Media Competition with Targeted Advertising: How Media Polarization Attenuates Political Polarization *PRELIMINARY AND INCOMPLETE - DO NOT CITE OR DISTRIBUTE*

Samarth Gupta* and Benjamin G. Ogden[†]

1-7-17

Abstract

Many posit that recent increases in media polarization and specialization are a cause for recent political polarization, particularly in the United States. We develop a model of endogenous media bias, voter learning, and voting to test this mechanism. We find that in the old media environment (with no ability to target advertising, and hence no media specialization), convergence did not necessarily indicate truth-telling: if a sufficient portion of the population has a taste for ideological media, firms would bias towards moderation in order to avoid alienating potential viewers on either side. The ability to target advertising (and specialize) will generate media polarization, but at least one media firm reports the truth and will capture ideologically unbiased voters. Therefore, the total level of information transmitted weakly increases, and expected political polarization decreases. The model demonstrates why modern media polarization is an unlikely culprit for political polarization, and may actually attenuate divergence.

---

*Boston University, Department of Economics

[†]Texas A&M, Department of Political Science & Université libre de Bruxelles, European Centre for Advanced Research in Economics and Statistics. Contact: bgogden@tamu.edu

1 Introduction

The media plays a central role in how voters acquire information about the effectiveness of policy. The complexity of policy changes often require voters to rely upon the simplified information provided by reporters and specialized forums, such as those found on the internet. Therefore, identifying the degree to which information will be conveyed remains of great interest to social scientists

In particular, one area of recent concern surrounds the rise of "targeted advertising", or the catering of ads to a specific viewer/group as a function of its online browsing activity. While this concept is not novel, it was much more, often prohibitively, difficult to implement with old media technology. With the emergence of online media, the concept has been turned into a practical aid for marketers and advertisers. By using profiles of online activity and browsing histories available via digital files known as "cookies", the cost of extracting information regarding a viewer are reduced. The profiles of particular viewers and demographic groups are then used to construct specially catered ads. A report by BIA/Kelsey estimated targeted advertisement to grow from $6.8bn in 2015 to $18.2bn in 2019, while the Network Advertising Initiative estimated that "behaviorally targeted advertising secured an average of 2.68 times as much revenue per ad as non-targeted runs of network ads" (Chamberlain (2015); Initiative (2010)).

This is of potential concern because the behaviors that are targeted by advertisers are heavily correlated with political ideologies. For example, 46% of consistent liberals prefer urban centers, while 41% of consistent conservatives prefer rural areas. Similarly 77% of consistent liberals prefer small houses closer to each other, while 75% of consistent conservatives prefer larger houses that a farther apart. Therefore, the benefits of targeted advertising increases the incentives of media, and particularly new media, to target viewers based upon their political inclination, which in turn could increase their incentive to skew reporting. Indeed, we have seen political polarization rise (e.g., McCarty, Poole and Rosenthal (2006)) at the same time as media polarization (e.g., Groseclose and Milyo (2005); Mullainathan and Shleifer (2005); Prior (2005); Gentzkow and Shapiro (2006); Duggan and Martinelli (2010); Gentzkow and Shapiro (2010); Puglisi (2011)). As noted by Bakshy, Messing and Adamic (2015), with the rise of Facebook's News Feed as a primary source of information, people actively encountered 70% less cross-cutting content. This is a concern that has also been raised both within the theoretical (e.g., Bernhardt, Krasa and Polborn (2008)) and empirical literature (e.g., DellaVigna and Kaplan (2007)), often with a focus upon cable news.

In order to test these concerns, we build a formal model in which citizens act as both consumers of news media and voters. What differentiates our model is that consumer/voters

have not only instrumental preferences over the media firms (i.e., they want to know the true state of the world in order to make informed voting decisions), but also intrinsic preferences over what state of the world is reported (i.e., confirmation bias), consistent with, e.g., Gentzkow, Shapiro and Sinkinson (2014). Media firms compete by committing to a bias in their reporting over the true state of the world, thereby allowing consumers to decide how much information they want reported and in what form (i.e. the precise state reported for each true state of the world). For example, while a liberal voter may prefer to know whether gun control policy works at all, they would also prefer to hear that gun control works very well to hearing that gun control works a little bit, regardless of the underlying state.

When firms simply try to maximize the number of consumers for their product (analogous to the old media), the firms will converge to offering the same expected bias. The convergent strategy will be to tell the truth if and only if the majority of consumers care more about updating their beliefs over the state of the world than their intrinsic preferences. In this case, both firms will report the true state and there will be efficient implementation of policy.

However, if most consumers are dominated by their intrinsic preferences, such that they want to hear that their preferred policy is the optimum for society, the two firms will each offer a mixed strategy. This will involve mixing between i) reporting the truth, ii) reporting with a slight skew in the direction of the majority of the population, and iii) reporting the median state regardless of the true state of the world. In particular, (iii) is consistent with media that avoid reporting any extremes out of fear of alienating some segment of the population. Therefore, there is less than efficient information revelation, as media are forced to worry about those voters with strong intrinsic preferences.

Now consider targeted advertising, which we assume, as described above, is correlated with voters' intrinsic preferences. If the incentives for targeting become strong enough (as we describe here), there will become an incentive for divergence of the firms. In particular, one firm will commit to a bias skewed in whichever direction features a plurality of the population. This is done as the firm seeks to capture exclusive rents over this plurality, which is worth more than $\frac{1}{2}$ the population elsewhere.

