

Media Bias

Sendhil Mullainathan and Andrei Shleifer*

June 18, 2002

Abstract

There are two different types of media bias. One bias, which we refer to as ideology, reflects a newspaper's private desire to affect reader opinions in a particular direction. The second bias, which we refer to as spin, reflects the newspaper's attempt to simply create a memorable story. We show that these two biases operate very differently. Whereas competition can eliminate the effect of ideological bias, it actually exaggerates the incentive to spin stories.

*We are extremely grateful to Ed Glaeser and Simon Johnson for comments.

1 Introduction

The descriptive accounts of the media industry (Mencken 1975, Goldberg 2002), as well as journalism textbooks (Jamieson and Campbell 2002) clearly recognize that media outlets do not just report plain facts. One commonly noted bias in reporting is ideological. During President Clinton's impeachment trial, the right-leaning media supported impeachment, arguing that even a small lie in a civil deposition constitutes a major crime. The left-leaning media, in contrast, saw the trial as a witch-hunt, treating the very same violations condemned by the conservative press as trivial. Another, perhaps less commonly noted, bias is not ideological at all, but rather results from the need to tell a memorable story. Thus both the conservative and the liberal media condemned Wen Ho Lee as a Chinese spy, without entertaining the possibility that he was scape-goated by dishonest law enforcement officials. The media were surely biased, and most likely wrong.

In this paper, we draw a sharp analytical distinction between these two kinds of biased reporting. We call the traditional left or right bias ideology, and the less traditional bias - one based on the need to tell a story - spin. Essentially, we see ideology as coming from the preferences of the suppliers of news, and spin coming from the consumer's demand for how information is presented. We then examine the consequences of both ideology and spin for the accuracy of news reporting. Our principal interest is how competition affects accuracy. We show that even with competition, the truth may not always come out, especially when spin is important.

We consider the case of private media outlets maximizing the welfare of their owners, which depends on ideology and profits. We begin with the case of a monopoly news supplier. We establish that, with Bayesian readers, only the ideology of the paper (or of any media outlet) is a source of bias in reporting. But we also consider the case of readers who are categorical thinkers - have limited memory, and can only remember the "bottom line." With categorical readers, it pays the monopolist newspaper to spin the news in order to make them memorable even in the absence of any ideology. This spin - the need to tell the story - is also a source of reporting bias.

The more interesting issues arise when we consider competition among the media outlets. Such competition is an important argument for free press: despite the ideological biases of individual newspapers, the truth comes out through competition. We show that, with Bayesian readers, this is indeed the case: competition undoes the biases from ideology. With readers who are categorical thinkers, however, the consequences of competition are more complex. We show that, in the absence of ideology, competition actually reinforces the adverse effects of spin on accuracy. Not only do the media outlets bias news reporting, but the stories tend to reinforce each other as the competitors strive to be remembered. When the first paper who uncovers the story is ideological, the possibility of spin exacerbates the influence of ideology. The reporting becomes more extreme when the discoverer knows that the followers will reinforce the story.

These results have significant implications for the accuracy in media. They explain, in particular, how the media in the aggregate are likely to get to the bottom of a news story with significant ideological dimension. Ideological diversity serves as a safeguard against spin. Our results are thus consistent with Richard Posner's (1999) highly favorable assessment of the press in the coverage of the Clinton affair. Our results also show why media bias is most severe in the cases where no or little ideological diversity bears on the story, such as the investigation of Wen Ho Lee. In this case, the bias comes from spin, and spin causes the followers to pile on. Competition among media outlets is not a solution to the problem of spin - indeed, it makes the problem worse. Our paper thus makes the case for extreme ideological diversity in the media - in such diversity lies the best hope against spin.

