

Media Bias in the Marketplace: Theory

Matthew Gentzkow*
Chicago Booth and NBER

Jesse M. Shapiro
Chicago Booth and NBER

Daniel F. Stone
Bowdoin College

January 2014

Abstract

We review the theoretical literature on market determinants of media bias. We present a theoretical framework that organizes many key themes in the literature, and discuss substantive lessons.

*In preparation for the *Handbook of Media Economics*, Simon Anderson, David Strömberg and Joel Waldfogel, eds. E-mail: gentzkow@ChicagoBooth.edu, jesse.shapiro@ChicagoBooth.edu, dstone@bowdoin.edu.

1 Introduction

Do media markets provide voters with accurate, timely information about politics? Survey evidence suggests many people are skeptical. Trust in media is low and falling (Ladd 2012). People on each side of the political spectrum believe news reporting is biased in favor of the other (Goldberg 2001; Alterman 2003). Pew (2011) reports that 77 percent of survey respondents in the US say news stories “tend to favor one side” and 63 percent of respondents agree news organizations are “politically biased in their reporting.”

Empirical evidence discussed in detail in the next chapter shows that different media outlets indeed select, discuss, and present facts differently, and that they do so in ways that tend to systematically favor one side of the political spectrum or the other. These differences can have large effects on voter behavior, and thus political outcomes. Concerns about bias have been a central driver of media regulation (Brocas et al. 2011). Understanding the economic forces that determine media content is thus of first-order importance.

In this chapter, we survey the theoretical literature on the market forces that determine equilibrium bias.¹ We begin by defining bias formally and discussing its normative implications. We then introduce a simple model that we use to organize the literature on the determinants of bias, focusing first on supply-side forces such as political preferences of media owners, and then turning to demand-side forces working through consumer beliefs and preferences.

2 What is bias?

Gentzkow and Shapiro (2006) discuss the following example. On December 2, 2003, American troops fought a battle in the Iraqi city of Samarra. Fox News began its coverage of 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.

¹Prat and Strömberg (2011), Andina-díaz (2011) and Sobbrío (2013) provide related surveys. See Zaller (1999), Sutter (2001) and Hamilton (2004) for early, less formal contributions that also cover some of the key related work from other fields (communications, sociology and political science). An important related early work from economics is Steiner (1952).

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.

And the English-language Web site 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.

These three news outlets base their stories on the same events, but manage to convey strikingly different impressions of what actually transpired. It is such systematic differences in the mapping from facts to news reports — that is, differences which tend to sway naive readers to the right or left on political issues — that we will call bias.²

Suppose that there is an unobserved state of the world $\theta \in \{L, R\}$, whose values we associate with outcomes favorable to the left and right sides of a one-dimensional political spectrum, respectively. Define the raw facts gathered by a news outlet to be a (possibly high dimensional) random variable $s \in S$ whose distribution depends on θ , and define a news report by $n \in N$. A *reporting strategy* σ is a possibly stochastic mapping from S to N .

We define bias as a partial order on reporting strategies. We say a strategy σ is *biased to the right (left) of* σ' (denoted $\sigma \succsim \sigma'$) if, loosely speaking, a consumer who believed a firm’s strategy was σ' would tend to shift her beliefs to the right (left) if the firm deviated to σ .

More precisely, let $\mu(n|\sigma)$ be a Bayesian consumer’s posterior probability that the state is R when she observes n from a firm believed to be playing strategy σ . Let $\lambda(\tilde{\sigma}|\sigma)$ be the distribution of μ when a consumer believes a firm is playing σ and it actually plays $\tilde{\sigma}$. Say two reporting strategies σ and σ' are *consistent* if for each of them n has the same support (i.e., the set of n that can be reported is the same) and they preserve the relative meaning of reports (i.e., $\mu(n|\sigma) > \mu(n'|\sigma) \iff \mu(n|\sigma') > \mu(n'|\sigma')$ for any $n, n' \in N$). We say σ is biased to the right (left) of σ' if

²This is equivalent to the definition of “slant” introduced in Gentzkow and Shapiro (2010).

(i) σ and σ' are consistent and (ii) $\lambda(\sigma|\sigma')$ first-order stochastically dominates (is dominated by) $\lambda(\sigma'|\sigma')$. Note that this definition is symmetric in the sense that σ is biased to the right of σ' if and only if σ' is biased to the left of σ .

This definition nests the leading forms of bias modeled in the literature, which differ in the definitions of S and N , and the constraints placed on the set of allowable σ . We divide these into two categories.³

The first category is outright distortion, which we model as a case where $N = S$. In this special case, we can think of reports as objective statements about the value of s , and it is meaningful to talk about a “truthful” strategy (one where $\sigma(s) = s$), and the “amount” of bias (some integral of the distance between $\sigma(s)$ and s). Note that this is a version of cheap talk, as in Crawford and Sobel (1982). Restricting attention to strategies consistent with the truthful strategy, our definition implies that σ is biased to the right of σ' if, at every s , it shifts probability toward false reports relatively favorable to R .

This notion of distortion also captures unnecessary omission of information. We model this as a case where one of the elements of S is a null signal that provides no new information about the state. For example, suppose that S consists of either *scandal_R* (“a scandal in the R party”), *scandal_L* (“a scandal in the L party”), or \emptyset (“no information”). Suppose further that strategies are constrained so either $n = s$ or $n = \emptyset$ (i.e., firms can hide information about scandals but cannot manufacture it). Note that this is related to disclosure games as in Milgrom and Roberts (1986). Then σ is biased to the right of σ' if it is relatively more likely to hide R scandals and relatively less likely to hide L scandals.

In the example above, if Fox News were to report that there were no civilian corpses at Samarra’s hospital, or Al Jazeera were to report definitively that US troops killed no Iraqi insurgents, we could conclude that at least one of the reports must have involved outright distortion. Papers modeling media bias as distortion include Mullainathan and Shleifer (2005), Baron (2006), Besley and Prat (2006), Gentzkow and Shapiro (2006), and Bernhardt et al. (2008).

The second category of bias is filtering of information, which we model as a case where N is coarser or lower-dimensional than S . This can capture selection, where s is a high-dimensional

³Prat and Strömberg (2011) divide bias into four types: facts bias, issue bias, framing bias, and ideological stand bias. Loosely, facts bias and framing bias are versions of what we call distortion, while issue bias and ideological stand bias are versions of what we call filtering.

vector and news outlets are constrained to choose one or a subset of the elements of the vector to report in n . It can also take the form of coarse summaries, where s is a continuous variable and reports n take on a finite number of values, such as when the report is an editorial endorsement of one candidate or the other. In the case of selection, our definition implies that σ is biased to the right of σ' if whenever σ selects an element s_i and σ' selects an element $s_j \neq s_i$, s_i is a fact more favorable to R than s_j . In the case of coarse summaries, σ is biased to the right of σ' if σ shifts probability mass from summary reports relatively favorable to L to reports relatively favorable to R . In the example, we can see selection at work, as Fox News is the only outlet to mention that eight Iraqis were captured, the *New York Times* is the only one that reported the angry reaction of Iraqis, and Al Jazeera is the only one that reported on the corpses at the hospital. We can also see coarsening in the reporting of the number of Iraqi insurgents killed, which Fox says was “at least 54” and the *New York Times* says was “as many as 54.” Papers modeling bias as selection include Strömberg (2004). Papers modeling bias as endorsements or other coarse summaries include Chan and Suen (2008) and Duggan and Martinelli (2011).

