Media (In)accuracy on Public Policy, 1980-2018

Stuart Soroka
University of Michigan
ssoroka@umich.edu

Christopher Wlezien
University of Texas at Austin
wlezien@austin.utexas.edu

Prepared for presentation at the Virtual Annual Meeting of the American Association of Public Opinion Research, 2020. Earlier versions of this paper was presented at the Annual Meetings of the American Political Science Association, Washington, DC, 2019, and the Southern Political Science Association, San Juan, 2020. The research is supported by National Science Foundation Grants SES-1728792 and SES-1728558.

Mass media are critical to representative democracy. This is well known, and acknowledged throughout modern political history, from the founding of democratic republics to the present day. The nature and accuracy of media coverage are central to effective accountability and responsive governments. Indeed, it is difficult to imagine how large-scale democracy could work without reasonably accurate media coverage of current affairs.

It is of some significance, then, that we are in the midst of both public and scholarly debate about the nature and quality of media coverage in the United States (US). The current academic debate is fueled by several factors. There is a growing body of work on journalists' mis-representation – and the public's corresponding misunderstanding – of scientific issues such as global warming and vaccinations (Boycoff and Boycoff 2004; Speers and Lewis 2005). There is a burgeoning literature on selective exposure and motivated reasoning, suggesting that even were media coverage to portray issues accurately, exposure and interpretation of that information would be subject to a range of pre-existing biases, and that this is enhanced in an increasing high-choice media environment (e.g., Prior 2007; Garrett 2009; Stroud, 2011; Arceneaux and Johnson 2013; van Aelst et al. 2017). And there are concerns about an increasingly ideologically-polarized electorate amplifying the likely impact of selective exposure and the systematically-biased media coverage that may accompany it (e.g., Abramowitz 2010; Ura and Ellis 2012).

Consternation over the accuracy of American mass media news has reached a fevered pitch in the wake of Russian interference in the 2016 US presidential election (e.g., Shane and Mazzetti 2018), the increasingly loose interpretation of facts by the Trump administration (e.g., Salam 2018) and ongoing claims of and concerns about "fake news" (e.g., Alcott and Gentzkow 2017; Kucharski 2016). There have been few moments since the rise of modern mass media during which information about current affairs was so suspect, not just by the public, but by media professionals as well.

The current climate in the US is in some ways relatively unique. The availability of inaccurate information is not one of those ways, however. There have of course been other periods of serious concern about media accuracy, including, if not especially, during the Vietnam War (see, e.g., Delli Carpini 1990). And there has always been variation in the quality and accuracy of media coverage over time and across issues and media outlets. There has always been accurate and inaccurate information about public policy, and there is accordingly a rich body of work, with well-developed theories and models that can help us examine instances in which media have facilitated or inhibited representative democracy.

This observation is the starting point for the work that follows. Our intention is partly to respond to very current concerns about the state of public affairs, especially in the role that mass media currently play in connecting policymaking and public preferences. But we also take seriously the possibility that, while there are times that media coverage is inaccurate, there are times when it is accurate as well. Our aim is to leverage existing theories, and over 40 years of data on public policy and media coverage, to better measure and understand both the successes and failures of mass media in modern representative democracy.

Media Bias & Misinformation, 1787-2018

The importance of mass media to large-scale representative democracy is well-known. At a minimum, the public must have at their disposal the information necessary to reward and punish governments in elections. Ideally, the public has the information required to hold governments accountable and to guide policy change between elections (Franklin et al. 2014). This is of course the objective of a mass media viewed as a Fourth Estate (Schultz 1998).

Given the obvious, long-recognized significance of mass media in large-scale representative democracy, it should come as no surprise that there are substantive bodies of work focused on the problems and/or biases in media coverage of public affairs. Consider for instance the literature on problems with media coverage of public policy (e.g., Dunaway 2011; Lawrence 2000; Baum and Groeling 2010) and scientific issues (e.g., Bennett 1988; Friedman et al. 1999; Schudson 2003; Stocking and Holstein 2009), and on sensationalism and negativity in news content (e.g., Altheide 1997; Cappella and Jamieson 1997; Lichter and Noyes 1995; Patterson 1994; Sabato 1991; Soroka 2014). Consider also work on "indexing" in news content (Bennett 1990), which argues that "that news content on political and public policy issues will generally follow the parameters of elite debate" (Lawrence 2014). A valuable body of recent work confirms that media regularly reflect government and elite sources (e.g., Entman 2003; Vliegenthart and Walgrave 2008; Walgrave and Van Aelst 2006). This likely serves to limit the ways in which media can be usefully critical of current affairs. (That said, indexing may also provide a certain degree of correspondence between news content and policy; it actually may decrease or increase the accuracy of news coverage, conditional on the accuracy of elite messaging.)