2 Model Setup

There are two players: the newspaper and the reader of the paper. The newspaper is either good or bad. A good newspaper has access to some information that the reader wants. A bad newspaper, on the other hand, has no useful information.¹ The reader

¹A bad newspaper cannot manufacture stories to appear good. We use these two types merely as a way to model the incentives of the good newspaper to create memorable stories. If no stories are remembered

uses the newspaper to improve his estimate of some value t . We begin by assuming that t is a single real number. This value could be many things: for example, the quality of a politician or the financial health of a firm. The common prior is $t \sim N(0, \sigma^2)$, i.e. t is normally distributed with variance σ^2 .

There are four periods. In the first period, the reader receives a signal r about the truth. This signal reflects the reader's prejudices about the topic before he reads the news-paper. This prejudice could be shaped by basic beliefs about how the world operates. For example, in the Wen Ho Lee case readers may have started with the presumption that the Chinese are likely to be spying on the US. Alternatively, these priors may themselves be the result of other news-paper stories.² The signal r is a noisy estimate of the truth, so

$$r = t + e_r$$

where e_r is distributed normally $N(0, \sigma_r^2)$. The reader uses this signal to update his beliefs.

The newspaper may also receive a signal. Good newspapers receive a signal for sure and bad newspapers receive no signal. The signal they receive will be denoted n and it equals the truth plus some noise: $n = t + e_n$, where e_n is normally distributed with variance σ_n^2 .

In the second period, the newspaper decides what to report. If it is a bad news-paper it has no news to report. A good newspaper, in contrast, has news n to report which it may modify before reporting it. Let s denote this modification. Thus the paper reports $\nu = n + s$. Modification is costly for the newspaper since it requires costly story-telling. Let $c(n - \nu)$ denote this cost function. We assume that $c(0) = 0$ and that $sign(c'(x)) = sign(x)$ and $c''(x) = 0$ and that $\lim_{x \rightarrow \infty} c(x) = \infty$. These assumptions mean that manipulation, in either direction, is costly to the newspaper.

In the third period, the reader reads the story. He then combines what he reads with his priors to form an assessment of t . Define this belief to be \hat{t} .

Finally, in the fourth period, the true value of t is revealed. The individual compares

the newspaper is perceived to be bad.

²We formally study this second case when we consider multiple newspapers.

this signal to *what he remembers* was in the newspaper. If he remembers no signal, he presumes it was a bad newspaper. If he remembers a signal, he presumes the paper is a good one.

2.1 Reader Psychology

The model turns on how we assume readers form beliefs and remember information. At one extreme, we might suppose people form Bayesian beliefs and have perfect recall. In this case, after reading the news ν , the individual updates and forms beliefs:

$$\hat{t} = r + k_1(\nu - r)$$

where k_1 is a constant that equals $\frac{\sigma_n^2}{\sigma_r^2 + \sigma_n^2}$.

An alternative perspective is to assume that readers think in broad categories rather than with precise fine-tuned priors of a Bayesian (Mullainathan 2000). The reader reads the story, combines it with what he knows and forms a general impression. What he carries away is this general impression or coarse summary of all his information. In our example, the reader may only maintain one of two beliefs: the truth is positive or the truth is negative. Thus Clinton is either a persecuted modern man or a villain. Wen Ho Lee is a spy or a victim. Specifically, define t_+ and t_- to be the positive and negative categories and assume these are symmetric in that $t_+ = -t_-$.³ Denote the categorical thinker's beliefs to be \hat{t}_c . We assume that he believes in the category closest to what the Bayesian would believe. In this context, that means that he believes t_+ when the Bayesian would have believed a positive t and he will choose t_- if the Bayesian would have expected a negative t . Formally:

$$\hat{t}_c = \begin{cases} t_+ & \text{if } \hat{t} > 0 \\ t_- & \text{if } \hat{t} \leq 0 \end{cases} \quad (1)$$

For example, suppose the reader has several pieces of evidence most of which suggest Wen Ho Lee is a spy but some of which suggest the government scape-goated him. A categorical thinker will walk away having placed Wen Ho Lee in the spy category, essentially collapsing his information into this simple summary.