We are not aware of any systematic empirical evidence on the relative prevalence of these forms of bias. But it seems clear that filtering, in the form of both selection and summary, plays a large role in the way media bias occurs in practice. Downs (1957), for one, saw selection as ubiquitous. He writes, “All reporting is biased because the reporter must select only some of the extant facts to pass on to his audience.... Because evaluation begins with emphasis upon—i.e., selection of—certain data in contrast to others, all such selection is evaluative to some extent. In short, there is no such thing as purely objective reporting of any situation or events” (207, 212).

3 Bias and welfare

Under our definition, bias is only defined as a relative concept: one outlet can be to the right of another, but only in special cases is it meaningful to talk about “unbiased” or “objective” reporting, about “more” or “less” bias, or about whether a given outlet is left or right of “center.” There is, in general, no inherent link between bias and welfare.

The benchmark for evaluating information content of signals is the Blackwell (1951) ordering. A reporting strategy σ is Blackwell more informative than σ' if σ' can be produced by combining

σ with random noise; in other words, if the distribution of σ' given σ does not depend on s . In this case, any decision maker who observes σ will have weakly higher expected utility than a decision maker who observes σ' , regardless of preferences. Conversely, if σ is not Blackwell more informative than σ' , there exist some preferences such that a decision maker would strictly prefer σ' to σ .

In the case of distortion bias, a truthful reporting strategy is weakly more informative (in the Blackwell sense) than any biased strategy, and so consumers cannot be better off seeing a biased report. This comparison can be strict, for example if s is a binary signal and σ sometimes sets $n = s$ and other times sets n to be the opposite of s . Or it may be weak, for example if $s \in \mathbb{R}$ and $\sigma(s)$ is equal to s plus a constant. In the latter case, a sophisticated consumer can simply invert the mapping σ to recover s . Of course consumers may not be this sophisticated, and even if they are they may or may not know the correct mapping σ . Even distortion of the “add a constant” form can entail welfare loss if consumers mistakenly interpret reports as if they were accurate.

In the case of filtering bias, on the other hand, two reporting strategies are not typically ordered by Blackwell informativeness. A selection strategy that always reports the facts most favorable to R is optimal for some decision problems and a selection strategy that always reports the facts most favorable to L is optimal for some others. The same applies to endorsement strategies that usually support right-wing candidates and endorsement strategies that usually support left-wing candidates. Bias that is relatively “centrist,” such as a selection rule that reports the mean of the facts in s , or an endorsement policy that endorses the right-wing and left-wing candidates with equal frequency, has no special normative status.

Of course, political information may have important externalities beyond the direct effect on consumer welfare. Assessing the effect of bias via these external effects is challenging. Many have argued that, all else equal, democracy will function better when citizens are better informed. If we accept this premise, distortion bias, which reduces information in the Blackwell sense, will often be bad from both a social and a private perspective. In the case of filtering bias, on the other hand, there is less reason to expect the social and private incentives to be well aligned.

4 A model of the market for news

We now introduce a stylized model of the news market, which we use to capture key points from the discussion above, and to summarize and frame our subsequent discussion of the literature. There is a continuum of consumers indexed by i , and either one or two media firms, each operating at most one outlet. As above, the state is $\theta \in \{L, R\}$, firms observe independent signals s , and firms make reports n . We restrict attention to the case of binary reports, $N = \{l, r\}$. The prior probability of $\theta = R$ is 0.5.

We consider particular forms of both distortion and filtering bias. In both cases, strategies will be represented by a single-dimensional index of bias $b \in [-1, 1]$, where a strategy with bias b is biased to the right of a strategy with bias b' if and only if $b \geq b'$. We assume firms announce and commit to a reporting strategy b before the signal s is observed.⁴

To model distortion bias, we assume that signals are binary, with $S = N = \{l, r\}$, and reports are like cheap talk in that firms can potentially report either l or r following any s . Signals are symmetric, with precision $\pi \equiv Pr(s = l|L) = Pr(s = r|R) > 0.5$. In the cases we consider below, a firm may find it optimal to distort either r signals or l signals (i.e., report $n = l$ when $s = r$ or report $n = r$ when $s = l$), but will never distort both with positive probability. We therefore let $b > 0$ be the probability of reporting $n = r$ following $s = l$, $b = 0$ represent truthful reporting, and $b < 0$ be (minus) the probability of reporting $n = l$ following $s = r$.

To model filtering bias, we assume that signals are continuous, with $S = \mathbb{R}$. We assume that s is normally distributed with standard deviation 1 and mean $m > 0$ in state R and $-m$ in state L . Here, we define precision $\pi \equiv Pr(s \geq 0|R) = Pr(s \leq 0|L)$. We restrict attention to strategies defined by a threshold τ such that $n = r$ if $s \geq \tau$ and $n = l$ otherwise. We let $b = Pr(n = r|s < 0)$ when $\tau < 0$, $b = -Pr(n = l|s > 0)$ when $\tau > 0$, and $b = 0$ when $\tau = 0$.

In a slight abuse of notation let $\sigma_l = Pr(n = l|L)$ and $\sigma_r = Pr(n = r|R)$, i.e., the probabilities of making a correct report in states L and R , respectively. Under both distortion and filtering, $\sigma_l + \sigma_r$ is continuous in b and maximized at 2π when $b = 0$. Also note that $\frac{\partial \sigma_l}{\partial b} < 0$ and $\frac{\partial \sigma_r}{\partial b} > 0$ for all b . With filtering bias, $\sigma_l + \sigma_r$ is inverted-U shaped, and thus $\frac{\partial(\sigma_l + \sigma_r)}{\partial b}|_{b=0} = 0$. With distortion bias, $\sigma_l + \sigma_r$ is inverted-V shaped.

⁴The strategy can be thought of as an editorial policy that cannot be easily changed (as explained by, for example, Gehlbach and Sonin 2013).

Consumers value information about θ , possibly because it has instrumental value for a real decision they must make (e.g., voting, labor supply, what claims to make at a cocktail party), or because they care intrinsically about having accurate beliefs. As a shorthand for these incentives, we assume consumers choose an action $A \in \{A_L, A_R\}$ whose payoff for i is $u_i(A|\theta)$, with $u_i(A_L|L) = \beta_L > 0$, $u_i(A_R|R) = \beta_R > 0$, $u_i(A_L|R) = u_i(A_R|L) = 0$. We let p denote a consumer's prior probability that $\theta = R$, which may be different from the true prior probability 0.5. We take $p = 0.5$ and $\beta_L = \beta_R = 1$ as baseline assumptions. Below, we consider cases in which both the prior p and the payoffs $\{\beta_L, \beta_R\}$ may differ across i , and we introduce the possibility that consumers may receive psychological utility from confirming their prior beliefs. We refer to the average (over i) of $E(u_i(A|\theta))$ as consumer welfare.