However, if targeted advertising becomes strong enough to incent divergence, it will always also incentivize at least one other firm to tell the truth. Therefore, given that voters can sort based upon expected bias, those voters that will actually make use of the information will acquire it. Therefore, there will always be a (weakly) stronger incentive for information revelation with targeted advertising than with the old media. Therefore, while media polarization will be greater, political polarization will actually be attenuated, as voters will rely less on their priors.

In addition, we consider a setting in which one or both of the media firms are trying

to influence the policy outcome of the election (e.g., what Roger Ailes has been accused of with Fox News). We show, consistent with models without intrinsic preferences (e.g., Oliveros (2015)) that as long as one firm is a profit-maximizer, policy-motivated firms cannot influence the election. If each extreme position has a firm trying to achieve implementation of its position, then if voter preferences are sufficiently risk-averse, the incentive for "radical moderation" with the old media will be greater. Therefore, while there may not be full information revelation with the new media (i.e., no one will necessarily report the truth), targeted advertising will still lead to weakly more revelation than without specialization.

The literature on media bias and political outcomes is vast and growing. For a general overview of the field, see Prat and Stromberg (2011). By focusing upon rational updating voters and firms which report information about the state of the world, we differentiate ourselves fully from one subset of the literature which focuses on media as endorsement heuristics for voters (e.g., Chiang and Knight (2011); Castaneda and Martinelli (2016)). Instead, we focus here on papers which examined endogenous media bias.

Similarly to our primary setting, these papers focus on firms which are profit maximizing and respond to demands for bias. Previous papers in this setting (e.g., Mullainathan and Shleifer (2005); Gentzkow and Shapiro (2006)) considered consumers with preferences for slant in news reporting, and generated media polarization (see Oliveros (2015) for an exception where consumer preferences are inversely related to media preferences). They considered neither the trade-off for voters between preference for slant and preference for information, the technology shift that would lead to a specialized media, nor the ex-post political outcome. Chan and Suen (2008) consider media that are exogenously limited in their ability to report the full truth, while Stromberg (2004) considers a similar setting to ours (increasing returns to scale), but without intrinsic and instrumental preferences, cannot generate divergence.

The other primary setting for endogenous media bias is on the supply-side. In this setting, it is easier to generate that political bias and efficiency is reducing in competition and specialization (e.g., Besley and Prat (2006); Gentzkow, Glaeser and Goldin (2006); Anderson and McLaren (2012); Sobbrio (2014); see Baron (2006) for an exception); however, inconsistent with this setting, Gentzkow and Shapiro (2010) found no evidence that editor or owner identity played a role in newspaper skew.

Our results stem from identifying two incentives that have been overlooked in the previous literature. First, with intrinsic preferences, it is possible that we may not get truth-telling from convergence. This is because sufficiently strong intrinsic preferences will deter firms from putting off consumers on either side of the preference distribution. This can help explain bland, uninformative news media which reports in the style of a press release. For

example, the Associated Press, a bastion of old media, has often been accused of reporting uncritically of both Presidents Obama ("Final jobs report for Obama presidency expected to be solid") and Trump ("Trump expected to embrace bold use of Twitter"). Second, there will be sorting as people know the bias of the media, so it is sufficient to have only one firm offering truth-telling. This is consistent with Groseclose and Milyo (2005) finding that while media generally exhibits a left-shift, there still exist truthful reporting, such as ABC's Good Morning America and PBS' Newshour with higher average net ratings. Therefore, specialization and competition will lead to better information acquisition, contra some of the models above without intrinsic preferences, but consistent with the more positive empirical evidence from, e.g., Gentzkow and Shapiro (2011); Gentzkow, Shapiro and Sinkinson (2014); Schroeder and Stone (2015). Therefore, while it is true that media and political polarization have been rising at the same time, media polarization is unlikely to be to blame for the rise, and in fact may be attenuating it.

We lay out the model in section 2. Sections 3 and 4 consider the equilibria without and with targeted advertising and specialization, respectively. Section 5 considers an extension with potential electorally-oriented firms. We discuss the implications of the model and conclude in sections 6 and 7.

2   The Model

2.1   Voters/Consumers

The population is composed of a continuum of voter-consumers. They have preferences over both policy (as a function of the state of the world) and the intrinsic properties of whatever media they consume.

There will be three states of the world which can be arrayed along a line, L, M, R, with each state occurring with equal probability. There will also exist a corresponding set of policies: l, m, r. These two collections together represent the efficiency of different policies in each state. For example, we could consider the relevant state of the economy, in which L represents when the economy is in severe recession, M represents full employment, and R represents an overheated economy. Therefore, Keynesian policies (represented here by l) will work well in L but are damaging in R, and vice-versa with Austerian policies (represented by r). Meanwhile, technocratic moderation (m) will be optimal in M.

Voters will be split into five groups based upon their ***ideological bias***. λ-type voters will be extremely biased, such that they prefer policy l in all states of the world, while ρ voters will always prefer policy r. μ_l and μ_r voters will be moderately biased, preferring l

(r, respectively) in states L (R) and M, and m otherwise. The remainder, μ (no subscript), will be unbiased, preferring to match the policy directly to the state. For future ease, we will refer to μ, μ_l, and μ_r voters as **updaters**, since their preferences are capable of changing based upon the state of the world.

The voter-consumers also feature intrinsic preferences over the type of news they consume, independent of the realized state. As each state of the world adheres to different properties of the world, given that a media source reports a state j, they will be reporting on a different set of topics and using a different tone. Therefore, different consumers will have preferences over hearing about these different states.