³Allowing more categories does not change the results.

The second assumption about categorical thinking is that recall is not perfect. Information consistent with a category is more likely to be remembered than information inconsistent with a category. The probability of recalling a news-story ν is $r(\nu, \hat{t}_c)$. We assume that recall probabilities are higher for news consistent with the reader's perceived category. That is a story is more likely to be remembered if $sign(\nu) = sign(\hat{t}_c)$. So, in the above example, once the categorical thinker places Wen Ho Lee in the spy category, he is less likely to remember the stories that suggested the government scapegoated him.

2.2 Newspaper Payoffs

A newspaper's payoff depends on several factors. First it depends on whether the newspaper is perceived to be good since presumably good newspapers sell better. The paper receives a payoff π if it is thought to be good and zero otherwise. Second, the newspaper may have an ideology τ which is either 0, +1 or -1 . A newspaper with an ideology of +1 wants readers to have a positive belief whereas a newspaper with an ideology of -1 wants readers to have a negative belief. We assume the newspaper suffers a cost β if it reports news that is inconsistent with its ideology. In this sense β measures the intensity of the newspaper's ideology. We further assume that this ideology is private knowledge and readers form inferences ignoring it.⁴

When readers are Bayesian, there is no confusion of good or bad newspapers since all stories are remembered. A good news-paper no matter what it reports will have its story remembered and be recognized as good. And by assumption a bad newspaper can never appear good. A good newspaper therefore receives a payoff of:

$$\pi - \beta(sign(\tau) \neq sign(\hat{t})) - c(n - \nu)$$

and a bad newspaper receives a payoff of:

$$-\beta(sign(\tau) \neq sign(\hat{t})) - c(n - \nu)$$

⁴This assumption is merely to simplify the calculations. At the opposite extreme readers would know and recognize these ideologies and would debias stories before incorporating them into beliefs. Media biases would have no effect on beliefs at all. Our assumption is a simple version of the more realistic case where readers partly but not fully recognize the bias.

When readers are categorical, however, good newspapers can be confused with bad ones. If a good newspaper’s story is inconsistent with priors, it will be forgotten and the newspaper will be thought of as a bad one. Readers presume that since they remember no story the paper must have been bad and reported nothing.⁵ A good newspaper’s expected payoff therefore depends on the probability of recall in this case:

$$\pi * r(\nu, \hat{t}_c) - \beta(\text{sign}(\tau) \neq \text{sign}(\hat{t})) - c(n - \nu) \quad (2)$$

A bad newspaper is never confused for a good one. Even when its stories are remembered they are seen as worthless. Consequently its payoffs are as before:

$$-\beta(\text{sign}(\tau) \neq \text{sign}(\hat{t})) - c(n - \nu)$$

With these payoffs, it is clear why categorical thinking provides a non-ideological reason for newspapers to manipulate stories. Newspapers want their stories to be memorable because if the reader does not remember any story he presumes the newspaper is a bad one and provides no information. To be memorable, the newspaper therefore has incentives to alter their stories so that they better fit the reader’s category.⁶ In the Wen Ho Lee case, unless the newspaper can change the reader’s mind, it has an incentive to report a story consistent with the reader’s categorization of Wen Ho Lee as a spy.

It is worth noting that this specific model is only one way to operationalize the element of reader psychology we want. What is critical to us is that the reader values stories which match their beliefs except when these stories are so large as to change beliefs. In our model this preference arises indirectly because of the differential memorability of stories, but there are other ways to model it. An alternative is that readers suffer from confirmatory bias: they dismiss stories that are inconsistent with what they believe. In fact, our model could be reinterpreted in this light. Suppose that readers find it easy to dismiss moderately inconsistent news but not grossly inconsistent news.

⁵One could get similar results by making the less drastic assumption that readers “rationally” infer that the paper might have been good and they simply forgot the story. Even in this case a forgotten story will increase the reader’s perceived likelihood that the paper is bad and thereby hurt its payoff.