Each consumer can choose to consume one report or no report. Consuming news has no monetary cost, but entails a time cost c_i which is distributed on $[0, \bar{c})$ with CDF $G(c)$, PDF $g(c)$, and $\bar{c} > 0$. We assume that c_i is independent of preferences and priors. If the consumer is indifferent between the two duopolists' reports, each firm is chosen with probability 0.5.

Each firm earns profit equal to the mass of consumers watching its report times per-consumer advertising revenue ρ . Below, we consider the possibility that firms may also care directly about consumer actions.

The timing of the game is:

1. Nature chooses the state of the world, which is unobserved to all players.
2. The firm or firms choose (and commit to) their biases, which are then known to consumers.
3. The firm or firms observe their signals privately.
4. Firms and consumers simultaneously report and (possibly) consume news, respectively.
5. Consumers update beliefs if they consumed news, choose their actions and receive their payoffs.

Before proceeding, we note that in the baseline case, in which firms only care about profit and consumers are homogeneous with centrist priors and preferences, there will be no media bias. Consumer i consumes news if expected utility from her action given news minus c_i is greater than

expected utility without news, which is 0.5. So i will consume news if

$$\begin{aligned} Pr(n = l|L)Pr(L) + Pr(n = r|R)Pr(R) - c_i &> 0.5 \\ 0.5(\sigma_l + \sigma_r - 1) &> c_i. \end{aligned} \tag{1}$$

Thus, the fraction of consumers who get news is $G(0.5(\sigma_l + \sigma_r - 1))$. Since $\sigma_l + \sigma_r$ is maximized when $b = 0$ (for both types of bias), the firm(s) will choose $b^* = 0$. Thus, news is unbiased and consumer welfare is maximized in both monopoly and more competitive markets, and total firm profits are the same in all market structures. Since consumers agree on wanting unbiased news, and the only incentive of firms is to provide consumers with what they want, consumers receive optimal news with or without firm competition, for either signal structure.

5 Supply-driven bias

One potential driver of bias is that firms may prefer consumers to take particular actions. Such preferences could arise from many sources. Another chapter in this handbook considers incentives arising from pressure by governments, advertisers, or other third parties. Here, we consider incentives internal to the firm. These could include direct preferences of media owners, or arise indirectly, from the preferences of reporters or editors.

The literature has reached three main conclusions about the implications of supply-driven bias. First, supply-side incentives can drive both filtering and distortion bias in equilibrium and can lead to manipulation of even rational, sophisticated consumers. If persuasive incentives are strong enough relative to the profit motive, firms will prefer to tilt their reports to affect consumer actions even at the cost of profits. Second, competition generally reduces distortion and moves outcomes back toward what would occur in the absence of persuasive incentives, even if the competing firms have similar biases. Intuitively, competition tends to align outcomes better with the demand of consumers, which in the baseline case is for unbiased news. Third, competition tends to increase consumer welfare, although it may increase or decrease total surplus once the ideological payoff of the owners is taken into account.

To illustrate these results in the context of our model, we assume that each firm has a preferred

action. In addition to its monetary profit, the firm earns a payoff equal to $\alpha > 0$ times the fraction of consumers taking its desired action. We call firms whose preferred action is A_R right-biased and firms whose preferred action is A_L left-biased.

As shown above, the fraction of (homogeneous, centrist) consumers who get news is $G(0.5(\sigma_l + \sigma_r - 1))$. Consumers who do not get news are indifferent between actions, so suppose they choose A_L with probability 0.5. The probability that a consumer who does get news chooses A_L is $\sigma_l Pr(L) + (1 - \sigma_r) Pr(R) = 0.5(\sigma_l + 1 - \sigma_r)$. Thus the expected fraction of all consumers choosing A_L is $0.5(\sigma_l + 1 - \sigma_r)G(0.5(\sigma_l + \sigma_r - 1)) + 0.5(1 - G(0.5(\sigma_l + \sigma_r - 1)))$.

Suppose there is just one left-biased firm. Then its objective function is α times the fraction of consumers choosing A_L plus ρ times the fraction of consumers consuming news:

$$\begin{aligned}\Pi_M &= \alpha[0.5(\sigma_l + 1 - \sigma_r)G(0.5(\sigma_l + \sigma_r - 1)) + 0.5(1 - G(0.5(\sigma_l + \sigma_r - 1)))] + \rho G(0.5(\sigma_l + \sigma_r - 1)) \\ &= \alpha(0.5(\sigma_l - \sigma_r)G(0.5(\sigma_l + \sigma_r - 1)) + 0.5) + \rho G(0.5(\sigma_l + \sigma_r - 1)).\end{aligned}\quad (2)$$

For distortion bias, the derivative of this expression with respect to b is negative when $b \geq 0$ for sufficiently large α and small ρ , and it is straightforward to show Π_m is not maximized when $b = -1$, so there exists an interior optimum, $b^* \in (-1, 0)$ if $\alpha(\rho)$ is large (small) enough. It can be shown with implicit differentiation that $|b^*|$ is increasing in α and decreasing in ρ . Notice that bias influences the distribution of consumer actions in equilibrium even though consumers are rational. This is an example of rational persuasion as in Kamenica and Gentzkow (2011).

To see the effect of competition, note first that all consumers will choose whichever firm has the lowest $|b|$. If there were a second firm with an opposing preference for consumer actions, then the two firms each offering biased news cannot be an equilibrium since if either firm had an equal or greater bias (than its competitor), then it could be strictly better off in terms of both profits and consumer actions by undercutting its competitor's bias. The only equilibrium is both firms offering unbiased news. More surprisingly, the same result occurs when both firms have the same preference for consumer actions. If one firm offered biased news, the other firm could obtain the entire market by undercutting this bias marginally, while only marginally decreasing the desired consumer actions. Thus there would always be an incentive to undercut if either firm offered biased news.

Since consumer welfare is decreasing in bias, clearly welfare is higher in duopoly. Results are very similar for filtering bias, except that a monopolist will choose some bias for any $\alpha > 0$. The following proposition summarizes these results.

Proposition 1. *Suppose firms strategically choose distortion (filtering) biases. Then:*

1. *If α is sufficiently large and ρ sufficiently small (for all $\alpha, \rho > 0$), then a biased monopolist will choose $|b^*| \in (0, 1)$, with $|b^*|$ increasing in α and decreasing in ρ .*
2. *In duopoly, both firms will choose $b^* = 0$.*
3. *Consumer welfare is higher in duopoly than monopoly.*

Even a monopolist with $\rho = 0$ will not choose maximal bias, for any α . Choosing total bias would alienate consumers, causing them to not consume news at all, and hence not be influenced (Gehlbach and Sonin stress this point). The literature on supply-driven bias extends these basic intuitions in a variety of ways.

Baron (2006) noted that since many news organizations are part of corporations, it seemed implausible that they would sacrifice profit for political bias. He showed that bias could enhance profit, however, despite the negative effect on readership, if journalists are willing to accept lower pay when allowed to bias the news, which they might want to do to advance their own career concerns or political preferences. In this case, it can be profitable for media owners to allow even distortion bias in exchange for pay cuts. In his model bias causes news prices to be lower due to lower demand, and has an ambiguous effect on welfare as it may increase profits, but the welfare effect is negative if demand for news is high. With competing firms, the relation between bias and profits is ambiguous.

