now consider the opposite extreme where the J firms are jointly owned, a case of particular interest in the policy debate. We assume that the firms are either all high quality (with probability λ) or all normal, so reputation is formed at the level of the owner. High quality firms always report the true state truthfully in both the scoop and feedback stages. We assume that the owner receives a continuation payoff $\sum_i f(\lambda_i)$ where λ_i is consumer i 's posterior on the firms' joint quality and $f(\cdot)$ is strictly increasing.

We first show that the unique equilibrium in the feedback stage is for normal firms to confirm the report of the scoop firm regardless of the true state. To see that this is an equilibrium, suppose consumers expect this strategy and consider a (normal) firm that is reporting in the feedback stage. Because both high quality and normal firms always confirm, consumers' beliefs about quality will remain strictly greater than zero if the firm follows its strategy of repeating the scoop firm's report and there is no exogenous feedback. This occurs with positive probability since $\mu < 1$. If the firm deviates and contradicts the report we assign consumers the belief that the firm is normal for sure regardless of the exogenous feedback. The firm will then strictly prefer to confirm since if it contradicts it receives the lowest possible reputation for sure. To see that the equilibrium is unique consider an alternative in which a feedback firm contradicts the scoop firm's report with positive probability. Consumers must assign a posterior of zero when they see a contradiction. A firm whose strategy called for it to contradict would therefore strictly prefer to deviate since consumers must assign it a posterior strictly greater than zero if it confirms and there is no feedback, an event that occurs with positive probability.

Next, observe that given this equilibrium in the feedback stage, increasing the number of firms has no effect on the scoop firm's decision. Both normal firms and high-quality firms always report the same thing in the scoop and feedback stages, so consumers never change their quality assessments based on the feedback-stage report. The decision of all J firms to jointly report either \hat{r} or \hat{l} is thus exactly equivalent to the monopoly problem analyzed earlier. We therefore have the following result:

Proposition 4 *Suppose all firms are jointly owned. Then in equilibrium firms in the feedback*

stage always confirm the scoop firm's report, and the scoop firm plays the equilibrium strategy of the monopoly game with feedback probability $\mu = \tilde{\mu}$. Equilibrium bias is therefore independent of the number of firms J .

The assumption that all firms have the same quality is important for this result. If we relax this assumption, the effect of joint ownership will depend on the relative weight of the two firms' reputations in the owner's continuation payoff.

In the analysis above, we have shown that with independent ownership, truthful feedback reporting can be supported by firms' desires to maintain their own reputations for quality. It is possible, however, that feedback firms might also have an incentive to *harm* the reputation of the scoop firm, say because the two firms' products are viewed as substitutes in the continuation game. Intuitively, such rivalrous incentives can have two competing effects on the quality of ex-post feedback. On the one hand, consider a case such as we have modelled above where feedback firms can misreport. If the feedback firms were solely concerned with damaging the reputation of the scoop firm, this would give them a motivation to lie and make it more difficult to sustain an equilibrium where their reports are fully revealing. On the other hand, suppose feedback reports are verifiable but costly to issue. Then a feedback firm interested in damaging the scoop firm's reputation wouldn't bother to pay the cost to confirm a *correct* scoop report, but might be willing to undertake effort to expose an incorrect one. Therefore the stronger is the incentive to harm the scoop firm's reputation, the less likely an incorrect report is to go unchallenged. These kinds of competitive incentives might exacerbate bias in a world of unverifiable feedback, but could help discipline it in a world where feedback is verifiable but costly to report.

5 Heterogeneous Priors and Market Segmentation

In the model thus far, 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. This seems a reasonable starting point for thinking about differences in slant across markets—for example, why Al Jazeera and CNN differ, or why local newspapers may differ in their political orientation. But in many key settings of interest, firms with different biases

compete in the same market. In this section, we extend the model to consider a case in which two firms report to consumers with different priors,²⁴ and illustrate why media firms might segment the market according to prior beliefs.²⁵

We will assume there are two equally large groups of consumers, denoted L and R , with prior beliefs about the true state $1 - \theta$ and θ , respectively, where $\theta > \frac{1}{2}$. For simplicity we suppose that the firm’s prior is $\theta^F = \frac{1}{2}$, the average prior belief of the consumers. Each consumer can only view one firm’s report, and we assume they choose to view the report that will provide the greatest expected utility—that is, that maximizes the probability of choosing the correct action A . In the event that two newspapers provide the same expected utility, a consumer will randomize, choosing each with equal probability. As before, all consumers receive feedback about the true state with probability μ .

We make the same assumptions about the two firms that we made about the monopolist above. Each firm is high quality with probability λ in which case it sees the true state and reports truthfully. Otherwise, a firm is normal, it sees a signal equal to the true state with probability π , and it can choose its report freely. Both types and signals are drawn independently across firms. Each firm receives a continuation payoff $\sum_i f(\lambda_i)$, defined over consumers’ quality assessments λ_i , where $f()$ is strictly increasing.²⁶

Consider first how a consumer’s decision of which paper to read depends on her beliefs about each firm’s equilibrium reporting strategy. Because we have modeled bias as pure distortion, we should expect it to reduce the value of a signal to all consumers. Importantly, however, the perceived cost of bias will depend on a consumer’s type. A firm that biases its signals to the right—that is, that sometimes distorts l signals but always reports r signals truthfully—will be relatively more

²⁴See Morris (1995) for a discussion of the role of heterogeneous priors in economic theory, and Morris (1994) for an application with heterogeneous priors.

²⁵Mullainathan and Shleifer (forthcoming) point out a different intuition for market segmentation, namely that price competition leads firms to avoid offering similar products.

²⁶We note three points about the assumption on continuation payoffs. First, nothing would change if we let the function $f()$ differ by firms so long as it was strictly increasing for both. Second, because we have assumed that consumers can only see one firm’s report, the results would also not change if we let one firm’s continuation payoff depend on a consumer’s belief about both its own and its competitor’s quality (the firm only cares about the consumers who see its report, and these consumers’ beliefs about its competitor will be constant). Finally, assuming $f()$ is strictly increasing *does* rule out some natural specifications of the continuation game. For example, suppose there were a second period in which consumers chose whichever firm they believed to be higher quality (and firms received a fixed advertising revenue per consumer). Then $f(\lambda_i)$ would be a step function with $f' = 0$ almost everywhere. What we would require in this case is some stochastic component in consumer utility—a cost of time, utility from reading the sports section, etc.—that would “smooth out” consumer demand.

attractive to R -type consumers than to L -type consumers, because L -type consumers expect l signals to arrive more often. We can see this formally by noting that a consumer who follows the signal of a normal firm that biases l signals with probability σ^* will choose the correct action with probability $\pi + (1 - \pi)\sigma^*$ if the true state is R and probability $\pi(1 - \sigma^*)$ if the true state is L . Using her prior belief that the state is R with probability θ , an R -type consumer would calculate that if she follows the firm's signal she will choose the correct action with probability $\pi - \sigma^*(\pi - \theta)$. An L -type consumer, on the other hand, would expect to choose the correct action with probability $\pi - \sigma^*[\pi - (1 - \theta)]$ which is strictly smaller. The same logic shows that a firm that biases its signals to the left will be relatively more attractive to L -type consumers.²⁷ If the two firms have equal bias in opposite directions—if one distorts l signals with probability $\sigma^* > 0$ and the other distorts r signals with the same probability—the R -type consumers will strictly prefer the right-biased firm and the L -type consumers will strictly prefer the left-biased firm.