There will similarly exist five types of single-peaked ordinal intrinsic preferences. Let a voter's intrinsic preference be represented by a super-script L if they prefer left-wing news, ML if they prefer moderate news to left-wing news to right-wing news, and so on. Therefore, a voter-consumer of type μ_l^{ML} will have a moderate ideological bias to the left, but will always prefer moderate news to left-wing news.

Voters rank making the "correct" policy decision lexiographically over their intrinsic news preferences.[1] Therefore, updaters will, when given the chance, always consume news which partitions the state space in a policy-relevant way; but, when indifferent between two news sources that do so, will choose the one that provides news that is consistent with their intrinsic preferences. Trivially, non-updaters will choose media based entirely upon their intrinsic preferences.

This stylization, admittedly made for tractability, initially appears stark; however, there is mounting evidence that voters do, in fact, consider the informativeness of a media source as of first-order importance relative to other desires. For example, Bruns and Himmler (2016) find that voters were willing to pay significant sums to switch to informative mass media if they share a common cause with other voters. This is consistent with the modeling assumption taken above that updaters, who will be the median voters, will try to acquire information before treating their own intrinsic preferences.

Abusing notation, let a voter's ideological type (e.g., λ) also represent the unconditional proportion of that type in the population. Therefore, $\lambda + \rho + \mu_l + \mu_r + \mu = 1$. Each voter's intrinsic media taste will then be drawn from a conditional distribution ζ_i, which is dependent on their ideological type i. In particular, if you order the ideological biases in the natural way (i.e., $\lambda < \mu_l < ... < \rho$) and the intrinsic preferences similarly, then when $i < j$, ζ_i is strictly first-order stochastically dominated by ζ_j.

---

[1]Note that we could make the viewer preferences continuous, and it would not change the results. For example, we could consider voter utilities which put a weight on both intrinsic preferences and instrumental value of information. As long as the state state remains discrete, we could determine cut-off values such that preferences would look *as if* they were lexiographic of the form we have here.

In order to make the problem interesting, let us make the following assumption:

Assumption 1: $\frac{1}{2} > \rho + \mu_r > \lambda + \mu_l$

Therefore, there are more right-wing voters than left wing voters, but neither make up a majority of the population without the aid of moderates. Note that Assumption 1 implies that an unbiased voter will always represent the median. Therefore, in a world with perfect information, the Condorcet Winner would be the policy that "matches" the state of the world, and such a policy would be enacted in any democratic equilibrium.

Additionally, for tractability, we will make the following assumption:

Assumption 2: $\zeta_\lambda(L) = 1 = \zeta_\rho(R)$, $\zeta_{\mu_l}(L) + \zeta_{\mu_l}(ML) = 1 = \zeta_{\mu_r}(R) + \zeta_{\mu_r}(MR)$, $\zeta_\mu(ML) + \zeta_\mu(M) + \zeta_\mu(MR) = 1$, $\zeta(M) < \frac{1}{2}$, and $\frac{\zeta_{\mu_l}(L)}{\zeta_{\mu_l}(L) + \zeta_{\mu_l}(ML)} = \frac{\zeta_{\mu_r}(R)}{\zeta_{\mu_r}(R) + \zeta_{\mu_r}(MR)}$.

Assumption 2 puts more specific constraints on the correlation between voter-consumers' policy and taste preferences. In particular, all left-wing extremists must prefer left-wing news, while left-wing updaters must, at the least, prefer left-wing news to right-wing news (and vice-versa for right-wingers). This is consistent with the idea that the desire to avoid cognitive dissonance incentivizes people to consume news that justifies what they know their ex-post policy preferences will recommend (e.g., Gentzkow, Shapiro and Sinkinson (2014)). The final piece of assumption 2 simply requires that moderately biased consumers on both sides of the political aisle feature strong intrinsic bias in the same proportions. This rules out, e.g., there being more moderate-left consumers with strong-left biases then moderate-right consumers with strong-right biases, despite there being more right-biased consumers on net.

Before any election, voter-consumers will choose one, and only one, media source of those available to consume. After updating their beliefs accordingly, they will then vote over a set of politicians who compete in a Downsian fashion. Therefore, voter type i's strategy is to pick media source $x_i \in \{1, 2\}$ and a voting profile as a function of their information partition $y_i(p)$ such that $y_i(\cdot) \in \{l, m, r\}$. The final implemented policy will be based upon the preferences of the median voter, which, by the assumption above, will be a moderate.

Note that this description of voter-consumers remains open to three different notions of polarization:

Definition 1: *An increase in λ and ρ (with a corresponding decrease in μ_l^L and μ_r^R), or an increase in μ_l^{ML} and μ_r^{MR}, is an increase in **ideological polarization**.*

Ideological polarization represents an increase in state-independent ideology on the part of voters, holding their intrinsic taste for types of news constant. An increase in ideological polarization implies that a society's preferences are becoming less tied to the state of the

world and more reliant upon (non-shared) ideology. Note that, under conditions of full information, ideological polarization would have no impact on the preferences of the median voter, and hence no impact on the outcome. It should be noted, in addition, that we have little evidence that ideological polarization has increased in the United States over the past half-century (e.g., Fiorina, Abrams and Pope (2005); Ozdemir and Ozkes (2014)).

Definition 2: *An increase in $\zeta_{\mu_l}(L)$ and $\zeta_{\mu_r}(R)$, or an increase in $\zeta_\mu(ML)$ and $\zeta_\mu(MR)$, is an increase in* **taste polarization**.