Anderson and McLaren (2012) study a model like ours in that media firms have preferences over consumer actions, and consumers have homogeneous political preferences and are fully rational and sophisticated about equilibrium firm behavior. The media's access to information is exogenous but the media can strategically withhold information (a version of distortion bias). Consumers can still be manipulated by bias since when political information is bad (from the firm's perspective), and is not reported, the consumers cannot know whether this is due to non-disclosure or the information being unavailable. Their model goes beyond ours in also considering prices, the

incentives of firms to merge, media costs, and allowing consumers to get news from more than one source. Consistent with our model, they find that competition is effective in reducing distortion, and preventing mergers can improve welfare. More subtly, they note that when competing media firms have opposing biases, mergers may not occur even when allowed, due to the owners' conflicting political motives. Moreover, joint duopoly profits can be higher than monopoly profit, as a single owner with a political bias cannot credibly commit to offering diverse, differentiated news.

Chan and Suen (2009) also study a model with rational consumers and biased outlets. Their model goes beyond ours by incorporating two-party political competition with endogenous platforms. Voter payoffs depend on the winning party's platform and an unobserved state of the world. The authors compare two cases: one in which the media simply report each party's platform, and another in which the media reports the platforms and also makes a cheap talk report on the state. They find that in the first case ("balanced reporting") the parties are undisciplined and choose extreme, polarized platforms, which is socially harmful. Platform convergence, which would be ideal in their model, is not an equilibrium because each party would have an incentive to deviate, and voters could not identify and punish a unilateral deviator. By contrast, in the second case of the model, the media report, even if biased, moderates the policies, since if one party was to deviate and propose a more extreme policy, then this party could be identified by voters given the media report. Greater competition enhances this moderation further. A subtle finding is that bias does not drive policy in its direction, but instead polarizes candidate platforms. Since the report is cheap talk it is a type of distortion bias. The degree of bias is constrained not by the need to keep the audience, as this is exogenous, but by the alignment between the media's and public preferences.

Brocas et al. (2011) consider several variations on a simple model of media competition with biased firms. One novel feature they consider is that in addition to biasing its own news, a firm can signal-jam a rival firm's news, making it uninformative. Consequently, duopoly competition is not sufficient to eliminate bias. However, when there are two firms supporting each of two "viewpoints" (political candidates), i.e. four firms in total, bias is completely eliminated. This is due to what they refer to as "informational Bertrand competition"—if two firms have the same viewpoint and one is slightly more informative, it obtains the entire market for that viewpoint.⁵ Hence, there cannot be an equilibrium in which one firm can become marginally more informative

⁵A similar effect is found in Ellman and Germano (2009).

and capture the whole market. The authors augment their model with experiments, which generally support their theoretical predictions.

6 Demand-driven bias

The other potential driver of bias is demand from consumers themselves. A robust conclusion of the empirical literature discussed in the next chapter is that consumers tend to choose media whose biases match their own preferences or prior beliefs—what we will refer to as confirmatory news. In the presence of such tendencies, even profit-maximizing media firms that do not care directly about influencing consumers' beliefs may choose biased reports.

An obvious question is why consumers would demand confirmatory news in the first place. The literature has identified three categories of incentives which could be at play, which we will call delegation, psychological utility, and reputation. We organize our discussion around these three categories.

The broad conclusions of the literature with regard to demand-driven bias differ in important ways from those when bias originates on the supply-side. First, demand-side incentives will drive filtering bias, but will usually not lead to distortion if consumers are rational and prefer more information (in the Blackwell sense) to less. The exception is when firms have reputational concerns and cannot commit to truthful reporting, as in Gentzkow and Shapiro (2006). Second, the effects of competition on bias are more ambiguous. This largely follows from standard Hotelling logic. For example, duopolists may adopt more extreme biases than a monopolist if they each cater to a single part of the ideological spectrum while the monopolist tries to appeal to both. Third, although it often remains true that competition improves consumer welfare, the broader effects of competition can be less clear-cut if more polarized media outlets generate negative political externalities.

6.1 Delegation

The first explanation for why consumers demand confirmatory news is specific to the case where bias takes the form of filtering. Recall that in this case, reporting rules are not typically Blackwell ordered, and so the value of a signal to a rational consumer will be a function of the consumer's priors and preferences. An observation originally due to Calvert (1985) and expanded on by Suen

(2004) shows that consumers whose preferences (or priors) favor state R will tend to get higher (subjective) expected utility from right-biased signals. Thus, even rational consumers will exhibit a preference for confirmatory news.

To see the intuition in the context of our model, note first that a signal only has value if it affects consumers' actions; we can thus focus attention on signals such that $A = A_R$ if $n = r$ and $A = A_L$ if $n = l$. Second, observe that basing one's action on such a signal is equivalent to delegating one's decision to the media outlet. The mapping from the underlying signal s to the consumer's action A becomes simply the reporting rule σ . Consumers will then clearly prefer reporting rules that are close to the mapping from s to actions they would have chosen if they could have observed s directly themselves. As either priors or preferences shift to the right, the consumer's optimal mapping chooses A_R for a larger set of s , so their preferred signal reports $n = r$ for a larger set of s as well.

Several results seem intuitive when delegation incentives drive bias. First, firms should never engage in outright distortion, since this strictly reduces the value of firms' products to consumers; the delegation incentive can only explain differences in filtering. Second, the analysis of competition will be similar to the analysis of differentiated product competition in other settings: increasing the number of firms tends to make consumers better off, because it makes it more likely someone will produce a product that caters specifically to their tastes. Third, welfare effects are subtly different when consumers have heterogeneous priors than when they have heterogeneous preferences, because in the former case at least some consumers are making systematic mistakes and so catering to their demands need not be optimal.

To show these results in the context of our model, we consider two cases. In the first, consumers have heterogeneous preferences: half are "leftists," with $\beta_L > \beta_R = 1$, and half are "rightists," with $\beta_R > \beta_L = 1$. In the second, consumers have heterogeneous priors: they all have $\beta_R = \beta_L = 1$, but half are "leftists" with prior $p_i = p_L < 0.5$ and half are "rightists" with prior $p_i = 1 - p_L > 0.5$.

To see that firms never engage in distortion bias, consider the case of a rightist consumer under heterogeneous preferences, with right-distortion biased news. The utility of a consumer who bases her action on the firm's report would be:

$$\sigma_l Pr(L) + \beta_R \sigma_r Pr(R) - c_i = 0.5(\pi(1 - b) + (\pi + (1 - \pi)b)\beta_R) - c_i. \quad (3)$$

This expression is decreasing in b if $\beta_R < \pi/(1 - \pi)$ and increasing in b otherwise. But in the latter case, it is easily shown that it cannot be utility-maximizing to consume news. Thus, for consumers who do choose to consume news in equilibrium, utility is decreasing in distortion bias, and so utility must be maximized with zero distortion bias. Since firms only want to attract as many consumers as possible, and can only do so by offering them greater utility, there will be no distortion bias in any equilibrium. This result holds for consumers with heterogeneous priors by equivalent reasoning. Intuitively, distortion bias makes reporting Blackwell less informative, reducing utility for any consumer.