We can then see immediately that there can be equilibria in which the consumer market is completely segmented by type. Suppose for the moment that all type- R consumers buy from firm 1 and all type- L consumers buy from firm 2. Then the equilibrium strategy of firm 1 will be that of a monopolist in a market where all consumers have prior θ —that is, to distort l signals with some probability and report r signals truthfully—and the equilibrium strategy of firm 2 will have the same bias but in the opposite direction. If the monopoly strategy calls for distortion with any positive probability, type- R consumers will indeed choose firm 1 and type- L consumers will choose firm 2.

This is not the only equilibrium of the game, however. Consider what happens when a firm has equal numbers of consumers of each type. Suppose consumers expect the firm to report honestly and that they see no *ex post* feedback. If the firm reports \hat{r} , the R -types will increase their estimate of its quality (say to λ_{high}) and the L -types will decrease their estimate of its quality (say to λ_{low}). Its continuation payoff per consumer will be $[f(\lambda_{high}) + f(\lambda_{low})]/2$. But because the problem is symmetric, this will also be its payoff if it reports \hat{l} —the only difference is it will then be the L -types who have posterior λ_{high} . The firm must therefore be indifferent about its report conditional on no feedback. Since it strictly prefers to report truthfully conditional on feedback, truthful reporting

²⁷The relationship between prior beliefs and the information value of binary signals is explored in detail by Suen (2004).

will be an equilibrium for any firm with equal numbers of each type. We can go further and show that *any* firm with equal numbers of consumers of each type must report honestly in equilibrium. To see this, suppose that a firm sometimes reports \hat{r} after seeing l . Relative to the case of honest reporting, this strictly reduces all consumers' posteriors after an \hat{r} report and strictly increases their posteriors after an \hat{l} report. Thus, the incentive to report \hat{l} after l would be even stronger than when consumers expect honest reporting, and the firm would prefer to deviate from its presumed strategy. The same argument shows the firm cannot report \hat{l} after r .

There are then two types of equilibria with honest reporting. First, suppose all consumers of both types will buy from firm 1. Then firm 2 will be indifferent about its report (since it has no consumers) and firm 1, as we just observed, will report honestly. If firm 2 plays any strategy with positive probability of distortion, all consumers will indeed prefer to buy from firm 1. Second, suppose consumers of both types divide evenly between the two firms. Then both firms will report honestly, consumers will be indifferent between them, and consumers will indeed divide evenly.

The proof of the following proposition establishes that these three types of equilibria are generically unique. The genericity argument is required because there is one other equilibrium configuration that is possible in principle: one consumer type strictly prefers one firm while the other consumer type is indifferent between the two. We show that if such an equilibrium does exist for some payoff functions $f()$, it will not exist for small perturbations of the payoff functions.

Proposition 5 *For generic continuation payoffs, the game with heterogeneous priors has three types of equilibria:*

1. A **segmented equilibrium**, in which one firm is read only by R-group consumers and biases its reports toward \hat{r} and the other is read only by L-group consumers and biases its reports toward \hat{l} . This type of equilibrium exists for any parameter values such that $\sigma_l^*(\hat{r}) > 0$.
2. An **effective monopoly**, in which one firm reports honestly and the other lies with positive probability, and all consumers read the report of the honest firm. This type of equilibrium exists for all parameter values.
3. An **honest equilibrium**, in which both firms report honestly and consumers divide evenly between them. This type of equilibrium exists for all parameter values.

Observe that in the segmented equilibrium, the firms' strategies are endowed with the comparative statics described in proposition 1, since each firm is playing the monopoly strategy for its respective consumer base. This means that as beliefs become more extreme (θ increases), the firms' reporting strategies tend to diverge, and that as feedback strengthens (μ increases) they tend to move back toward honest reporting. Indeed, for high enough μ , the monopoly strategies call for honesty, so that in these cases all (generic) equilibria involve consumers receiving as much information as possible.

6 Evidence on the Determinants of Bias

In the model above, we argued that two factors should play a key role in determining the direction and strength of bias: *ex post* feedback and competition. In this section, we review existing evidence and present new evidence on both predictions.

6.1 Feedback

As proposition 1 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

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

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 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 1 above.²⁸

Evidence from financial reporting also suggests a role for feedback in limiting media bias. Lim

²⁸A 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 our supplemental appendix, 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.

(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.

6.2 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, provides one example of 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. 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 a third 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.²⁹

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 $bush_i$ and $gore_i$ denote the number of seconds given to each candidate 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 table 1 shows, this is indeed the case. Column (1) of the table indicates that 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 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 $\log(\text{population})$ 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,

²⁹ 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).

which could also drive differences across locations in the competitiveness of the news market.³⁰

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.³¹ 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.

7 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 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 4 offers one 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 proposition 3 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

³⁰The estimated effect of competition is also robust to controlling for the average total amount of candidate speaking time aired by stations in each market.

³¹See also Firth and Gift (1999).

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 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] Christopher Avery and Margaret Meyer. Designing hiring and promotion procedures when evaluators are biased. *Oxford University Mimeograph*, May 2003.
- [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] Timothy Besley and Andrea Prat. Handcuffs for the grabbing hand? media capture and government accountability. *London School of Economics Mimeograph*, April 2004.
- [10] Bryan L. Boulier and H. O. Stekler. Predicting the outcomes of National Football League games. *International Journal of Forecasting*, 19:257–270, 2003.
- [11] 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.