⁶An alternative interpretation is that reporters want to write memorable stories since that will raise their profile within the paper.

In this case, we can reinterpret the recall function as a dismissal function. When the reader dismisses news as fallacious, he presumes the paper is bad which is isomorphic to presuming that a newspaper whose story has been forgotten is bad. Similarly, the fact that very big stories cannot be dismissed is equivalent to the fact that categorical thinkers change their mind to big enough news and hence remember the story.⁷ This re-interpretation produces the same results as ours. Once again newspapers with small amounts of inconsistent news would bias stories so as not to be dismissed. But for large amounts of news, they do not fear dismissal and would accurately report it.

We have chosen the model with recall and categorization for two reasons. First, it better matches descriptive accounts of the media, which suggest that newspapers are very interested in writing memorable stories. Second, it provides a better framework for thinking about what kinds of spin will occur. Newspapers will attempt to spin stories so that they better fit a particular category. Nevertheless, beyond interpretation, the results we provide below would be the same with the alternative model as well.

3 Results

3.1 Two Types of Bias

We begin with the case where the reader is Bayesian. Here, the newspaper has only one reason to try and distort the news: if it has an ideology it wants to pursue. If it has no ideology, it gains nothing from manipulating the news. The first proposition formalizes this idea.⁸

Proposition 1 *If the newspaper has no ideology ($\tau = 0$), then it does not manipulate news: $\nu = n$. If it has an ideology it wants to manipulate news in the direction of its ideology so that $\text{sign}(\nu - n) = \text{sign}(\tau)$. The magnitude of manipulation is increasing in the strength of the ideology (β) and decreasing in the size of the news ($|n|$) and in*

⁷Specifically one would replace the recall function in equation 2 with a dismissal function that is non-linear. The probability of dismissal would be increasing as the news is further from priors until it gets large enough at which point the probability of dismissal would diminish.

⁸All proofs are in the appendix.

the cost of manipulation ($c(\cdot)$).

This proposition suggests one source of bias: the political ideology of the newspaper. But when agents are categorical thinkers, another bias may arise. Newspapers may manipulate news not because of their political ideology but because they cater to the reader's interests. The paper wants to have its story remembered and, therefore, may manipulate news stories to make it more memorable. The following proposition formalizes this notion of *spin*.

Proposition 2 *When the agent is a categorical thinker the newspaper biases the news even when it has no ideology, when $\tau = 0$. When the size of the news ($|n|$) is small enough it biases the news towards priors: $|r - \nu| < |r - n|$. When $|n|$ is large enough, however, it reports the truth.*

A simple example illustrates this proposition. Suppose the reader believes 2 and the newspaper receives a signal -1 . Suppose the newspaper reports the truth in this case. Since -1 is not enough to switch beliefs, the categorical reader continues to believe t_+ . But since the story that the newspaper reported is negative, it contradicts the final category and therefore has a good chance of being forgotten. The newspaper therefore has an incentive to alter the story to make it positive and thereby more likely to be remembered. In other words, if the news is not enough to change the reader's mind, the paper has an incentive to spin it to match the readers prior convictions.

But if the signal is large enough, the newspaper may want to go against beliefs. In the above example, if the newspaper sees a signal of -3 it knows that by reporting the truth it can change the reader's mind. Once the reader's mind is changed the news story and category match and therefore the story will be remembered. In this case the newspaper no longer has the incentive to spin the story to increase memorability.

These two propositions emphasize that underlying the notion of media bias there are really two distinct phenomena. One is driven by the political ideology of the newspaper and arises from an attempt to convince the public of a particular viewpoint. The other bias is driven by an attempt to appeal to the reader in our case by

creating a memorable story. This kind of “spin” is merely an attempt by the newspaper to cater to the reader’s psychological make-up.