For filtering bias, given a consumer's action follows the firm's report, the derivative of utility with respect to b is $0.5 \frac{\partial(\sigma_l + \sigma_r \beta_R)}{\partial b}$, which must be strictly positive when $b = 0$ if $\beta_R > 1$, since as mentioned above, $\frac{\partial(\sigma_l + \sigma_r)}{\partial b}|_{b=0} = 0$, and $\frac{\partial \sigma_r}{\partial b} > 0$ for all b . Thus, for a rightist with $\beta_r > 1$, utility must be maximized for some $b > 0$, and a firm that targets just one consumer segment will offer the segment its optimal bias.

A monopolist that operates one outlet may find it optimal to do this—to cater to just one segment—if the segments are sufficiently polarized. Otherwise, a monopolist would find it optimal to appeal to both segments with $b = 0$. If we assume that $g' \leq 0$, it is then unambiguous that the share of consumers gained on one side of the spectrum from increasing bias is less than the share of consumers lost on the other side. But for any degree of consumer heterogeneity, duopolists would have incentives to differentiate via filtering bias. If both duopolists offered $b = 0$, then they would split the monopoly audience. They would then each face an incentive to at least marginally change b , since by doing so they would continue to attract half the monopoly audience (all the consumers consuming news with $b = 0$ whose preferences are on the side of the firm's new bias), plus they would attract new consumers only willing to consume biased news. Given that one duopolist offers $b \neq 0$, the best response for the other duopolist is to offer the utility maximizing filtering bias to the other consumer segment. In equilibrium, then, each firm delivers one segment of consumers its optimal bias. The bias cannot be extreme ($|b| \neq 1$) because then news would be uninformative. Let $0.5(\hat{\sigma}_l + \hat{\sigma}_r)$ denote the unconditional probability of a correct signal with these optimal thresholds (this is the same for the optimal threshold of either a leftist or rightist). Note that a monopolist that operates two outlets would behave in the same way as the duopolists.

It is intuitive, and easily verified, that consumer welfare is higher in duopoly than monopoly,

if consumers have heterogeneous preferences. This is not necessarily the case if consumers have heterogeneous priors. Let \hat{c}_D denote the the cost above which a consumer will not consume news even if it is optimally biased. Then, using our assumption that the objective prior is the mean of subjective priors (0.5), welfare in duopoly, W_D , can be written as

$$W_D = 0.5(\hat{\sigma}_l + \hat{\sigma}_r)G(\hat{c}_D) + 0.5(1 - G(\hat{c}_D)). \quad (4)$$

In monopoly with one outlet with $\tau = 0$, letting \hat{c}_M be the analog to \hat{c}_D , welfare W_M is

$$W_M = \pi G(\hat{c}_M) + 0.5(1 - G(\hat{c}_M)). \quad (5)$$

Duopoly minus monopoly welfare can be written

$$W_D - W_M = \underbrace{(\pi - 0.5)(G(\hat{c}_D) - G(\hat{c}_M))}_{\text{market size effect} > 0} + \underbrace{(0.5(\hat{\sigma}_l + \hat{\sigma}_r) - \pi)G(\hat{c}_D)}_{\text{information effect} < 0}. \quad (6)$$

Since $\hat{c}_D > \hat{c}_M$, the “market size effect,” the growth in consumers getting news due to differentiation, is strictly positive, while the “information effect,” the news being incorrect more often than necessary due to filtering bias, is strictly negative. Which effect is larger simply depends on $G()$ and $\hat{\sigma}_l + \hat{\sigma}_r$. The market size effect is larger when $G()$ is steeper and $g()$ is flatter, and differentiation expands the news audience substantially because getting news is costly for many consumers. The information effect is larger when consumers are more polarized and filtering bias is more costly (m is larger). Since either effect can be shrunk arbitrarily close to zero, either effect could dominate and the difference could take either sign.⁶

We summarize these results as follows.

Proposition 2. *Suppose consumers have heterogeneous preferences (priors). If β_L and β_R (p_L and p_R) are sufficiently close to 1 (0.5), then:*

1. *Under distortion bias, $b^* = 0$ for all firms under either monopoly or duopoly.*
2. *Under filtering bias, a monopolist chooses $b^* = 0$ and duopolists choose $b^* \in (0, 1)$, $b^* \in (-1, 0)$ that maximize (subjective) expected utility for each segment. Consumer welfare is*

⁶Chan and Suen (2008) and Chan and Stone (2013) decompose the effects of competition in a similar way.

(may or may not be) higher under duopoly.

If β_L and β_R (p_L and p_R) are sufficiently large (far from 0.5), then

- 1. Under distortion bias, no consumer will consume news for any level of bias.*
- 2. Under filtering bias, a monopolist will offer one segment its (subjective) optimal bias. Duopolists will offer each segment its (subjective) optimal bias, and welfare is strictly greater than in monopoly.*

We now discuss a few papers from the literature that relate to this case and expand on it.

Chan and Suen (2008)'s model is similar to ours, but incorporates political competition. In their model, all consumers are interested in news on one issue, a one-dimensional state of the world, but have different preferences on which party to vote for in a given state. The media can only report a binary signal although the true state is continuous, and so consumers have different preferences over reporting strategies—some want the media to only endorse the leftist party in extreme states (as these are the only states in which those voters prefer to vote leftist), some want the rightist party only endorsed rarely, and some want a more even mix. This leads to endogenous filtering bias, as in our model. Unlike in our model, however, due to the continuum of consumer heterogeneity in their model, competing duopoly outlets take the same position. This is akin to the classic Hotelling result of convergence in duopoly with horizontal differentiation and no price competition, first applied to media markets by Steiner (1952). Like Steiner, Chan and Suen point out that if the two firms have the same owner, this can cause differentiation. However, Chan and Suen's results are more complex as their model incorporates political platform competition. This causes the relationship between media differentiation and welfare to be more ambiguous than in our model.

Burke (2008)'s basic model with heterogeneous consumer priors is also very similar to ours, with a binary state and news report, and firms that just want to maximize market size. He considers a type of filtering bias in which firms can choose σ_l and σ_r , but the sum of the two, $\sigma_l + \sigma_r$, is constant (the foundations of this reporting technology are not modeled). He shows this set-up implies that if a consumer's prior is at least marginally non-centrist, then the consumer thinks that she gets maximal information value from extremist media, and consequently, duopolists maximally

differentiate. A monopolist may still take a centrist position, as it does not want to alienate either side of its audience. In his model, given the filtering technology there is no information loss (or gain) from filtering bias, so welfare unambiguously improves with differentiation due to market expansion. Burke extends the analysis by looking at a dynamic setting, showing results hold even if consumers are initially unbiased—being forward looking and knowing the future value of biased news, they prefer biased news from the start. He also shows results are robust to incorporating price competition.

Strömberg (2004) develops a model in which duopoly newspapers do not care about manipulating consumers and only want to maximize profits. The firms are *ex ante* identical. His model goes beyond ours in incorporating news production costs (assuming an increasing returns to scale technology), multiple dimensions of news on different government programs, and political competition. Filtering bias in his model takes the form of firms focusing coverage on some programs more than others. He finds that (unlike in our model) the two firms behave the same in equilibrium, both focusing coverage on groups that are larger, more valuable to advertisers, and care more about the policy outcome, due to the increasing returns to scale. He also models political behavior and shows that this filtering bias leads these consumers to receive better policy outcomes.