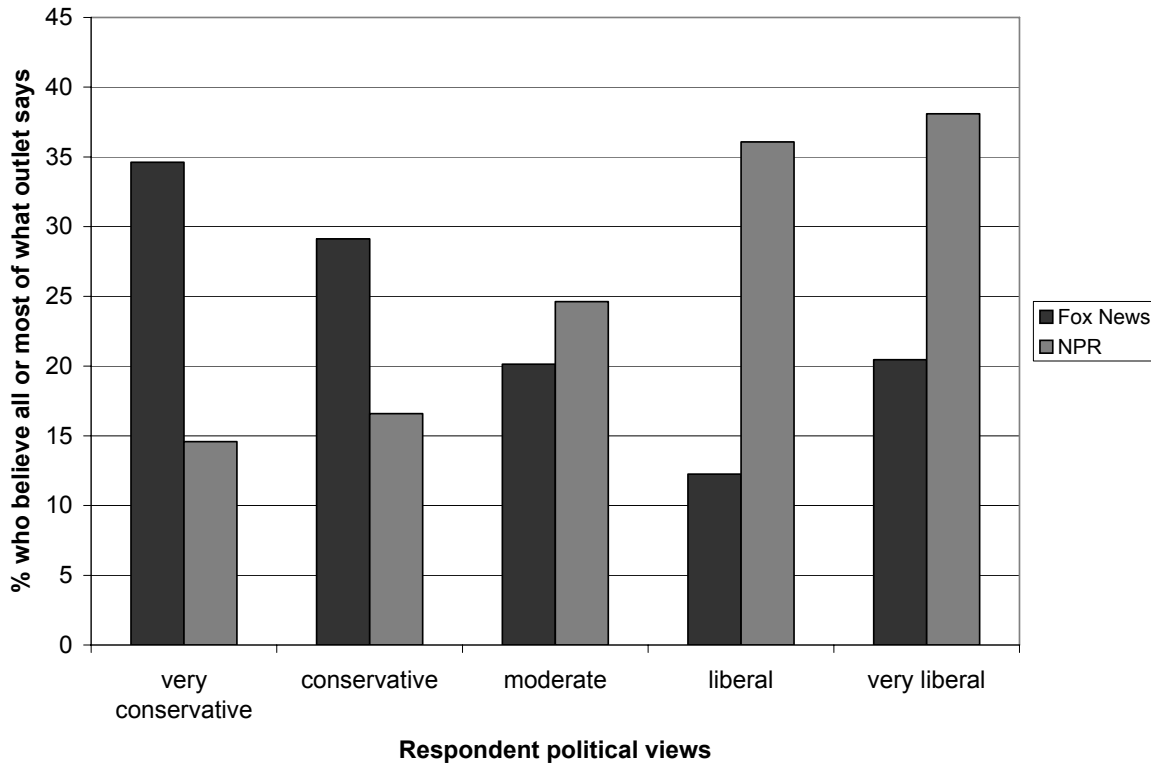
- [12] Joel Campagna. Between two worlds: Qatar’s Al-Jazeera satellite channel faces conflicting expectations. *CPJ Press Freedom Reports*, October 2001. <http://www.cpj.org>.
- [13] Mark Cooper, Gene Kimmelman, and Cheryl Leanza. Comments of Consumers Union, Consumer Federation of America, Civil Rights Forum, Center for Digital Democracy, Leadership Conference on Civil Rights and Media Access Project. December 2001.
- [14] Ann Coulter. *Slander : Liberal Lies About the American Right*. Three Rivers Press, September 2003.
- [15] 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.
- [16] Somnath Das, Carolyn B. Levine, and K. Sivaramakrishnan. Earnings predictability and bias in analysts’ earnings forecasts. *Accounting Review*, 73(2):277–294, April 1998.
- [17] Simeon Djankov, Caralee McLiesh, Tatiana Nenova, and Andrei Shleifer. Who owns the media? *Journal of Law and Economics*, 46(2), October 2003.
- [18] Alexander Dyck and Luigi Zingales. The media and asset prices. *Harvard Business School mimeograph*, August 2003.
- [19] Matthias R. Effinger and Mattias K. Polborn. Herding and anti-herding: A model of reputational differentiation. *European Economic Review*, 45:385–403, 2001.
- [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 Gentzkow and Jesse M. Shapiro. Media bias and reputation. *University of Chicago Mimeograph*, January 2005.
- [27] 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.
- [28] Matthew A. Gentzkow and Jesse M. Shapiro. Media, education, and anti-Americanism in the Muslim world. *Journal of Economic Perspectives*, Summer 2004. Forthcoming.

- [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. *Quarterly Journal of Economics*. Forthcoming.
- [31] Bernard Goldberg. *Bias : A CBS Insider Exposes How the Media Distort the News*. Perennial, January 2003.
- [32] Simon Grant, Atsushi Kajii, and Ben Polak. Intrinsic preference for information. *Journal of Economic Theory*, 83:233–259, 1998.
- [33] Tim Groseclose and Jeff Milyo. A measure of media bias. *Quarterly Journal of Economics*. Forthcoming.
- [34] 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.
- [35] Paul Heidhues and Johan Lagerlöf. Hiding information in electoral competition. *Games and Economic Behavior*, 42:48–74, 2000.
- [36] Andrew Heyward. CBS statement on bush memos. *CBSNews.com*, September 20 2004.
- [37] 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.
- [38] Martin Kaplan and Matthew Hale. *Local TV coverage of the 2000 general election*. Norman Lear Center, USC Annenberg School of Communication, February 2001.
- [39] 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.
- [40] 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.
- [41] Jonathan Koehler. The influence of prior beliefs on scientific judgments of evidence quality. *Organizational Behavior and Human Decision*, 56(1):28–55, 1993.
- [42] Andrew Kohut. Media myopia: More news is not necessarily good news. *New York Times*, July 11 2004.
- [43] David Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27:253–279, 1982.
- [44] Terence Lim. Rationality and analysts’ forecast bias. *Journal of Finance*, 56(1):369–385, February 2001.
- [45] 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.

- [46] 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.
- [47] Paul Milgrom and John Roberts. Predation, reputation and entry deterrence. *Journal of Economic Theory*, 27:280–312, 1982.
- [48] 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.
- [49] Stephen Morris. Trade with heterogeneous prior beliefs and asymmetric information. *Econometrica*, 62(6):1327–1347, November 1994.
- [50] Stephen Morris. The common prior assumption in economic theory. *Economics and Philosophy*, 11:227–253, 1995.
- [51] Stephen Morris. Political correctness. *Journal of Political Economy*, 109(2):231–265, 2001.
- [52] Sendhil Mullainathan and Andrei Shleifer. The market for news. *American Economic Review*. Forthcoming.
- [53] Sendhil Mullainathan and Andrei Shleifer. Media bias. *Harvard Institute of Economic Research Discussion Paper*, (1981), October 2002.
- [54] 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.
- [55] Allan H. Murphy and Robert L. Winkler. Reliability of subjective probability forecasts of precipitation. *Applied Statistics*, 26(1):41–47, 1977.
- [56] Allan H. Murphy and Robert L. Winkler. Probability forecasting in meteorology. *Journal of the American Statistical Association*, 79(387):489–500, September 1984.
- [57] Wojciech Olszewski. Informal communication. *Journal of Economic Theory*, pages 180–200, 2004.
- [58] The Gallup Organization. *The 2002 Gallup Poll of the Islamic World: Executive Report*. The Gallup Poll Tuesday Briefing. Princeton, NJ, February 2002.
- [59] 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.
- [60] Marco Ottaviani and Peter Norman Sørensen. Professional advice. *Journal of Economic Theory*, 2005.
- [61] Marco Ottaviani and Peter Norman Sørensen. Reputational cheap talk. *RAND Journal of Economics*, 2005.
- [62] 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.

- [63] Canice Prendergast. A theory of "yes men". *American Economic Review*, 83(4):757–770, September 1993.
- [64] Associated Press. Arab media muted in coverage of beheading; many papers play down story, others ignore it. May 12 2004.
- [65] Jana Smith Raedy, Philip Shane, and Yanhua Yang. Horizon-dependent underreaction in financial analysts' earnings forecasts. *University of Colorado Mimeograph*, September 2003.
- [66] Dan Rather. Dan Rather statement on memos. *CBSNews.com*, September 20 2004.
- [67] Donald H. Rumsfeld. Secretary rumsfeld remarks at stakeout outside ABC TV. *Department of Defense News Transcript*, October 28 2001.
- [68] Robert Satloff. Rethinking public diplomacy in the Middle East. *Weekly Standard*, 8(46), August 2003.
- [69] Wing Suen. The self-perpetuation of biased beliefs. *Economic Journal*, 114:377–396, April 2004.
- [70] 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.
- [71] 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*



