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 $\bar{\pi}()$ and $\pi()$ may be either PMFs or PDFs so long as any mass points of $\bar{\pi}()$ are also mass points of $\pi()$.

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 λ 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:

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

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

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

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:

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

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

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

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

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:

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

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

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

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 $\Delta()$ 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 $f()$ 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.

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

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

C Sports picking by *Boston Globe* columnists, 1983-1994

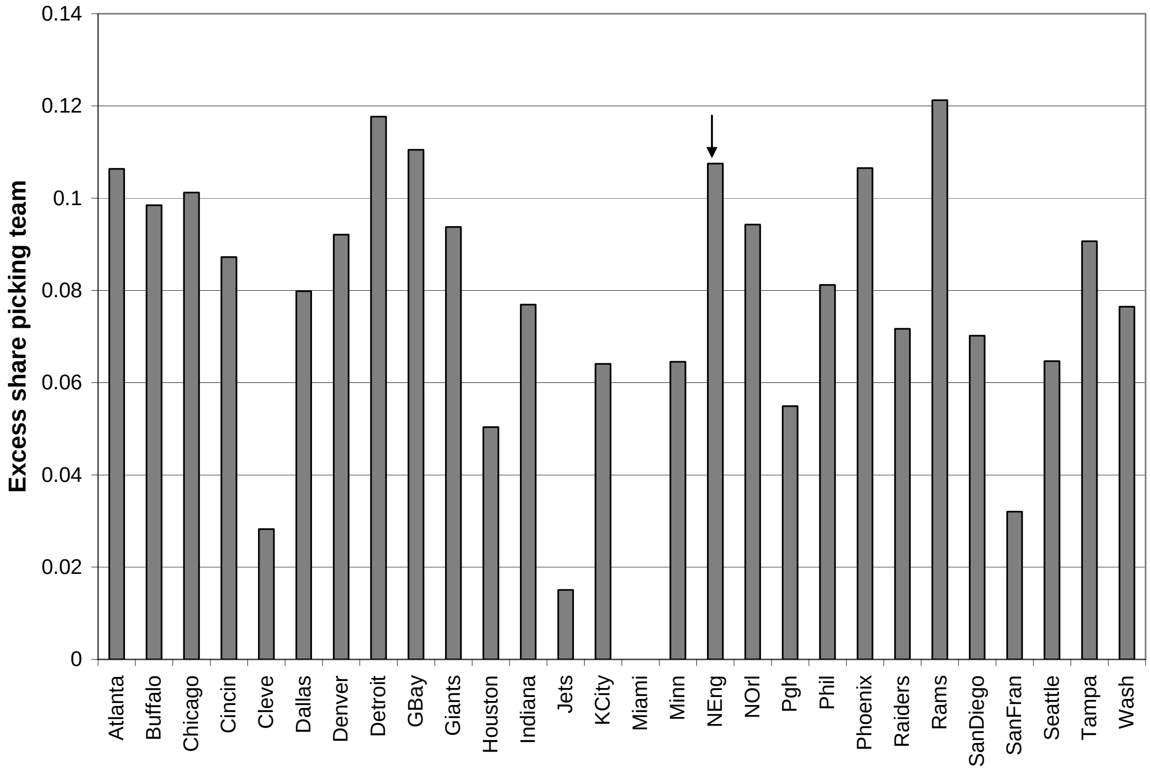

Notes: Data from Avery and Chevalier (1999). Dataset contains information on the picks of *Boston Globe* sports columnists for NFL games in the 1984-1994 seasons, as well as the outcome of the game and the opening betting line. The bar for team i represents the estimated coefficient $\hat{\delta}_i$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

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

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.