Duggan and Martinelli (2011) assume that media are constrained to make a one-dimensional report on a two-dimensional policy space. In their model, an election winner chooses levels of an income tax and a public good, and consumers choose whether or not to make an investment necessary to obtain a high paying job. Voters are informed about the incumbent's platform but only know the challenger prefers more of the tax and public good, but not the exact platform. Duggan and Martinelli analyze the case of the media having political preferences, but also analyze optimal filtering for voters, which is the filtering that would occur if the media just wanted to maximize profits. If the return on jobs is high enough, optimal bias (for consumers) is no information on the tax since consumers will invest in the high paying job anyway, and full information on the public good so consumers can vote for the challenger if her public good proposal is optimal. If the return on jobs is low, optimal bias balances information on taxes to induce consumers to invest, while providing some information on the public good to improve voting on that dimension. However, "no bias" (equal reporting on both dimensions) is never optimal, consistent with the principle that biased preferences causes demand for biased information.

6.2 Psychological utility

We next consider the case in which consumers get direct utility from news whose bias matches their own prior beliefs. This “psychological” utility is a function only of bias and consumers’ priors, and is separable from the instrumental value of the information in a news report. The foundations for preferences over bias have been developed in both psychology and economics (Lord et al. 1979; Nickerson 1998; Rabin and Schrag 1999; Tirole 2002). Such a taste could result from a preference for self-image, esteem or consistency (it may “feel good” or be more entertaining to receive news confirming one’s beliefs are correct, and others are wrong), a desire to avoid complexity (updating away from the prior may involve cognitive costs), or other reasons.

We denote psychological utility by $\phi(b; p_i)$. Assume $\phi(0; p) = 0$ for all p , and $\phi(b; p = 0.5) = 0$ for all b (that is, this utility is only received if news is biased and the prior is not perfectly centrist). If $p_i > 0.5$, then $\phi()$ is increasing in b at a decreasing rate, and $\phi()$ is symmetrically decreasing in b for $p_i < 0.5$.⁷ We maintain the assumption that $g' \leq 0$.

Assume consumers have heterogeneous priors, and again consider first the case of a monopoly outlet. Utility without news for a leftist is $1 - p_L$, so a necessary condition for a leftist to consume news is:

$$\begin{aligned} (1 - p_L)Pr(n_l|L) + p_L Pr(n_r|R) + \phi(b; p_L) - c_i &> 1 - p_L \\ (1 - p_L)(\sigma_l - 1) + p_L \sigma_r + \phi(b; p_L) &> c_i, \end{aligned} \quad (7)$$

The condition for rightists is symmetric. If these conditions are satisfied for some $c_i > 0$, then the monopolist’s payoff is

$$\Pi_M = 0.5(G((1 - p_L)(\sigma_l - 1) + p_L \sigma_r + \phi(b; p_L)) + G((1 - p_R)\sigma_l + p_R(\sigma_r - 1) + \phi(b; p_R))). \quad (8)$$

It can be shown that this expression decreases as b moves away from zero, and so $b^* = 0$ is a local maximum, for either bias type. It may also be optimal to offer one of the two segments its optimal bias (if p_L, p_R are sufficiently far apart).

Thus far the analysis for this case has been basically the same as that of the delegation case. The

⁷It is not necessary for our analysis to specify different properties of $\phi()$ for distortion and filtering bias, so we omit this notation for simplicity.

situation is different though for distortion bias in duopoly. Expected utility from news for a leftist is increasing in left distortion bias if $\pi + p_L - 1 < \left| \frac{\partial \phi(b; p_L)}{\partial b} \right|$; that is, if the marginal psychological utility of bias is greater than the marginal cost of lost information. This means that if $\pi + p_L - 1 < \left| \frac{\partial \phi(b; p_L)}{\partial b} \right|_{b=0}$, then in duopoly the firms will differentiate and offer each segment of consumers its optimal bias, whether bias is distortion or filtering bias (the logic for why differentiation is an equilibrium is the same as that discussed above). Henceforth assume this condition holds.

The comparison between monopoly and duopoly welfare is ambiguous, for the same reason as the case without psychological utility—there is a tradeoff between more consumers getting news, and consumers getting less informative news. We summarize these results as follows.

Proposition 3. *Suppose consumers have heterogeneous priors and get psychological utility from news. If p_L and p_R are sufficiently close to 0.5, then:*

1. *A monopolist will choose $b^* = 0$ under distortion or filtering bias.*
2. *In duopoly, under either type of bias, firms will choose $b^* > 0$, $b^* < 0$ that maximize subjective expected utility for each segment. Welfare may be higher or lower in duopoly than monopoly.*

If p_L and p_R are sufficiently far from 0.5, then

1. *Under distortion or filtering bias, a monopolist will offer one segment its optimal bias.*
2. *Under distortion or filtering bias, duopolists will offer each segment its optimal bias, and welfare is strictly greater than in monopoly.*

Note that for this case, unlike the delegation case, duopolists might prefer to use distortion bias, rather than filtering bias if distortion bias is less costly. It is also worth noting that this is the only case that we analyze in which total bias, i.e. $|b| = 1$, is possible, and news would become completely uninformative. This would occur if psychological utility were important enough.

Mullainathan and Shleifer (2005) study a model in which the state of the world is a real number, utility is lost when the news is further from the consumer's prior mean, and media firms only want to maximize profits. Their model also incorporates price competition. They show that uninformative distortion bias can exist in many types of equilibria. When consumers are homogeneous,

duopoly reduces price but not bias (if consumers have “biased priors” i.e. a prior mean not equal to zero); when consumers are heterogeneous, duopolists differentiate and become even more extreme than the most extreme consumers, as this softens price competition. Mullainathan and Shleifer also examine an extension in which consumers who wish to can cross-check news in duopoly, and show these consumers can become more informed this way than they would be in monopoly.

Gabszewicz et al. (2001) analyze a similar Hotelling-style model of politically differentiated duopoly newspaper competition. They assume consumers prefer media content to be closer to their political ideologies. Gabszewicz et al. also consider price as an endogenous variable, and confirm that with quadratic transportation costs, in equilibrium the firms maximally differentiate to soften price competition. The authors add to this result by showing that if firms then sell their readership to advertisers, this moderates differentiation and may drive prices to zero.

Bernhardt et al. (2008) study a voting model in which media consumers are rational information processors, but partisan voters get utility from negative news about the opposing party. Although voters draw correct inferences from their news, they choose news to maximize this non-instrumental, psychological utility from news, given that their votes are not pivotal, and so are often not as informed as they could be. Their paper thus highlights the importance of knowledge externalities, which are not incorporated in our model. Again, consumers demand information-destroying distortion bias, although in this case it is distortion by omission. Since voters do not become as informed as possible, sub-optimal electoral outcomes may occur; these are more likely when the distribution of voter preferences is asymmetric and the median voter is partisan. Even when the median voter is non-partisan and receives unbiased news there can be electoral problems though, as partisan voters may vote (ex post) incorrectly due to distortions in their news.