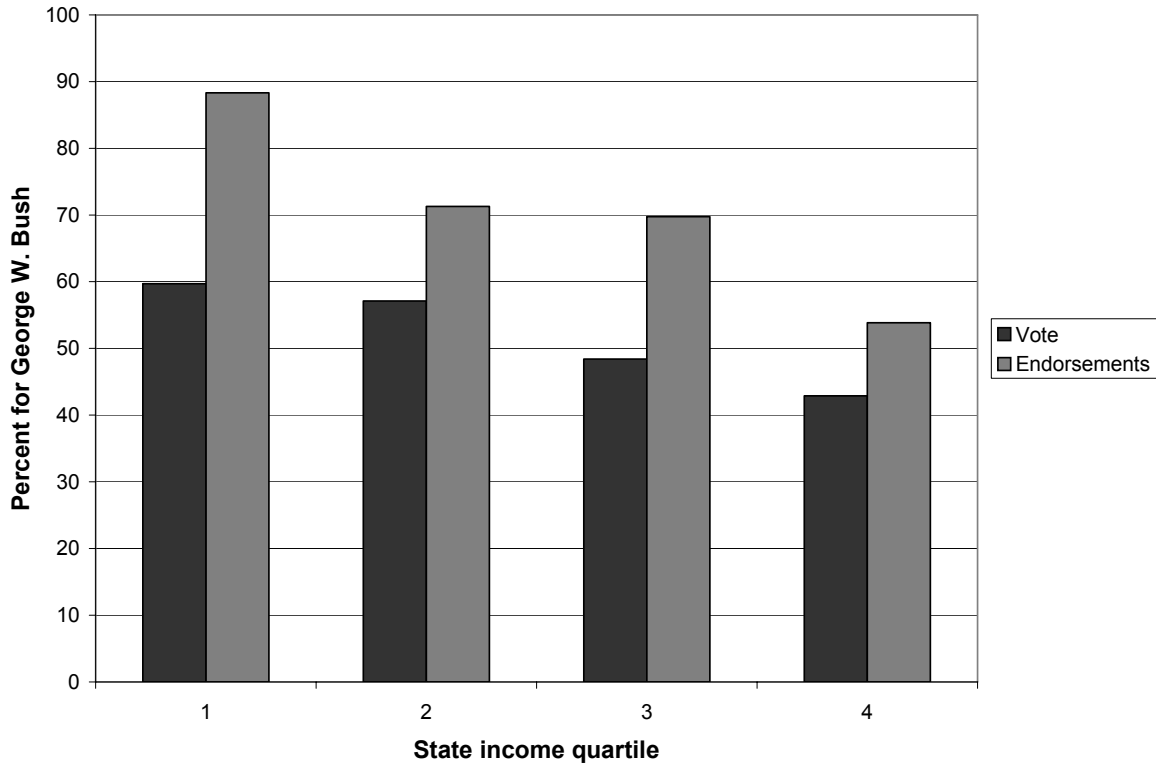
Notes: Data from the Pew Research Center for the People and the Press's 2002 News Media Believability Survey. Exact wording for survey question on respondent political views:

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 *The timing of the monopoly game*

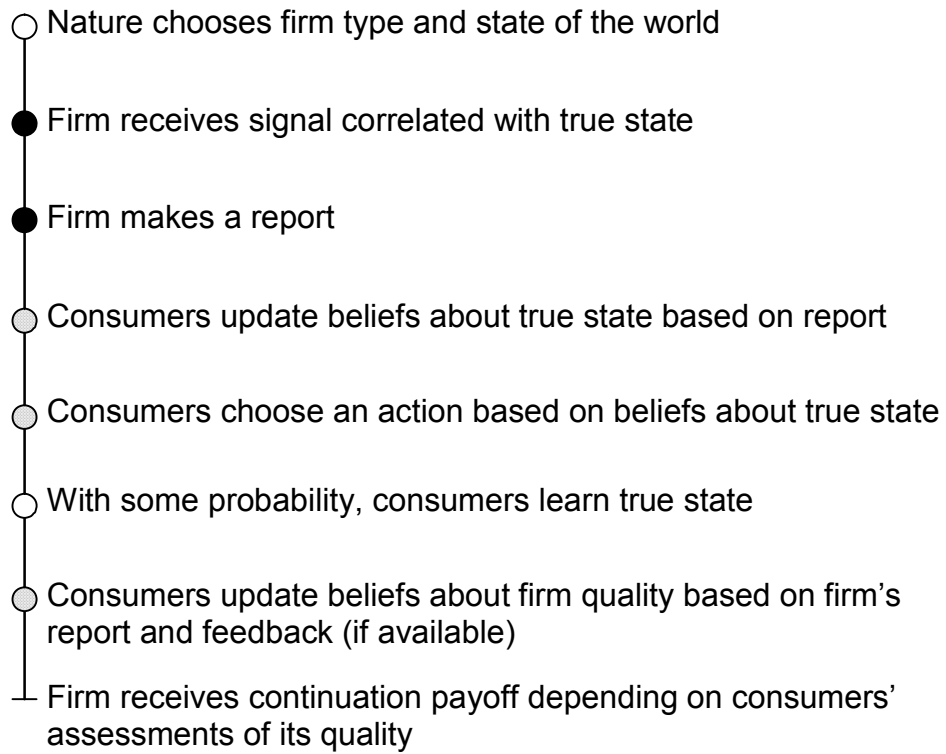
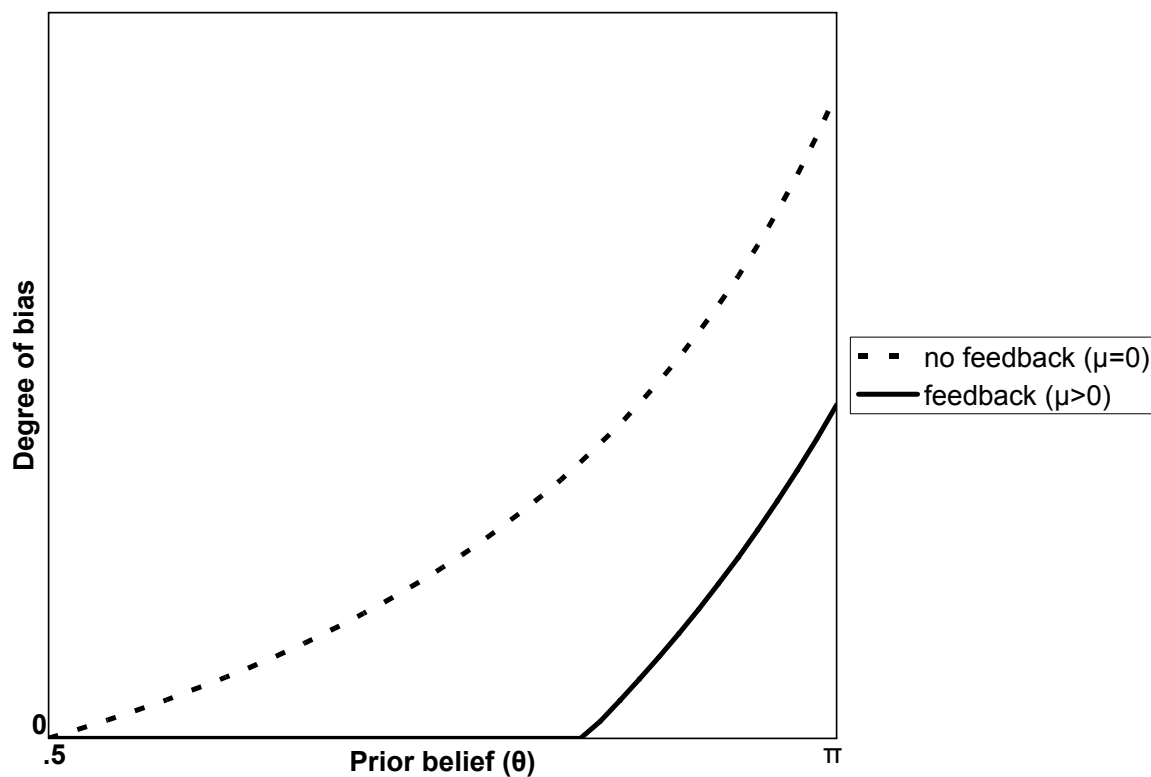
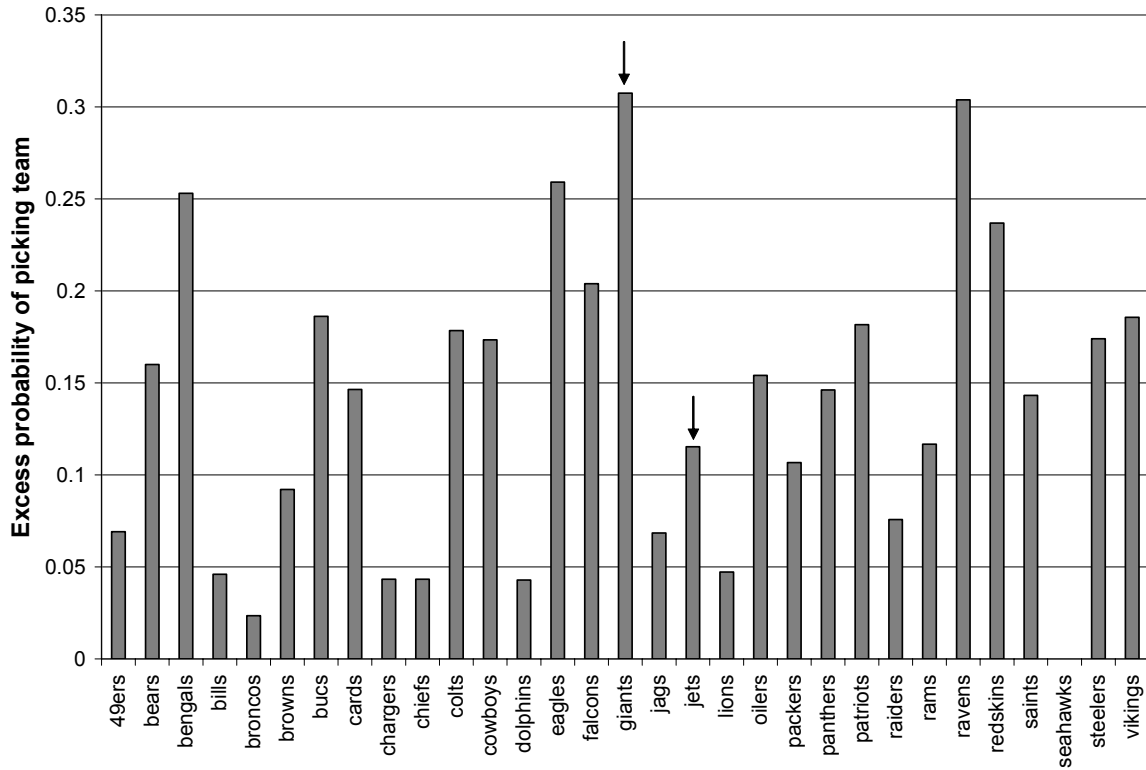