Taste polarization represents an increase in the intrinsic preference for news of a certain type independent of their underlying policy preference. There is more evidence of this type of polarization as voters begin to sort more into media that fits their political priors (e.g., Gentzkow and Shapiro (2006)). This also makes sense considering the evidence of Tesler (2016) that voters symbolic political tastes form early in the life cycle based upon life experiences and backgrounds, which are becoming increasingly divergent in a globalized world.

These first two forms of polarization are primitives of the model. The following form of polarization is an outcome:

Definition 3: *A distribution of ex-ante policy preferences has greater* **political polarization** *if it has greater variance.*

Political polarization represents an increase in the *revealed* political preferences of voters. This will, of course, be an equilibrium outcome of the interaction between voter preferences and their beliefs formed after viewing the available media. Indeed, this is the variable of outcome that has been on the clear rise in the United States over the last fifty years and which has often been blamed on the targeting and specialization of media firms.

2.2 Media Firms

Each of two media firms seek to maximize their profit π, which is a strictly increasing function of the number of voter-consumers who view their programming. Their strategy space is to pre-commit to a skew $\sigma_j \in \{\hat{L}, L\hat{M}, \hat{M}, \hat{M}R, \hat{R}, T\}$. We can think of T as committing to truth-telling, and therefore reporting the true state of the world in each state. Each of the other skews represents a commitment to a particular bias in reporting. $\hat{L}$, $\hat{M}$ and $\hat{R}$ each represent reporting their respective state in each possible revelation of ω; by contrast, $L\hat{M}$ and $\hat{M}R$ represent a slight skew in line with the policy preferences represented by μ_l and μ_r.

Therefore, we can fully rank the preferences over (non-dominated) media sources for each of the seven relevant types:

- λ (ρ) voters will always prefer a more left-biased (right-biased) source; whether they prefer T or $\hat{M}$ will depend upon whether they are risk-averse (in which case they prefer $\hat{M}$) or risk-seeking (in which case they prefer T).

- μ_l^L (μ_r^R) voters will most prefer $L\hat{M}$ $(\hat{M}R)$, followed by T. Barring one of those, they will prefer a more left-biased (right-biased) news source.

- μ_l^{ML} (μ_r^{MR}) will be the same, except preferring $\hat{M}$ to $\hat{L}$ $(\hat{R})$.

- μ^{ML} (μ^{MR}) will prefer T to their preferred mildly biased news source to their less-preferred mildly biased source. They will then prefer $\hat{M}$ to the respective extremely-biased news sources.

- μ^M will be similar, except that they will be indifferent between the mildly-biased, and then the extremely-biased sources, respectively.

Based upon these effective preferences, the media will commit to a skew in order to maximize the number of voters who will choose to view their programming.

In particular, we will consider two particular forms of π, though the model is open to testing different forms of profit-making. First, we will consider **non-targeted advertising**. We can think of this as the old media, when firms were less specialized to certain types of voter-consumers and the extent of advertising was the impact of getting eyes to see the particular advertising. In this case, π_j will simply be a linear function of the number of voters who choose j.

Within the new media environment, however, advertising firms are able to microtarget their advertising to subgroups based upon such items as political preference. They do this, in particular, by creating on-going relationships with consumers through e-mail lists, social media connections, and broadly "maximizing clicks". Of course, such advertising must go deeper than standard advertising, and therefore this can only be done for one group at a time. Therefore, we will also consider **targeted advertising** in which π_j will be linear in all by the largest ideological group K, where K is a voter-consumer's first-choice among intrinsically-preferred media. With respect to K, media firms will gain a value $\nu > 1$.

We can also define one final form of polarization:

Definition 4: *Let $\hat{\omega}_j(\omega)$ be what is reported by media firm j in state ω. If $|\hat{\omega}_1(\omega) - \hat{\omega}_2(\omega)|$ increases for all ω, this is an increase in **media polarization**.*

Media polarization is also an equilibrium outcome, and remains another phenomenon observed over the last 30 years within the United States that stands to be explained. In addition, the connection between media polarization and political polarization has often been posited, but has been insufficiently examined on a logical level.

2.3 Equilibrium

Given a distribution of voter-consumer preferences and a media profit function π, we will first seek the Nash equilibrium of the game such that firms pick their media skew σ_1 and σ_2, and then voter-consumers choose their media consumption strategy as a function of their type i, x_i. This will be the **media equilibrium**, and directly determines media polarization.

Given an equilibrium of the media consumption game, we will be able to determine the **political equilibrium**. The voters, now empowered with a (potentially) finer partition of information, will vote for a policy based upon the information they have: $y_i(p_i)$. This will determine political polarization, as well as the effective outcome, which will be the preference of the median voter type, μ^M.

Recall that the full-information Condorcet winner would be to pick the analogous policy to the state ω. Therefore, we will call an equilibrium **inefficient** if it does not result in the analogous policy always being chosen, and **less efficient** than another equilibrium if it results in the "correct" policy being chosen less often.