3.2 Competition

An economist might reason that competition between media outlets should eliminate media bias and lead to the revelation of the truth. To investigate this idea, we consider the case where there are multiple newspapers instead of just one paper publishing news. Suppose there are N of them. Each newspaper releases news in a sequential manner with newspaper indexed 1 releasing first, followed by the one indexed 2 and so on. Each newspaper receives an independent signal about the truth and reports on this signal.

The reader reads all of the newspapers and aggregates the stories.⁹ This assumption is intended to capture the idea of a reader who is exposed to multiple media outlets either because he generally uses multiple sources or because he searches them out because he is interested in a particular story. Assuming exposure to multiple sources is a necessary feature of any model of competition requires because that is the source of competition. If readers are in fact isolated and read only one source, then one would simply apply the results of the previous section and assume it is a case of local monopoly. Thus in what follows when we speak of the number of newspapers we mean the number that the reader is actually exposed to. Finally, we assume that newspaper payoffs are independent. That is a newspaper receives the same payoff π for being perceived as good independent of the number of other newspapers. This means that the effects of competition we discuss arise purely from information aggregation rather than from industrial structure.

Now suppose that agents are Bayesian. If newspapers have different enough ideologies, the average news story is an unbiased estimate of the truth. Each newspaper may bias news towards its ideology but these differences wash out in the aggregate.

⁹When the reader is categorical, we assume that he forms his category after reading all the stories rather than sequentially.

Proposition 3 *Suppose agents are Bayesian and the number of newspapers having a positive ideology is the same as the number of newspapers having a negative ideology. Then in aggregate there is no bias in beliefs. That is the reader's beliefs after reading all the newspaper stories (ν) are the same as if they could see the true news (n) themselves.*

So if there are enough left-wing and enough right-wing papers, the truth about Clinton will emerge. Each side may exaggerate the story in its preferred direction but these exaggerations will cancel. The key feature is the canceling ideologies between newspapers.

But while competition eliminates the bias caused by ideology, it does little to affect spin, the bias caused by categorization. To understand why this is the case, note that each newspaper reporting on a story still feels the need to go along with reader's priors. But now that effect is reinforced because each newspaper knows that future newspapers will also cater to the reader's beliefs at the time. Since each of the prior newspapers has biased the news, the reader is now stronger in his beliefs. Each newspaper, therefore, feels an even greater pressure to bias toward priors.

Proposition 4 *Suppose agents are categorical and newspapers have no ideology. Then even as $N \rightarrow \infty$, newspapers continue to bias news towards priors as long as each individual news-paper is receiving small enough pieces of news. In this case, agents continue to believe in false facts even with and arbitrarily large numbers of news-stories and total information. Moreover, the newspaper has greater incentives to spin the story with competition than without.*

This Proposition applies to stories where newspapers have no ideology. For example, newspapers may have had little ideological interest in characterizing Wen Ho Lee as a spy. In this case we would interpret what happened as a piling on by the newspapers. If readers began with a moderate prior that the Chinese are spying on the US, then the first newspaper had an incentive to spin the story in this direction. This reinforces the prior and increases the incentive of the next newspaper to spin their story in the same direction further reinforcing the belief. Each newspaper therefore piles on to this particular spin, making it more likely that successive newspapers follow suit. Unless

a newspaper receives some extremely compelling evidence that Wen Ho Lee was not a spy, the piling on will result in a very biased outcome.

The proposition also applies to some cases where newspapers have an ideology if the reporters themselves do not have one and the editors exercise little control. Most notably on smaller political issues, even if newspaper editors have particular ideologies, editors will not find it worthwhile to control exactly how each story is written.

To summarize, competition does not necessarily remove bias. If the bias is due to the ideology of newspapers, then increasing the diversity of opinions helps get the truth out. But if the bias is one of spin, then competition only exaggerates it and keeps aggregate information away from the truth.