Figure 4 *A numerical example*



Notes: Vertical axis shows equilibrium probability of reporting \hat{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 ($\mu = .5$). Drawn for the case of $f(x) = x$, $\pi = .7$, and $\lambda = .3$.

Figure 5 *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

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

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.

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: Proofs

PROOF OF PROPOSITION 1

We show first that it cannot be an equilibrium for the firm to distort r signals with positive probability (i.e. for the firm to play $\sigma_r(\hat{l}) > 0$). Recall from the text that the gain from reporting \hat{l} after r is strictly smaller than the gain from reporting \hat{l} after l . If the firm sometimes reports \hat{l} after r it must therefore always report \hat{l} after l . A normal firm would then be more likely to correctly match L states than R states, implying $\lambda(\hat{r}, R) > \lambda(\hat{l}, L)$ and thus, since $\theta^F(r) > \frac{1}{2}$, $\Delta^f(r) > 0$. Since $\lambda(\hat{r}, 0) > \lambda(\hat{l}, 0)$ when the firm reports honestly, and increasing $\sigma_r(\hat{l})$ increases $\lambda(\hat{r}, 0)$ and decreases $\lambda(\hat{l}, 0)$, we also know $\Delta^{nf} > 0$. The firm would therefore strictly prefer to deviate and report \hat{r} after r . Any equilibrium must therefore have at most garbling of l signals into \hat{r} .

Observe next that it cannot be an equilibrium for the firm to *always* garble l signals into \hat{r} . Suppose that consumers expected normal firms to always report \hat{r} and consider a firm that sees a signal of l . If the firm deviates and reports \hat{l} , a consumer that sees no feedback or that receives feedback that the state is L will have posterior $\lambda(\hat{l}, L) = 1$, since only high quality firms report \hat{l} . In these cases, the firm would prefer to deviate. The event a consumer sees \hat{l} and receives feedback that the state is R has zero probability in the proposed equilibrium, and so beliefs at this node can be arbitrary. Suppose, in the worst case for the firm, that they assume that if this node is reached the firm is normal for sure ($\lambda(\hat{l}, R) = 0$). Payoffs from following the strategy are then:

$$(1 - \mu) f(\lambda(\hat{r}, 0)) + \mu [\theta^F(l) f(\lambda(\hat{r}, R)) + (1 - \theta^F(l)) f(0)].$$

Payoffs from deviating are:

$$(1 - \mu) f(1) + \mu [\theta^F(l) f(0) + (1 - \theta^F(l)) f(1)].$$

Because $\lambda(\hat{r}, 0)$ and $\lambda(\hat{r}, R)$ are both strictly less than one and $\theta^F(l) < \frac{1}{2}$ (since $1 - \pi < \theta^F < \pi$), the firm strictly prefers to deviate.

Any equilibrium must therefore have a probability of distortion $\sigma_l^*(\hat{r}) \in [0, 1)$. We now show that the equilibrium $\sigma_l^*(\hat{r})$ exists and is unique. Note first that the payoff to reporting \hat{r} after l , $\Delta(l)$, is strictly decreasing in the garbling probability $\sigma_l(\hat{r})$. This follows from the fact that increasing $\sigma_l(\hat{r})$ decreases the posterior on quality after an \hat{r} report and increases the posterior after an \hat{l} report (for any feedback), and that this is strict for no feedback or feedback that matches the firm's report. Now consider two cases. First, suppose that when consumers expect truth-telling, $\Delta(l) \leq 0$. Then truthful reporting is an equilibrium and it is unique because at any candidate $\sigma_l(\hat{r}) > 0$ we would have $\Delta(l)$ strictly negative and so the firm would prefer to deviate and always report \hat{l} after l . Second, suppose that when consumers expect truth-telling, $\Delta(l) > 0$. Then the facts that (i) $\Delta(l)$ is strictly decreasing in $\sigma_l(\hat{r})$, and (ii) $\Delta(l) < 0$ at $\sigma_l(\hat{r}) = 1$, imply that there is a unique $\sigma_l^*(\hat{r}) \in (0, 1)$ such that $\Delta(l) = 0$ when consumers expect $\sigma_l^*(\hat{r})$. This mixed strategy is then the unique equilibrium.

To derive the comparative statics with respect to θ , observe that for any $\mu \in (0, 1)$, truthful reporting of r signals, and $\sigma_l(\hat{r}) \in [0, 1)$, lemma 1 implies that $\Delta(l)$ is strictly increasing in θ . Therefore the probability $\sigma_l^*(\hat{r})$ must be increasing in θ in the mixed strategy range.