3 Equilibria without Targeted Advertising

Without targeted advertising, competitive media face a situation similar to a Hotelling framework. Voters have various demands for bias, and the media firms, playing a zero-sum game, face a pressure to converge towards offering the same coverage (in expectation). Therefore, in the world before targeted advertising and the ability to monetize specific consumer groups, we should expect that there is little difference between mass media sources:

Proposition 1: *Without targeted advertising, each media source will offer the same expected bias.*

Due to the fact that media have no carrying capacity (i.e., they can always allow more viewers with no marginal cost), competition without specialization draw them to offer the same message. Therefore, there is no incentive for voters to make any conscious choice between sources, and they are able to simply "turn on whatever":

Corollary 1: *Without targeted advertising, every viewer watches each media source with equal probability, and features the same expected information partition.*

Therefore, with targeted advertising, any differences between voters will be driven directly by their ideology, not differences in beliefs. They will tend to agree on the facts of the world as it exists, as they draw information from the same sources, but there will still remain some disagreement driven by ideology. Taken together, this allows for a relatively simple analysis of the media and political interaction without targeted advertising:

Theorem 1: *Without targeted advertising:*

- *If updaters are a majority and/or voters are risk-loving, the unique media equilibrium features $\sigma_1 = \sigma_2 = T$. The unique political equilibrium results in policy that matches the state of the world.*

- *If updaters are a minority and voters are risk-averse, the unique equilibrium involves both media sources mixing between T, $\hat{M}$, and $\hat{M}R$. The unique political equilibrium involves picking the policy which matches the reporting.*

- *Otherwise, both equilibria exist.*

As in any standard model of zero-sum, costless competition, when consumers have preferences that satisfy single-crossing, there will exist a single median preference which is the dominant strategy for each firm. The key question is when such a scenario obtains in the media setting. In this case, if a majority voters want at least *some* information, then truthtelling will be the median preference of all voters. Any attempt to report with bias would lead to losing more voters than the firm would gain (by Assumption 1), while an attempt to move towards being a "radical moderate" ($\hat{M}$) would lose all the updaters. This would also be true if voters are risk-loving, as $\hat{M}$ is dominated by T in such a world.

Trivially, since the media is constrained towards reporting the truth, all voters will have perfect and complete beliefs about the true state of the world. Therefore, each voter will make the ex-post optimal vote, and the policy which matches the true state of the world will be implemented. This scenario looks similar to what is often described by those lamenting the lack of accurate reporting in the current, specialized media space: voter-consumers acquire all the information they need, and the politically efficient outcome is achieved despite ideological differences.

Note, however, that with risk-averse voters who have mostly made up their minds, population preferences will no longer satisfy single-crossing. In such a setting, we can and will have a cycle. Suppose you are a media firm. If you expect the other firm to report the truth, you will have an incentive to offer a "radically moderate" position, as this will pick up all non-updaters who do not want to witness news that is radically different from their preconceived preferences. However, if your competitor was offering $\hat{M}$, then you would have an incentive to offer mildly right-wing biased news, as this would pick up both right-wing biased voters and unbiased viewers who are looking for at least some information. Finally, if your competitor were offering this policy, then you would want to be unbiased in order to pick up both left-wing voters and unbiased viewers who want a full revelation of the state

of the world. Therefore, the only equilibrium will involve playing a mixed strategy between these three strategies.

There are two incentives that pull the media firms away from reporting the news truthfully. The simplest involves an incentive to "skew" towards the plurality bias in the population (in this case, a right-wing bias). If your competitor is offering decidedly moderate coverage (say, $\hat{M}$), then you will want to skew your coverage in the direction of the majority of the population by biasing your coverage to the right. This can be done in such a way that picks up the unbiased voters as well, as biased-yet-informative is better than moderate-but-uninformative.

The more interesting incentive involves the pull towards "radical moderation". Since a majority of viewers are not updaters and watch mass media purely for consumption purposes, a majority of viewers would prefer to avoid watching media that disagrees strongly with their priors. Since, however, no ideology is large enough to make up a majority of the population, the best strategy for a media firm must take this into account by biasing their coverage towards reporting the moderate state in order to avoid making any particular group uncomfortable. This is consistent with the colloquial notion of milquetoast "news" that is useless as a source of information as it must avoid upsetting any particular viewer who holds sufficiently strong prejudices.

This pressure to (potentially) pull away from truth-telling indicates that (ex-ante) convergence of media firms does not necessarily mean fully informative reporting. In particular, while the mixed strategies will be such that rational voters will vote as if the reporting is truthful, they will be receiving potentially inaccurate information, leading to (depending on the state of the world and the draw of the media firms' mixed strategies) overly moderate policies in extreme states of the world, and overly right-wing policies in moderate states of the world.

Corollary 2: *Without targeted advertising, if updaters are not a majority of the population with risk-averse voters:*

- *An increase in ideological polarization will increase the probability of an inefficient political outcome, and increase the probability that such an inefficient outcome is too extreme.*

- *An increase in taste polarization will have no effect on the political outcome.*

- *An increase in the skew of the ideological distribution will decrease the probability of an inefficient political outcome, and increase the probability that such an inefficient outcome is too extreme.*

Unsurprisingly, the more ideologically polarized a polity is, the more incentive exists for the media firms to diverge farther from truth-telling, as they must worry more about putting off the extremes of the political distribution. Additionally, taste polarization has no impact without targeted advertising.

More surprisingly, an increase in the skew of the ideological distribution will actually decrease the probability of an inefficient outcome. Recall that the truth-telling equilibrium is broken by a fear of alienating either of the two extremes. Therefore, as one extreme becomes more pronounced, media firms only need to worry about that particular side of the distribution; therefore, there is less disincentive to tell the truth. Conditional on not telling the truth, however, inefficient outcomes will tend to be more extreme (in the direction of the skew).