3.3 Interaction between Spin and Ideology

In a competitive environment, spin and ideology interact in interesting ways. First, when agents are categorical, the bias from ideology can be *exaggerated* by spin. Consider the case where the first newspaper reporting the story has an ideology and other newspapers do not. Then the first newspaper knows that by manipulating its story, not only is it affecting beliefs today, it is also affecting future reporting. The following proposition formalizes this.

Proposition 5 *Suppose that newspaper 1 has an ideology $\tau \neq 0$ and that later newspapers do not. If agents are categorical thinkers and each newspaper receives only a small bit of news, then the aggregate bias is higher when there are more newspapers.*

This result emphasizes the case where a single ideological newspaper can bias stories in a particular way and exploit later newspapers' desire to be consistent with priors. Proposition 5 has one important implication, namely the possibility for politicians themselves to control spin. If a politician can get out his story to a friendly newspaper first, and thus create a particular image of himself or his proposals, future newspapers will be influenced by the priors created by this first story. As a consequence, the image will stick, competition notwithstanding.

On the other hand, when the follower newspapers have their own ideologies, they do not passively accept the early newspapers' bias. The following proposition shows that when in fact newspapers' ideologies are very different, then the incentives for newspapers to spin stories falls.

Proposition 6 *Suppose the number of newspapers with a positive ideology equals the number with a negative ideology and that agents are categorical thinkers. As the strength of ideologies (β) increases, newspapers' incentives to spin stories decreases and the total amount of bias in reporting decreases.*

Proposition 6 is perhaps the crucial result of the paper. It shows that newspaper ideology, far from being a problem to be lamented, is in fact the primary force toward accuracy in media. Without ideological diversity, spin causes newspapers to pile on, and sacrifice accuracy in order to spin stories. When newspapers have different ideologies, however, the competitive pressures to follow the leader are countered by the pressures of ideology. Ideological diversity is an antidote to spin.

4 Conclusion

One of the central questions in discussing the media is when can careful consumers of competitive media sources come away with an accurate picture of facts. In this paper, we try to answer this question by focusing on two sources of media bias: ideology and spin. We show that, in the absence of spin, competition undoes the biases from ideology. However, in the absence of ideology, competition only promotes the biases from spin.

These results help shed light on the kinds of issues where we expect accuracy in the media, at least in the aggregate. For news that have significant right-left ideological dimensions, such as presidential elections or race issues, competitive media are likely to produce accurate reporting, on average. On the other hand, for news that do not have such a dimension, such as cases involving law enforcement or many aspects of foreign policy, spin rather than ideological diversity is likely to dominate competitive

reporting. For such stories, competition does not eliminate the media bias, and there is no good reason to expect accuracy in media.

References

Goldberg, Bernard, 2002 *Bias: a CBS Insider Exposes How the Media Distort the News*, Washington, DC: Regnery Pub.

Jamieson, Kathleen Hall and Karlyn Kohrs Campbell, 2002, *The Interplay of Influence*, Belmont, CA: Wadsworth/Thompson Learning.

Mencken, H. L. 1975, *A Gang of Pecksniffs*, New Rochelle, NY: Arlington House Publishers

Mullainathan, Sendhil, 2000, "Thinking Through categories," MIT mimeo.

Posner, Richard 1999, *An Affair of State: The Investigation, Impeachment, and Trial of President Clinton*, Cambridge, Mass. : Harvard University Press

Appendix

Proof of Proposition 1

Proof: Since agents are Bayesian, they remember everything they read. Consequently a good news-paper with no ideology receives a payoff of:

$$\pi - c(n - \nu)$$

Since this is strictly decreasing in $|n - \nu|$, there is no incentive to manipulate. Hence $\nu = n$.

When the newspaper has an ideology it receives a payoff:

$$\pi - \beta(\text{sign}(\hat{t}) \neq \text{sign}(\tau)) - c(n - \nu)$$

If $\text{sign}(\hat{t}) \neq \text{sign}(\tau)$, it pays a cost of β for reporting the truth. Let ν' be the news required to make $\text{sign}(\hat{t}) = \text{sign}(\tau)$. If β is larger than $c(n - \nu')$, the newspaper biases news towards its ideology. Thus the incentive to manipulate is clearly decreasing in the size of the news n . ■

Proof of Proposition 2.