For the comparative statics on μ , suppose $\theta \in (\frac{1}{2}, 1)$, that consumers expect truthful reporting of r signals, and that we are at an equilibrium value of $\sigma_l(\hat{r})$ —that is, either $\sigma_l(\hat{r}) = 0$ or $\sigma_l(\hat{r}) > 0$ and $\Delta(l) = 0$. Then $\Delta(l)$ is strictly decreasing in μ . To see this, note first that the payoff to

reporting \hat{r} after l conditional on feedback ($\Delta^f(l)$) must be strictly negative, since the firm believes L to be the most likely state and the reward to correctly matching the L state is greater than the reward for correctly matching the R state. Second, the payoff conditional on no feedback (Δ^{nf}) must be strictly positive. This follows from the fact that $\lambda(\hat{r}, 0) > \lambda(\hat{l}, 0)$ if consumers expect truth-telling, and the fact that $\Delta(l) = 0$ if consumers expect a positive probability of distortion. The signs of $\Delta^f(l)$ and Δ^{nf} also imply that if consumers expect truth-telling, $\Delta(l)$ is strictly negative for $\mu = 1$ and becomes strictly positive as $\mu \rightarrow 0$. There is therefore a unique $\mu^* \in (0, 1)$ such that the equilibrium is truth-telling for $\mu \geq \mu^*$ and a mixed strategy for $\mu < \mu^*$, and the probability of distortion in the mixed strategy range is strictly decreasing in μ .

For the limit result on π , note that if consumers expect the firm to report honestly, it is straightforward to show that $\lim_{\pi \rightarrow 1} \Delta^{nf} = 0$ and $\lim_{\pi \rightarrow 1} \Delta^f(l) < 0$, which by continuity implies that $\lim_{\pi \rightarrow 1} \sigma_l^*(\hat{r}) = 0$. ■

PROOF OF PROPOSITION 5

We first show that each firm must distort at most one signal with positive probability. As in the monopoly case, the expected payoff to reporting \hat{r} is strictly greater after a signal r than after a signal l . (The only way that seeing r rather than l affects the firm's expected continuation payoff is that it shifts the firm's posterior on the true state toward R , increasing the expected return to reporting \hat{r} conditional on consumers receiving feedback.) Therefore a firm that sometimes reports \hat{r} after l must always report \hat{r} after r . Similarly, the return to reporting \hat{l} is strictly greater after a signal l and so a firm that sometimes reports \hat{l} after r must always report \hat{l} after l .

Now, observe that there can logically be four types of equilibria:

1. Type- R consumers strictly prefer one firm and type- L consumers strictly prefer the other firm
2. Both consumer types strictly prefer the same firm
3. Both consumer types are indifferent between the firms
4. One consumer type strictly prefers one firm and the other is indifferent between the firms

The discussion in the text established (i) that segmented equilibria exist whenever $\sigma_l^*(\hat{r}) > 0$ and are the only possible equilibria of the first type; (ii) that effective monopoly equilibria exist and that these are the only equilibria of the second type; and (iii) that honest reporting equilibria exist and that these are the only equilibria of the third type. To complete the proof we show that equilibria of the fourth type will not exist generically.

Suppose without loss of generality that the L -type consumers have a strict preference for firm 1 and the R -type consumers are indifferent between the firms. Then firm 1's customers are $2/3$ L types and $1/3$ R types while firm 2's customers are all R types. Therefore, firm 2 must be playing the monopoly strategy for R types, which will involve distorting l signals with some probability, possibly 0. Denote this probability σ^* . If $\sigma^* = 0$, both types would either be indifferent (if firm 1 reports honestly) or strictly prefer firm 2 (if firm 1 reports with bias), which contradicts our assumption. Therefore we can restrict attention to cases where $\sigma^* > 0$.

Note first that firm 1 cannot be playing the same strategy as firm 2, because then L types would also be indifferent between the two firms. Firm 1 also cannot distort l signals with a different non-zero probability because then both types would either strictly prefer firm 1 (if the probability was lower) or strictly prefer firm 2 (if the probability was higher). Firm 1 must therefore distort r signals with some positive probability σ' . From the expression for consumer value derived in the

text, we can see that to make the R types indifferent, σ' will be defined uniquely by the equation

$$\pi - \sigma' [\pi - (1 - \theta)] = \pi - \sigma^* (\pi - \theta). \quad (1)$$

Note that this immediately implies $\sigma' > 0$. In order for the strategy σ' to be an equilibrium, firm 1 must be indifferent about its report when consumers expect it to play σ' and it observes a signal r . Letting $\Delta^f(r)$ denote the relative payoff to reporting \hat{r} after r conditional on feedback, and letting $\lambda^R(\hat{s}, 0)$ and $\lambda^L(\hat{s}, 0)$ represent the posteriors of consumers of types R and L respectively, we must have:

$$\frac{2}{3} \left[f(\lambda^R(\hat{r}, 0)) - f(\lambda^R(\hat{l}, 0)) \right] + \frac{1}{3} \left[f(\lambda^L(\hat{r}, 0)) - f(\lambda^L(\hat{l}, 0)) \right] = \Delta^f(r). \quad (2)$$

Equations 1 and 2 are a system with a single unknown, σ' , so there will not be a solution for generic $f()$. ■

Supplemental Appendix to “Media Bias and Reputation”

Matthew Gentzkow
Graduate School of Business
University of Chicago

Jesse M. Shapiro*
University of Chicago and NBER

September 14, 2005

A 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 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 $S \in \{L, R\}$. Information sources, which may be high or low quality, make a report $\hat{s} \in \mathcal{D}$, where \mathcal{D} is some set of possible reports. The density of a high-quality report conditional on the state S is $\bar{\pi}(\hat{s}; S)$ and the density of a low-quality report is $\pi(\hat{s}; S)$. Here $\bar{\pi}(\cdot)$ and $\pi(\cdot)$ may be either PMFs or PDFs so long as any mass points of $\bar{\pi}(\cdot)$ are also mass points of $\pi(\cdot)$.

We say that a value \hat{s} *supports* R if $\bar{\pi}(\hat{s}; R) > \bar{\pi}(\hat{s}; L)$ —i.e. if seeing report \hat{s} 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:

$$\begin{aligned} \frac{\bar{\pi}(\hat{s}; R)}{\bar{\pi}(\hat{s}; L)} &> \frac{\pi(\hat{s}; R)}{\pi(\hat{s}; L)} \text{ if } \hat{s} \text{ supports } R; \\ \frac{\bar{\pi}(\hat{s}; L)}{\bar{\pi}(\hat{s}; R)} &> \frac{\pi(\hat{s}; L)}{\pi(\hat{s}; R)} \text{ if } \hat{s} \text{ supports } L. \end{aligned} \tag{1}$$

Suppose that a consumer has prior probability θ that the true state is R , and prior probability $1 - \theta$ that the source is high quality. The following proposition characterizes how the report \hat{s} influences the consumer’s posterior estimate of quality $\lambda(\hat{s})$.

Appendix Proposition 1 $\lambda(\hat{s})$ is strictly increasing in θ if \hat{s} supports R and strictly decreasing in θ if \hat{s} supports L .

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

$$\mathcal{L} = \frac{\bar{\pi}(\hat{s}; L)(1 - \theta) + \bar{\pi}(\hat{s}; R)\theta}{\pi(\hat{s}; L)(1 - \theta) + \pi(\hat{s}; R)\theta}$$

*E-mail: gentzkow@ChicagoGSB.edu, jmshapir@uchicago.edu.