4   Equilibria with Targeted Advertising

With the availability of targeted advertising, the media game is no longer strictly zero-sum. In particular, since each firm can target a particular group, there exists a potential gain to be made from specialization. Indeed, as you make the returns from targeting a specific group sufficiently large, there will be divergence in equilibrium. Indeed, as Theorem 2 shows, this specialization will take a specific form:

Theorem 2: *With targeted advertising, if v is sufficiently large and voters are risk averse, there will exist two media equilibria, each of the form $\sigma_i = T$ and $\sigma_{-i} = \hat{M}R$.*

The intuition behind the existence portion of theorem 2 is straightforward. Imagine that we are in the state of the world where $\nu \approx 1$ and there is vanishingly little room for targeting. The equilibrium that emerges will be the same as that which appeared in Theorem 1, and each firm will target its advertising towards whichever group is the biggest.

As the value of specialization increases, so too will the value of deviation towards whichever consumer group is largest. Suppose the largest group of consumers prefers right-biased news; then instead of simply mimicking the other firm, it will eventually make sense to instead differentiate yourself and begin offering news that's biased to the right. This will allow the deviator to capture the plurality of voter-consumers, and target their advertising to them. This will occur if and only if the value of gaining the other half of the plurality (which they were previously splitting in the convergent equilibrium) outweighs the value of splitting the majority. Hence, there exists the incentive (as expected) to open Fox News when the benefit of targeted advertising towards right-wing consumers becomes high enough.

Note, however, that conditional on one group deviating, the other will no longer have an

incentive to follow. Why? They will now be able to offer the truth, capture the majority of viewers, and specially target whichever sub-group is better represented between the remaining two. Since the other group will deviate only when the value of specializing in, and capturing all of, the right-wing consumers outweighs splitting all groups equally, it will also be the case that the best response to group $-i$'s deviation is to offer truth-telling. This will capture all remaining voter-consumers (a majority), and it must be the case that specializing in advertising targeting the largest remaining group will be more profitable than trying to fight with $-i$ for market share in right-wing news:

Corollary 3: *Targeted advertising leads to greater media polarization.*

Without targeted advertising, there was only the possibility for media polarization when both groups were mixing, and even then the polarization is only probabilistic. With sufficient targeted advertising, however, there will now be polarization in two of the three states of the world (L and M). Therefore, as expected, the ability of media firms to specialize will also lead to larger differences in *how* news is reported in any given state of the world.

However, it is not ex-ante obvious that this media polarization is a bad thing. In particular, with targeted advertising, there always exists one firm which commits to full truth-telling, something which is not necessarily true without specialization. Since voter-consumers know which media firm will do this, left-wing and unbiased updaters will sort into that firm (as it provides them with the information they need), while right-wing updaters will sort into the right-wing news (as they still receive the coarser information needed by them). Hence, all voter-consumers will now receive the information they need to make the correct decision. As the median voter is an unbiased updater, it will be the case in this setting that the equilibrium will be efficient:

Corollary 4: *If v is sufficiently large, there will be full information revelation in equilibrium, and the political equilibrium will be efficient.*

How does this compare to the case with media convergence? If updaters are a majority of the population, it is also the case without targeted advertising that information would be fully revealed. Therefore, there will be no difference between the outcomes under both institutions:

Proposition 2: *If updaters are the majority of the population, the political equilibrium will be the same with and without targeted advertising.*

Since political outcomes are the same when updaters are the majority of the population, the only potential difference originating from the availability of targeted advertising will be the (potential) increased media polarization. However, this polarization will have no impact

upon either political polarization nor the implemented policy outcome. This is because voters' preferred policies remain the same in all states of the world in either case.

By contrast, if updaters are a minority within the population (alternatively, when most voters care about media for purely intrinsic reasons), then there will be differences between the two settings. In this case, the lack of targeted advertising led to a mixed strategy equilibrium which prevented full information revelation. Therefore, the amount of political polarization was greater and the political outcome less efficient than under full revelation.

With targeted advertising, however, information is fully revealed and the efficient policy outcome can be reached. In addition, political polarization will be reduced despite the increase in media polarization, as voters will have more relevant information on net, and therefore will rely less on voting their priors:

Proposition 3: *If updaters are not the majority of the population, political polarization decreases with targeted advertising.*

Note that this means that we should expect media polarization and political polarization to be (weakly) inversely related. In the US context, for example, this is not true in a time series setting. However, we are not observing the counterfactual in which media specialization was deterred. Proposition 3 implies that, without media specialization, there may be *even greater* increases in political polarization in the erecent past. The competition between the New York Times and the Wall Street Journal may make all newspapers on average more informative to their respective readers.

Why is this the case? Media specialization becomes weakly beneficial because the voters who consume the skewed media find that truth-telling media are inframarginal to their voting decision; they only need to differentiate L from M and R to be able to make their decision. Therefore, they get all the information they need from skewed media.

By contrast, unbiased updaters need full truth telling to make relevant decisions, and by construction these consumers are pivotal in the elections. Therefore, they are potentially harmed without media specialization by the lack of a fully unbiased media source. However, with competition and specialization, there will be an incentive for at least one firm to offer this. Therefore, they will receive the greater information they need by being able to sort into this firm in the targeted advertising setting.

5   Electorally-Motivated Firms

TBA

6 Discussion

The results of the model are driven by the assumption of both *instrumental* and *intrinsic* preferences over the information provided by media firms. We remain agnostic among the possible microfoundations for this structure. One possibility, in the vein of Carrillo and Mariotti (2000), is to consider a time-inconsistent voter, who cannot guarantee that the preferences of his future self will be the same as his present self, and therefore will want to remain "strategically ignorant". For example, consider a voter who only owns an SUV. He currently does not care whether his fuel consumption affects global warming, but he cannot guarantee that his future self wouldn't feel shame about it and change his behavior. Therefore, he may prefer news media to report, ceteris paribus, that global warming is not a severe problem. Such an incentive could also be captured by a behavioral model of shame.