Proof: A good newspaper worries about being remembered. When it has no ideology, its payoff is

$$\pi * (r(\nu, \hat{t}_c) - c(n - \nu))$$

Suppose the news is inconsistent with the readers priors. Specifically the sign of the news contradicts the sign of the priors. The cost of manipulating the news so that its sign matches the priors is approximately $c(n)$ since the newspaper will have to introduce a bias of n to change its sign. So as the size of the news ($|n|$) decreases, the newspaper can change its sign with arbitrarily low cost. But the benefits of having it match beliefs are discrete and fixed since information consistent with a category is always better remembered than information inconsistent with a category. As the story gets smaller and smaller, it will not affect beliefs so that belief matches priors: $sign(\hat{t}_c) = sign(r)$. Consequently, for small news stories the newspaper cannot hope to change priors. Instead it is more profitable to manipulate the story to fit these prior beliefs. So if $sign(n) \neq sign(r)$, small stories are biased so that their sign matches the sign of the prior beliefs.

As news gets very large, however, the costs of changing sign are growing while the benefits remain fixed. Consequently the newspaper does not find it profitable to manipulate large stories and reports the truth. ■

Proof of Proposition 3.

Proof: Let b_+ be the bias of the newspapers with a positive ideology and b_- be the bias of the newspapers with a negative ideology. By homogeneity, the extent of bias depends only on the ideology of the newspapers. By symmetry, $b_+ = -b_-$. Consequently, when the reader add together all the signals, the biases will cancel exactly. ■

Proof of Proposition 4.

Proof: Suppose each newspaper receives small enough news. Consider the last newspaper. The situation faced by this paper is exactly the same

as that of the one news-paper case. By proposition 2 we know that it will bias the news towards the agents beliefs before reading it. Proceeding inductively, each news-paper has incentives to bias towards the reader's priors since it knows that the newspaper which follows biases the story in the same direction. Since each story is biased, the truth does not emerge as the number of stories gets large.

Finally, note that as the number of newspapers increases, each news-paper has a diminishing effect on the reader's beliefs. Agents are getting stronger and stronger in their beliefs as the news stories accumulate. Each news-paper then has less individual effect in changing the reader's mind. This means that the case where the news-paper can report the truth and change beliefs is less and less likely. It therefore has a greater incentive (than in the single news-paper case) to bias the news toward priors. ■

Proof of Proposition 5.

Proof: Suppose newspaper 1 reports ν_1 . The situation after this report is exactly the situation in Proposition 4. Each of these newspapers will therefore bias stories towards the sum of prior beliefs and ν_1 . The incentive of the first newspaper to exaggerate are therefore are the same as in Proposition 1 since whatever it reports will be reinforced and the readers final perceptions will be determined by what it reports. Consequently if ideology is strong enough the newspaper will bias towards its ideology. And later newspapers will spin their stories in this direction. Thus the total amount of stories will be even more biased than in the single newspaper case. ■

Proof of Proposition 6.

Proof: Consider the incentive of the last news-paper. It faces an incentive to spin the story as well as incentive to bias for ideological reasons. The incentive to spin depends on the priors of the reader after having read all previous newspaper. Ideological bias, however, is independent of priors. Consequently, as β increases, the news-paper is less likely to spin, indepen-

dent of what previous newspapers have done. Proceeding inductively, as β increases, each prior newspaper is less likely to spin stories for two reasons. First, they know that the newspapers which follow are less likely to respond to the priors created by the spin. Second, they themselves have an ideology. Consequently, the aggregate amount of spin has diminished.

Now since roughly half the newspapers are biased in one way and half are biased in the other, the bias induced by ideology cancel as in Proposition 3. Consequently, total bias decreases. ■