The derivative $d\mathcal{L}/d\theta$ has the same sign as:

$$\begin{aligned} & [\pi(\hat{s}; L)(1 - \theta) + \pi(\hat{s}; R)\theta] [\bar{\pi}(\hat{s}; R) - \bar{\pi}(\hat{s}; L)] \\ & \quad - [\bar{\pi}(\hat{s}; L)(1 - \theta) + \bar{\pi}(\hat{s}; R)\theta] [\pi(\hat{s}; R) - \pi(\hat{s}; L)] \\ = & \pi(\hat{s}; L)\bar{\pi}(\hat{s}; R) - \pi(\hat{s}; R)\bar{\pi}(\hat{s}; L) \end{aligned}$$

The result then follows from (1) above. ■

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 and then make a binary report to consumers. Bias arose in this context as pure distortion—firms sometimes reporting \hat{r} when their signal was l . In this section, we extend the model to the case where firms' underlying information is a continuous rather than binary signal. We continue to assume that they make a binary report and consider the case of a monopoly firm with homogeneous consumer beliefs. Assuming a continuous signal captures the idea that media firms' must take a large amount of underlying information and summarize or filter it into a much simpler report for consumers. Note in particular that it can be seen as an approximation to the case where firms receive a large number of underlying binary signals that are either r or l and must choose one of these signals to report to consumers.

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. As before, the presence of ex-post feedback will tend to discipline this incentive and therefore to reduce the amount of equilibrium bias.

Suppose a normal firm receives a signal $s \in (-b, b)$ with $b \in (0, \infty]$ whose distribution function $G(\cdot)$ depends on the state of the world. (Here we use $b = \infty$ to denote the case in which $(-b, b) = \mathbb{R}$.) After observing this signal, the firm has the option of reporting either \hat{r} or \hat{l} . (We continue to assume that the high type always reports the true state.) We assume that $G(\cdot)$ has full support on $(-b, b)$, and that higher values of s indicate a greater likelihood that the true state is R . More precisely, we assume that

$$\frac{g(s | R)}{g(s | L)} \tag{R1}$$

is strictly increasing in s , where $g(\cdot)$ is the (continuous and differentiable) probability density function associated with $G(\cdot)$. We will consider the case where where the firm's prior θ^F is equal to $\frac{1}{2}$.

We also impose the following restrictions:

$$\lim_{s \rightarrow -b} \frac{g(s | R)}{g(s | L)} = 0 \tag{R2}$$

$$\lim_{s \rightarrow b} \frac{g(s | R)}{g(s | L)} = \infty \tag{R3}$$

$$g(0 | R) = g(0 | L) \tag{R4}$$

$$1 - G(0|R) = G(0|L) > \theta \tag{R5}$$

Restrictions (R2) and (R3) imply that as the value of s 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 on the informativeness of the firm's signal by guaranteeing that consumers in either group would rather take action R when $s > 0$ and L when $s < 0$ 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, where we assume for simplicity that the firm reports \hat{r} when it is indifferent:

Appendix Proposition 2 *There exists a cutoff $k^* \in (-b, 0]$ such that in any equilibrium the firm reports \hat{r} if and only if $s \geq k^*$. The cutoff k^* is weakly increasing in μ and weakly decreasing in θ , strictly whenever $k^* < 0$.*

Proof. The first step is to show that any equilibrium strategy must take the cutoff form. Let $C \subset (-b, b)$ be the set of signals such that the firm reports \hat{r} . It is sufficient to show that if $s' \in C$, the firm will strictly prefer to report \hat{r} after seeing any $s'' > s'$. Note that conditional on consumers receiving exogenous feedback, increasing the signal s increases the firm's posterior on the true state $\theta^F(s)$ and so strictly increases the expected gain to reporting \hat{r} rather than \hat{l} . Conditional on no feedback, increasing s does not change the expected gain to reporting \hat{r} . Thus, a firm that weakly preferred reporting \hat{r} after s' must strictly prefer reporting \hat{r} after s'' .

We now show that the cutoff k^* exists and is unique. Suppose consumers expect the firm to play a strategy with cutoff k' . Write the firm's expected gain to reporting \hat{r} rather than \hat{l} given a signal s and consumer expectations k' as:

$$\Delta(s, k') = (1 - \mu) \Delta^{nf}(k') + \mu \Delta^f(s, k').$$

The argument in the previous paragraph shows that $\Delta()$ is strictly increasing in s for a given k' . Therefore a necessary and sufficient condition for a cutoff k^* to be an equilibrium is that $\Delta(k^*, k^*) = 0$. $\Delta()$ is also strictly increasing in k' for a given s , since increasing the cutoff makes normal firms more likely to report \hat{l} and less likely to report \hat{r} (thus decreasing the posteriors on quality after the former report and increasing them after the latter). The facts that $\Delta()$ is strictly increasing in both arguments and that by (R2) and (R3) we have

$$\begin{aligned} \lim_{x \rightarrow -b} \Delta(x, x) &< 0 \\ \lim_{x \rightarrow b} \Delta(x, x) &> 0 \end{aligned}$$

mean that such a k^* exists and is unique.

To see that $k^* \leq 0$, suppose first that consumers expect a cutoff $k' = 0$ and that the firm sees a signal $s = 0$. Then (R4) implies that consumers' posteriors on quality will be the same as in the binary model with $\pi = G(0|L)$. The fact that $\theta^F(0) = \frac{1}{2}$ then implies that the firm will be indifferent about its report conditional on feedback. The gain to reporting \hat{r} conditional on no feedback, $\Delta^{nf}(0)$, will be zero if $\theta = \frac{1}{2}$ and strictly positive if $\theta > \frac{1}{2}$. Therefore $\Delta(0, 0) \geq 0$, which implies that $k^* \leq 0$, with $k^* < 0$ whenever $\theta > \frac{1}{2}$.

To see the comparative static on θ , recall that the only terms in the firm's expected payoffs that change with θ are consumers' posteriors on quality conditional on no feedback. Lemma 1 implies that increasing θ increases the posterior after \hat{r} and decreases the posterior after \hat{l} , strictly if $k^* < 0$. $\Delta(s, k')$ is therefore increasing in θ for any s and k' , which means the equilibrium k^* is decreasing in θ (strictly if $k^* < 0$).

To see the comparative statics on μ , note that at $\theta = \frac{1}{2}$, $\Delta^{nf}(0) = \Delta^f(0, 0) = 0$. Thus the equilibrium at this point is independent of μ . For the case where $k^* < 0$, it is possible to show that

(R1) and (R5) imply $\Delta^{nf}(k^*) > 0$ and $\Delta^f(k^*, k^*) < 0$. Since $\Delta(k^*, k^*) = 0$, increasing μ makes $\Delta(\cdot)$ strictly negative and so k^* must increase to restore equilibrium. ■

A.3 Allowing for a dishonest high type

In the model presented in the body of the paper, we assume that a high-type firm both knows the true state of the world *and* always reports its signal 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 future profits, which we now assume are given by the same continuation payoff function $f(\cdot)$ that determines the payoff of normal firms. 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 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. To see why, note that by definition normal types will be willing to play the strategy defined in proposition 1 given that high types are reporting honestly, since that is the assumption that is maintained throughout proposition 1. To see that in this case the high type will be willing to report honestly, observe that the only difference in the reporting incentives of the high and normal types come through the firm’s posterior on the true state. The presence of feedback therefore guarantees that the high type always has strictly more incentive than the normal type to report honestly. Since normal firms always weakly prefer to report \hat{r} given a signal of r , high types must strictly prefer to report \hat{r} in this case since then they are assured of matching the feedback. Additionally, since the normal firm either strictly prefers to report \hat{l} given a signal of l or is indifferent to its report in that case, the high type will strictly prefer to report \hat{l} . Therefore it is an equilibrium for the high type to report honestly and for the normal type to play the strategy characterized in proposition 1.