Schulz and Weimann (1989) were, as far as we know, the first to study a model of media competition incorporating this type of non-standard consumer psychology. They assume consumers prefer to avoid cognitive dissonance, and for this reason prefer newspapers with ideologies similar to their own. In their model, two political parties and two newspapers both choose ideological locations and a separate, vertical aspect of information, while newspapers choose prices as well. The newspapers and parties do not directly interact in the model. The authors find that both parties and newspapers differentiate, but newspapers differentiate more, but not maximally, due to the presence of the vertical type of information.

6.3 Reputation

The last reason we discuss that consumers may receive confirmatory news is that they trust the firms that provide it more. A rational consumer who is uncertain about an information source's accuracy will tend to judge it to be higher quality when its reports match the consumer's priors. Thus, firms can have incentives to produce confirmatory news to improve their reputations, increasing future demand. This observation was made by Prendergast (1993) and Brandenburger and Polack (1996) in the context of labor and financial markets, respectively, and applied to media markets by Gentzkow and Shapiro (2006).

There are two main takeaways for this case. First, if firms cannot commit to their strategies ahead of time, their incentives to build a reputation for quality can lead them to engage in distortion bias in equilibrium, even though this strictly reduces the value of their reports to consumers, and can make both them and consumers worse off. Second, the relationship between bias and competition involves a new element. As above, Hotelling-type forces may or may not lead duopolists to adopt more biased positions than monopolists. Additionally, however, consumers can use the reports of one outlet to cross-check the other, and this may discipline firms' incentives to bias their reports. We illustrate the first of these implications formally, and refer the reader to Gentzkow and Shapiro (2006) for more detail on the second.

To capture reputation in the context of our model, we extend it to include two types of media firms: a high-quality type that observes the state directly and reports it honestly, and a normal type that observes a noisy signal and reports strategically (as above). We focus on the case of distortion bias, where the normal firms observe $s \in \{r, l\}$ and can freely choose $n \in \{r, l\}$, and we drop the assumption that such a firm commits to its strategy *ex ante*. (If we kept this assumption, consumers could infer the firm's type directly from its announced strategy.) We assume there is a monopoly firm, and assume the *ex ante* probability it is a high type is λ . We assume consumers are homogeneous with prior belief $p \in (0.5, 1)$.

Firm profits depend on both current advertising revenue and future profits; we capture the latter in reduced form by simply assuming they are increasing in consumers' posterior belief that the firm is high-quality. Thus, since the normal-type firm cannot commit to its reporting strategy, it will always choose whichever strategy maximizes consumers' posteriors.

A consumer who observes $n = r$ will have a posterior belief that the firm is a high type of

$$Pr(high|r) = \frac{Pr(r|high)Pr(high)}{Pr(r)} = \frac{p\lambda}{p\lambda + [\sigma_r p + (1 - \sigma_l)(1 - p)](1 - \lambda)},$$

and a consumer who observes $n = l$ will have posterior belief

$$Pr(high|l) = \frac{Pr(l|high)Pr(high)}{Pr(l)} = \frac{(1 - p)\lambda}{(1 - p)\lambda + [(1 - \sigma_r)p + \sigma_l(1 - p)](1 - \lambda)}.$$

where σ_r and σ_l now denote $Pr(n = r|R)$ and $Pr(n = l|L)$ respectively from the perspective of the consumer given her beliefs about the firm's strategies. The firm will report r whenever $Pr(high|r) > Pr(high|l)$, and will report l whenever $Pr(high|r) < Pr(high|l)$.

Suppose, first, that consumers expected the firm to report its signal truthfully. In this case, $\sigma_r = \sigma_l$, and so $Pr(high|r) > Pr(high|l)$. The firm would prefer to deviate and always report r . Thus, truthful reporting cannot be an equilibrium.

There also cannot be an equilibrium in which the firm biases its reports left, nor one where it only reports r , because in this case $\sigma_r = 1$, $\sigma_l = 0$, and $Pr(high|r) < Pr(high|l)$.

The unique equilibrium must therefore be in mixed strategies, with an interior bias $b \in (0, 1)$ chosen such that $Pr(high|r) = Pr(high|l)$. This occurs at the point where

$$\frac{p}{1 - p} = \frac{1 - \sigma_l}{1 - \sigma_r}.$$

Substituting $\sigma_r = \pi + (1 - \pi)b$ and $\sigma_l = \pi(1 - b)$, we can solve for the equilibrium bias:

$$b^* = \frac{(2p - 1)(1 - \pi)}{p(1 - \pi) + \pi(1 - p)}.$$

It is straightforward to show that this expression is increasing in p , so that bias is greater the more extreme are consumer beliefs. It also happens that it is decreasing in π , so that firms with more accurate signals engage in less bias.

Bias can occur in equilibrium even when it makes all market participants strictly worse off. Bias always reduces consumer welfare. More surprisingly, it can also decrease firm profits. To see this, suppose that low-type firms were required to report their signals truthfully. This would

unambiguously increase first-period profits because the value of signals would be higher and so more consumers would read the firm's report. The effect on continuation payoffs is ambiguous. However, Gentzkow and Shapiro (2006) show that if the baseline share λ of high-type firms is sufficiently low, the first-period effect dominates and firm profits fall.

Gentzkow and Shapiro extend this basic analysis in several ways. First, they allow consumers to receive feedback on the news before firms receive their reputation payoffs. They show that in the case of homogeneous but non-centrist priors, in the unique equilibrium a monopolist distorts its report towards the prior. The degree of distortion decreases as the chance of feedback increases, disappearing when feedback is certain. This helps explain why empirically bias is lower for news on events where "truth" is revealed relatively quickly and clearly, such as sports, weather, and election outcomes (as opposed to events where it is difficult to immediately evaluate truth, e.g., the effects of fiscal stimulus on GDP or climate change). Second, as already mentioned, they study a specific form of media competition, in which the presence of multiple firms allows consumers to cross-check the firms' reports, and they show that this form of competition reduces bias. The role of competition here is closely related to the observation in other settings that competing advocates may be more effective at eliciting truth than a single non-partisan (e.g., Dewatripont and Tirole 1999).

Stone (2011) analyzes a related model with distortion bias, and consumers trusting like-minded media to deliver more accurate news. In his model, both consumers and reporters have information processing biases. Consumers think of themselves as unbiased, however, causing them to infer that reporters with similar ideologies are also unbiased. This makes consumers believe (often falsely) that ideologically similar media is less biased and more informative. Consequently, consumers unwittingly demand distortion biased confirmatory news. Increased media differentiation resulting from greater competition can cause media to become more biased and less informative. This can reduce welfare, but need not since the differentiation due to competition generates a positive market size effect.

7 Conclusion

Bias can persist in commercial media markets for a variety of theoretical reasons. Supply-side bias persists when media management or labor are willing to sacrifice profits for political gain. Demand-side bias persists when consumers perceive biased media to be more informative or more enjoyable. Bias caused by reputation concerns persists when feedback on truth is weak.

Despite its negative connotations, bias as we have defined it need not be socially harmful. In some cases—outright distortion or needless omission—the consequences for consumers, at least, will be unambiguously negative, and it is meaningful to talk about a benchmark of “unbiased” reporting. In other cases—filtering or selection—an efficient market may supply different varieties of bias, and there is no meaningful notion of unbiasedness.