Another potential microfoundation involves rational ignorance (e.g., Austen-Smith (1991); Lupia and McCubbins (1998)). Since voters know that their marginal vote is highly unlikely to be pivotal, they will value the immediate consumptive value of news an order of magnitude greater than the informative value. In this setting, we would expect that those voters who do value information do so for intrinsic, sociotropic reasons. Therefore, it should be unsurprising to see such a high value put on confirmation bias.

Note that with targeted advertising, media firms face a tradeoff between maximizing the number of viewers or the intensity of viewers. Indeed, this is consistent with evidence from the Pew Research Foundation that viewers of, e.g., potentially biased cable news networks tend to be much more involved in politics and provide higher value targets for advertisers (*Political Polarization in the American Public* (2014)).

One implication of the model is that as voters become more ideologically polarized, targeted advertising actually becomes more useful for information revelation, as it reduces the quality of information available in the old media environment. Therefore, it is exactly in a polarized political environment such as that facing the United States today that specialized media will be most useful. Indeed, when specialized media is less useful is when the distribution is skewed; i.e., when one side dominates. Therefore, as long as the US remains divided fairly equally, then specialized media should not be a cause of concern.

Discuss empirical testable implications

7 Conclusion

TBA

References

Anderson, Simon P. and John McLaren. 2012. "Media Mergers and Media Bias with Rational Consumers." 10(4):831–59.

Austen-Smith, David. 1991. "Rational Consumers and Irrational Voters: A Review Essay on Black Hole Tariffs and Endogenous Policy Theory." *Economics and Politics* 3(1):73–92.

Bakshy, Eytan, Solomon Messing and Lada A. Adamic. 2015. "Exposure to ideologically diverse news and opinion on Facebook." *Science* 348 (6239):1130–1132.

Baron, David P. 2006. "Persistent Media Bias." 90(1):1–36.

Bernhardt, Dan, Stefan Krasa and Mattias Polborn. 2008. "Political Polarization and the Electoral Effects of Media Bias." 92(5):1092–1104.

Besley, Timothy J. and Andrea Prat. 2006. "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability." 96(3):720–736.

Bruns, Christian and Oliver Himmler. 2016. "Mass Media, Instrumental Information, and Electoral Accountability.".

Carrillo, Juan D. and Thomas Mariotti. 2000. "Strategic Ignorance as a Self-Disciplining Device." 67:529–544.

Castaneda, Alejandro and Cesar Martinelli. 2016. "Political Economics of Broadcast Media.".

Chamberlain, Lauryn. 2015. "Geo-Targeted Mobile Ad Revenues to Reach $18.2 Billion by 2019.".
URL:*http://www.geomarketing.com/geo-targeted-mobile-ad-revenues-to-reach-18-2-billion-by-2019*

Chan, Jimmy and Wing Suen. 2008. "A Spatial Theory of News Consumption and Electoral Competition." 75(3):699–728.

Chiang, Chun-Fang and Brian Knight. 2011. "Media Bias and Influence: Evidence from Newspaper Endorsements." *Review of Economic Studies* 78:795–820.

DellaVigna, Stefano and Ethan Kaplan. 2007. "The Fox News Effect: Media Bias and Voting." 122(3):1187–1234.

Duggan, John and Cesar Martinelli. 2010. "A Spatial Theory of Media Slant and Voter Choice." *Review of Economic Studies* 78:1–29.

Fiorina, Morris P., Samuel J. Abrams and Jeremy C. Pope. 2005. *Culture Wars? The Myth of the Polarized Electorate.* Pearson Longman.

Gentzkow, Matthew, Edward L. Glaeser and Claudia Goldin. 2006. *The Rise of the Fourth Estate. How Newspapers Became Informative and Why It Mattered.* University of Chicago Press pp. 187–230.

Gentzkow, Matthew and Jesse M. Shapiro. 2006. "Media Bias and Reputation." *Journal of Political Economy* 114:280–316.

Gentzkow, Matthew and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence from U.S. Daily Newspapers." *Econometrica* 78:35–71.

Gentzkow, Matthew, Jesse M. Shapiro and Michael Sinkinson. 2014. "Competition and Ideological Diversity: Historical Evidence from US Newspapers.".

Gentzkow, Michael and Jesse M. Shapiro. 2011. "Ideological Segregation Online and Offline." 126(4):1799–1839.

Groseclose, Tim and Jeffrey Milyo. 2005. "A Measure of Media Bias." *The Quarterly Journal of Economics* 120:1191–1237.

Initiative, National Advertising. 2010. "Study Finds Behaviorally-Targeted Ads More Than Twice As Valuable, Twice as Effective as Non-Targeted Online Ads.".
URL: *https://www.networkadvertising.org/pdfs/NAI$_{Beales_Release}$.pdf*

Lupia, Arthur and Mathew D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* Cambridge University Press.

McCarty, Nolan, Keith T. Poole and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches.* MIT Press.

Mullainathan, Sendhil and Andrei Shleifer. 2005. "The Market for News." *American Economic Review* 95(4):1031–1053.

Oliveros, Santiago. 2015. "Demand for Slant: How Abstention Shapes Voters' Choice of News Media." 125(587):1327–1368.