Other equilibria are also possible, however. Given the continuation payoffs we assume, the normal type always wishes to emulate the high type’s reporting strategy. If the high type is not being perfectly honest, in general the strategy defined in proposition 1 will not be an equilibrium, because 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’s strategy involves randomization given some signal, 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 \hat{r} when its signal is l , then high-type firms must be indifferent between reporting \hat{r} and reporting \hat{l} given consumers’ beliefs about the strategies of the two types. But then a small increase in the probability of the high type reporting \hat{r} will increase the incentives for the high type to do so. This in turn will lead high-type firms to move towards reporting \hat{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 \hat{r} , this reduces the incentive for them to say \hat{r} , so that behavior has a tendency to return to the equilibrium point.

To define stability formally, let $q \in \{0, 1\}$ index whether a firm is high-type (with $q = 1$ denoting a high-type firm), and let $\sigma_s(\hat{s}; q) \in [0, 1]$ be the probability that type q reports \hat{s} given a signal of s . Analogously, let $\Delta(s; q)$ be type q ’s net return to reporting \hat{r} given a signal of s . We will say that an equilibrium is *stable* if for all q and s , either $|\Delta(s; q)| > 0$ or $\Delta(s; q) = 0$ and $\partial\Delta(s; q)/\partial\sigma_s(\hat{r}; q) < 0$. That is, an equilibrium is stable if for each signal s and type q , either the

type strictly prefers its equilibrium report, or it is indifferent between reports and an increase in its probability of its reporting \hat{r} *strictly decreases* its return to doing so. This definition captures the idea that when a type’s behavior is perturbed, it ought to have an incentive to move back to the equilibrium point.

Finally, the model also permits equilibria in which both normal and high quality firms always make the same report regardless of their signal. For example, if consumers expect both normal and high types to always report \hat{r} , their beliefs about quality will be unchanged when they see \hat{r} regardless of the exogenous feedback. Seeing \hat{l} , on the other hand, is a zero probability event so we could assign consumers the belief that if this node is reached the firm is normal for sure. This means all types would strictly prefer to report \hat{r} so this would be a stable equilibrium. This is not a particularly interesting equilibrium, however, because the firm’s report would have no value to consumers. We will refer to equilibria in which both types of firms always make the same report as *degenerate* and focus on the set of non-degenerate equilibria. We also implicitly ignore the equilibrium in which the high type plays a pure “lying” strategy—i.e. always reports \hat{l} when the state is R and vice-versa—since this is equivalent to the equilibrium in proposition 1 up to a relabeling of the reports.

We now have the following result:

Appendix Proposition 3 *There exists a unique non-degenerate stable equilibrium in which the high type reports honestly, and the normal type plays the equilibrium strategy defined in proposition 1.*

Proof. We have already shown that these strategies constitute an equilibrium. To see that it is stable observe that all types except possibly the normal type who has seen a signal l strictly prefer to make the report called for in the equilibrium. We showed in proposition 1 that in this equilibrium $\partial\Delta(l; 0)/\partial\sigma_l(\hat{r}; 0) < 0$ whenever $\sigma_l(\hat{r}; 0) > 0$, so that the equilibrium is stable.

Proposition 1 established that this equilibrium is unique in the class of equilibria in which the high type reports honestly. Therefore to complete the proof we need only show that there exists no stable equilibrium in which the high type misreports with positive probability. If the high type never randomizes, it must be the case that either: (i) the high type misreports in both states with probability one, which is equivalent to a relabeling of the equilibrium in proposition 1; or (ii) the high type misreports one state with probability one and reports the other state honestly, which would be a degenerate equilibrium. Suppose, then, that for some signal s' the high type randomizes. Then we must have $\Delta(s'; 1) = 0$. But an increase in the probability of high-type firms reporting \hat{r} will lead to an increase in the incentive to report \hat{r} , i.e. that $\partial\Delta(s; 1)/\partial\sigma_s(\hat{r}; 1) > 0$, so any such equilibrium fails to meet the definition of stability. ■

B 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. 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.¹ 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.

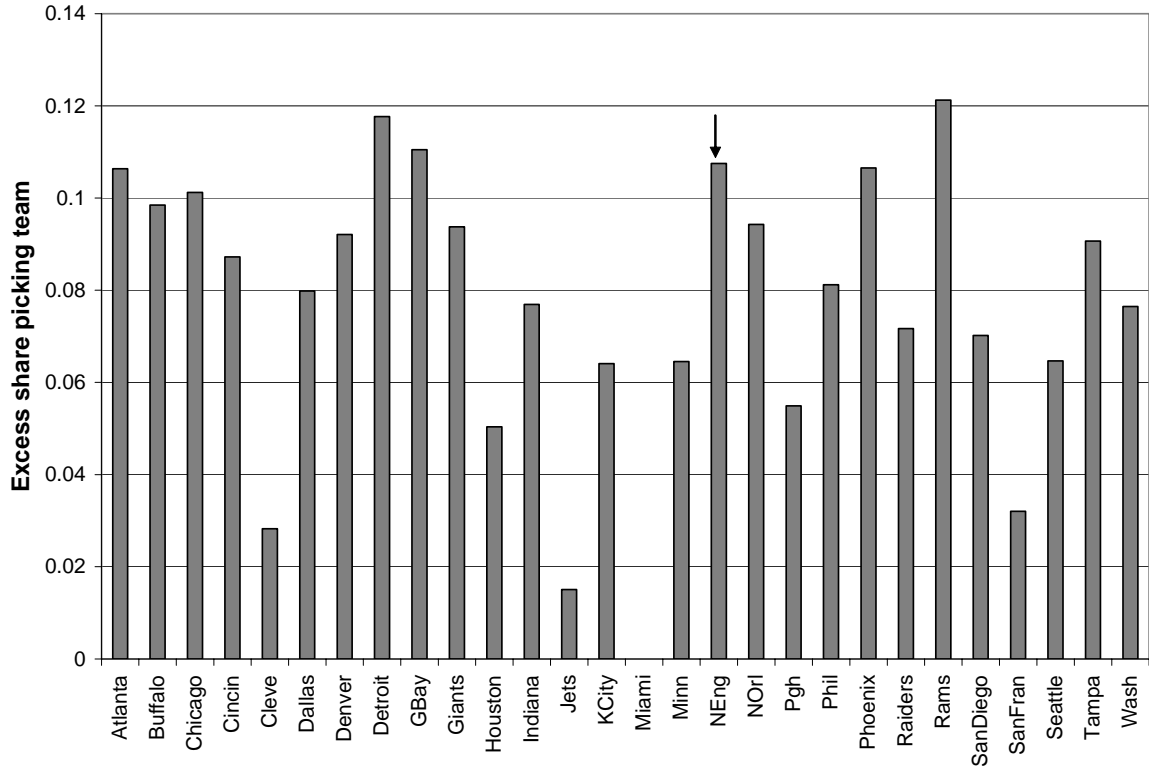
¹Sample sizes by country are as follows: Pakistan (2,043), Iran (1,501), Indonesia (1,050), Turkey (1,019), Lebanon (1,050), Morocco (1,000), Kuwait (790), Jordan (797), and Saudi Arabia (754). Other than a slight oversampling of urban households, the samples are designed to be representative of the 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		Overall CNN quality rating	
	(Mean = .33, SD = .33)		(Mean = .10, SD = .24)	
Importance of religion (Mean = .76, SD = .30)	-0.1711 (0.0132)	-0.1520 (0.0132)	-0.0418 (0.0101)	-0.0291 (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.

C 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

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

where win_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.