The effects of competition on bias and consumer welfare are in general ambiguous. Competition tends to sharpen firms’ incentives to give consumers what they want. When supply-side incentives lead firms to distort their reports, competition tends to lessen bias and improve welfare. When consumers themselves demand bias, competition may lead to more extreme biases that cater to these tastes. This often improves welfare if consumers are rational and information-seeking, but otherwise may not. When distortions originate in firms’ incentives to build a reputation for quality, the availability of information from competing sources may allow consumers to distinguish true quality more accurately, and so reduce bias by softening these reputational incentives.

References

- Alterman, Eric. 2003. *What Liberal Media?: The Truth About Bias and the News*. NY: Basic Books.
- Anderson, Simon P. and John McLaren. 2012. "Media Mergers and Media Bias with Rational Consumers." *Journal of the European Economic Association* 10(4): 831–859.
- Andina-díaz, Ascensión. 2011. "Mass Media in Economics: Origins and Subsequent Contributions." *Cuadernos de Ciencias Económicas y Empresariales* no. 61: 89-101.
- Baron, David P. 2006. "Persistent Media Bias." *Journal of Public Economics* 90(1-2): 1–36.
- Bernhardt, Dan, Stefan Krasa, and Mattias Polborn. 2008. "Political Polarization and the Electoral Effects of Media Bias." *Journal of Public Economics* 92(5–6): 1092–1104.
- Besley, Timothy and Andrea Prat. 2006. "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability." *American Economic Review* 96(3): 720-736.
- Blackwell, David. 1951. "Comparison of Experiments." *Second Berkeley Symposium on Mathematical Statistics and Probability*: 265-272. University of California Press.
- Brandenburger, Adam and Ben Polak. 1996. "When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See." *The RAND Journal of Economics* 27(3): 523-541.
- Brocas, Isabelle, Juan D. Carrillo, and Simon Wilkie. 2011. "Media Ownership Study No. 9: A Theoretical Analysis of the Impact of Local Market Structure on the Range of Viewpoints Supplied." Mimeo, Federal Communications Commission.
- Burke, Jeremy. 2008. "Primetime Spin: Media Bias and Belief Confirming Information." *Journal of Economics & Management Strategy* 17(3): 633–665.
- Calvert, Randall L. 1985. "The Value of Biased Information: A Rational Choice Model of Political Advice." *The Journal of Politics* 47(2): 530-555.
- Chan, Jimmy and Daniel F. Stone. 2013. "Media Proliferation and Partisan Selective Exposure." *Public Choice* 156(3-4).
- Chan, Jimmy and Wing Suen. 2008. "A Spatial Theory of News Consumption and Electoral Competition." *The Review of Economic Studies* 75(3): 699-728.
- Chan, Jimmy and Wing Suen. 2009. "Media as Watchdogs: The Role of News Media in Electoral Competition" *European Economic Review* 53(7): 799–814.
- Crawford Vincent P. and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50(6): 1431-1451.
- Dewatripont, Mathias and Jean Tirole. 1999. "Advocates." *Journal of Political Economy* 107(1): 1-39.
- Downs, Anthony. 1957. "An Economic Theory of Political Action in a Democracy." *Journal of Political Economy* 65(2): 135-150.
- Duggan J. and C. Martinelli. 2011. "A Spatial Theory of Media Slant and Voter Choice." *Review of Economic Studies* 78(2): 640-666.
- Ellman, Matthew and Fabrizio Germano. 2009. "What do the Papers Sell? A Model of Advertising and Media Bias." *The Economic Journal* 119(537): 680–704.
- Gabszewicz, Jean J., Dider Laussel, and Nathalie Sonnac. 2001. "Press Advertising and the Ascent of the 'Pensée Unique'." *European Economic Review* 45(2): 641-651.
- Gehlbach, Scott and Konstantin Sonin, 2013. "Government Control of the Media." Working Paper.

- Gentzkow, Matthew and Jesse M. Shapiro. 2006. "Media Bias and Reputation." *Journal of Political Economy* 115(2): 280–316.
- Gentzkow, Matthew and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence from US Daily Newspapers." *Econometrica* 78(1): 35-71.
- Goldberg, Bernard. 2001. *Bias: A CBS Insider Exposes How the Media Distort the News*. Washington DC: Regnery Publishing.
- Hamilton, James T. 2004. *All the News That's Fit to Sell: How the Market Transforms Information Into News*. Princeton University Press.
- Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101 (6): 2590–2615.
- Ladd, Jonathan M. 2012. *Why Americans Hate the Media and How it Matters*. Princeton, NJ: Princeton University Press.
- Lord, Charles G., Lee Ross, and Mark R. Lepper. 1979. "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence." *Journal of Personality and Social Psychology* 37(11): 2098-2109.
- Milgrom, Paul and John Roberts. 1986. "Price and Advertising Signals of Product Quality." *Journal of Political Economy* 94(4): 796-821.
- Mullainathan, Sendhil and Andrei Shleifer. 2005. "The Market for News." *American Economic Review* 95(4): 1031-1053.
- Nickerson, Raymond S. 1998. "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises." *Review of General Psychology* 2(2), 175-220.
- Pew Research Center for the People and the Press. 2011. *Views of the News Media: 1985-2011: Press Widely Criticized, But Trusted More than Other Information Sources*. Accessed at <<http://www.people-press.org/files/legacy-pdf/9-22-2011%20Media%20Attitudes%20Release.pdf>> on December 16, 2013.
- Prat, Andrea and David Strömberg. 2011. "The Political Economy of Mass Media." CEPR Discussion Paper No. DP8246.
- Prendergast, Canice. 1993. "A Theory of 'Yes Men'." *The American Economic Review* 83(4): 757-770.
- Rabin, Matthew and Joel L. Schrag. 1999. "First Impressions Matter: A Model of Confirmatory Bias." *The Quarterly Journal of Economics* 114 (1): 37-82.
- Schulz, Norbert and Joachim Weimann. 1989. "Competition of Newspapers and the Location of Political Parties." *Public Choice* 63(2): 125-147.
- Sobbrio, Francesco. 2013. "The Political Economy of News Media: Theory, Evidence and Open Issues". *Handbook of Alternative Theories of Public Economics*. Cheltenham: Edward Elgar Press. Forthcoming.
- Steiner, Peter. 1952. "Program Patterns and Preferences, and the Workability of Competition in Radio Broadcasting." *Quarterly Journal of Economics* 66 (2): 194-223.
- Stone, Daniel F. 2011. "Ideological Media Bias." *Journal of Economic Behavior & Organization* 78(3): 256-271.
- Strömberg, David. 2004. "Mass Media Competition, Political Competition, and Public Policy." *Review of Economic Studies* 71 (1): 265-284.
- Suen, Wing. 2004. "The Self-Perpetuation of Biased Beliefs." *The Economic Journal* 114(495): 377-396.

- Sutter, Daniel. 2001. "Can the Media Be So Liberal? The Economics of Media Bias." *Cato Journal* 20(3): 431.
- Tirole, Jean. 2002. "Rational Irrationality: Some Economics of Self-Management." *European Economic Review* 46(4-5): 633-655.
- Zaller, John. 1999. *A Theory of Media Politics*. Unpublished manuscript.