Ozdemir, Ugur and Ali Ihsan Ozkes. 2014. "Measuring Public Preferential Polarization." Ecole Polytechnique Centre National De La Recherche Scientifique Working Paper.

Political Polarization in the American Public. 2014.

Prat, Andrea and David Stromberg. 2011. "The Political Economy of Mass Media.".

Prior, Markus. 2005. "News vs. Entertainment: How Increasing Media Choice Widens Gaps in Political Knowledge and Turnout." *American Journal of Political Science* 49:577–592.

Puglisi, Riccardo. 2011. "Being the New York Times: the Political Behaviour of a Newspaper." *The B.E. Journal of Economic Analysis & Policy* 11:1–34.

Schroeder, Elizabeth and Daniel F. Stone. 2015. "Fox News and Political Knowledge." 126:52–63.

Sobbrio, Francsco. 2014. "Citizen-editors' Endogenous Information Acquisition and News Accuracy." 113:43–53.

Stromberg, David. 2004. "Mass Media Competition, Political Competition, and Public Policy." 71(1):265–284.

Tesler, Michael. 2016. *Post-racial or most-racial? Race and politics in the Obama era.* University of Chicago Press.

A   Proofs

A.1   Theorem 1

Suppose updaters are a minority and voters are weakly risk-averse. Let us look for equilibria.

First, note that an extreme bias is always dominated by a moderate bias, so $\hat{L}$ and $\hat{R}$ will never be played in equilibrium. Then, note that reporting $\hat{M}L$ is strictly dominated by reporting the truth. Therefore, we can limit ourselves to looking for strategies that mix between T, $\hat{M}$, and $\hat{M}R$. This can be written as a normal-form game:

	T	$\hat{M}$	$\hat{M}R$
T	$\frac{1}{2}, \frac{1}{2}$	$\upsilon, \lambda + \rho$	$\mu + \mu_l + \lambda, \mu_r + \rho$
$\hat{M}$	$\lambda + \rho, \nu$	$\frac{1}{2}, \frac{1}{2}$	$\lambda + \mu_l, \mu + \mu_r + \rho$
$\hat{M}R$	$\mu_r + \rho, \mu + \mu_l + \lambda$	$\mu + \mu_r + \rho, \lambda + \mu_l$	$\frac{1}{2}, \frac{1}{2}$

Note that if the other firm is reporting truthfully, you can profitably deviate by reporting $\hat{M}$ in all states of the world, since $\nu < \frac{1}{2}$. However, if the other firm is reporting $\hat{M}$, then you can profitably deviate by having only a slight right-wing bias. But if the other firm is reporting $\hat{M}R$, then a deviation to telling the truth is profitable. Therefore, we must look for an equilibrium in mixed strategies.

Let p be the probability of reporting truthfully, and let r be the probability of reporting $\hat{M}$. Then the payoff to truth telling is $\frac{p}{2} + r(\mu + \mu_l + \mu_r) + (1 - p - r)(\mu + \mu_l + \lambda) = \frac{p}{2} + (1 - p)(\mu + \mu_l + \lambda) + r(\mu_r - \lambda)$, the payoff from $\hat{M}$ is $p(\lambda + \rho) + \frac{r}{2} + (1 - p - r)(\lambda + \mu_l) = (1 - r)(\lambda + \mu_l) + p(\rho - \mu_l) + \frac{r}{2}$, and the payoff to $\hat{M}R$ is $p(\mu_r + \rho) + r(\mu + \mu_r + \rho) + \frac{1 - p - r}{2}$.

Therefore, the mixed strategy requires $r = \frac{\frac{1}{2} - \rho - \mu_r}{\frac{1}{2} - \mu_r - \mu_l}$, $1 - p - r = \frac{\rho + \lambda - \frac{1}{2}}{\frac{1}{2} - \mu_r - \mu_l}$, $p = \frac{\frac{1}{2} - \mu_l - \lambda}{\frac{1}{2} - \mu_r - \mu_l}$.

A.2 Theorem 2

Regardless of v, the best response to an extreme bias is either truth-telling or a moderate bias in the same direction, as this delivers a payoff of $v * max\{\mu + \mu_l^M + \mu_r^M, \mu_r^R + \rho\} + \mu_l^L + min\{\mu + \mu_l^M + \mu_r^M, \mu_r^R + \rho\}$ in the left-wing case, and similarly in the right-wing.

The best response to a moderate bias is truth-telling, as this delivers $\nu * max\{\lambda + \mu_l^L, \mu + \mu_l^{ML} + \mu_r^{MR}\} + min\{\mu_l^L, \mu + \mu_l^{ML} + \mu_r^{MR}\}$ in the right-wing case, and similarly in the left-wing.

The best response to $\hat{M}$ is either T ($\nu * max\{\mu + \mu_l^{ML} + \mu_r^{MR}, \mu_r^R\} + \mu_l^L + min\{\mu + \mu_l^{ML} + \mu_r^{MR}, \mu_r^R\}$) or $\hat{M}R$ ($\nu * max\{\mu + \mu_r^{MR}, \mu_r^R + \rho\} + min\{\mu + \mu_r^{MR}, \mu_r^R + \rho\}$).

Let v get arbitrarily large. Then the best response to T is $\hat{M}R$, which delivers $\nu * max\{\mu_r^{MR}, \mu_r^R + \rho\} + min\{\mu_r^{MR}, \mu_r^R + \rho\}$. Therefore, for v sufficiently large, the unique equilibrium will be $\{T, \hat{M}R\}$.

A.3 Theorem 3

TBA

B Extension to N-firms

TBA