Importantly, past work makes clear that some individuals can and do learn about policy from news content (e.g., Barabas and Jerit 2009; Barabas et al. 2014; Druckman 2005; Eveland 2001, 2002; Jerit et al. 2006; Neuner et al. 2019). The implicit finding in this work is that media content will sometimes, and perhaps even often, reflect public policy. However, the story of media influence on public policy issues is not black and white, but gray. Media can, and sometimes do, provide the information required for both electoral accountability and dynamic public responsiveness to policy change. But media also sometimes fail in this regard.

How might we characterize the nature of that failure? There are vast literatures on the ways in which media coverage mischaracterizes minority groups, pays undo attention to violent crime, prioritizes scandal, and so on. Our primary interest in this paper is in media coverage of year-to-year changes in budgetary policy, which helps to narrow the kinds of failure that we can observe.

In some cases, the nature of failure is mainly about the lack of information. There are some policy domains that do not reliably find their way into media coverage. In these cases, public responsiveness seems very unlikely – there simply is no (or little) information to which people can respond.

In other cases, media coverage is a source of misinformation, and consequently misperceptions. Note that we might regard all inaccurate information about policy as misinformation, whether it is purposeful or inadvertent, whether it is expressly false or simply misleading. Misinformation is in our view a lack of correspondence between what media content suggests is happening and what is actually happening. Viewed in this way, misinformation is not a peculiar consequence of the social-media-fueled and politically-polarized information environment. It is and has been a regular feature of mass media coverage since the dawn of mass media.

Note that our definition of misinformation is in line with recent work by Lazer et al. (2018: 1094), who distinguish between fake news ("fabricated information that mimics news media content in form but not in organizational process or intent"), disinformation ("false information that is purposely spread to deceive people") and misinformation ("false or misleading information"). Fake news and disinformation are in this view subsets of the over-arching category of misinformation. There are reasons to believe that social media, and online information more generally, have contributed to a marked increase in the flow of both fake news and disinformation. Increasing political polarization, similarly fueled in part by changes in media technology, likely also contributes to the ready spread of pro-in-partisan/anti-out-partisan fake news and disinformation. These trends have led to a burgeoning and fascinating body of work focused on misinformation and misperceptions. (There is a lot of terrific work in this area, but see, e.g., Allcott and Gentzkow 2017; Bennett and Livingston 2018; Bode and Vraga 2015; Del Vicario et al. 2016; Flynn et al. 2017; Garrett et al. 2016; Nyhan and Reifler 2010; Pasek et al. 2015; Scheufele and Krause 2019; Southwell and Thorson 2015; Tandoc et al. 2018; Thorson 2016; Weeks 2015.)

We take no issue with this literature; indeed, we find ourselves deeply motivated by it. We nevertheless want to emphasize that, because misinformation is not a new phenomenon, there are long-standing bodies of literature, and large bodies of data, that can provide some perspective on past, present, and future trends in misinformation. Our aim in the sections that follow is to develop techniques that identify (in)accuracies in media content. In future work, we (and others) can then use those techniques to better understand the effects of those (in)accuracies, not just for public opinion but for policy as well.

Estimating a "Media Policy Signal"

We focus here on accuracies and inaccuracies in media coverage of defense spending from 1980 to 2018. We do using a method developed in Soroka and Wlezien (2018), and used in both Neuner et al. (2019) and Dun et al. (2019). And we rely on three corpora: one on newspapers, one on television transcripts, and one on public affairs-focused Facebook posts.

The Newspaper Corpus

Our full-text newspaper corpus is drawn from Lexis-Nexis, using the Web Services Kit (WSK), which facilitates the downloading of several hundred thousand stories, formatted in xml, in a single search request. The corpus has been discussed in Soroka and Wlezien (2019), so we will not go into much detail here. We use a search that combined pre-coded subjects and full-text keywords, as follows: STX001996 or BODY(national defense) or BODY(national security) or BODY(defense spending) or BODY(military spending) or BODY(military procurement) or body (weapons spending). (STX001996 is the "National Security" index term, one of five sub-topics with the "International Relations and National Security" topic.)

Our selection of newspapers is based on availability, alongside circulation, with some consideration given to representing different regions. Our working database relies on the following newspapers: *Arizona Republic*, *Arkansas Democrat-Gazette*, *Atlanta Journal-Constitution*, *Boston Globe*, *Chicago Tribune*, *Denver Post*, *Houston Chronicle*, *LA Times*, *Minneapolis Star-Tribune*, *New York Times*, *Orange County Register*, *Philadelphia Inquirer*, *Seattle Times*, *St. Louis Post-Dispatch*, *Tampa Bay Tribune*, *USA Today*, and *Washington Post*. Not all newspapers start in 1980 – most enter the dataset in the early 1990s. All are gathered up to the end of the 2018 fiscal year. These are 17 of the highest-circulation newspapers in the US, three of which aim for national audiences, and seven of which cover considerably large areas in the northeastern, southern, midwestern, and western parts of the country. Combining these newspapers offers, we think, a reasonable representation of the national news stream, at least where newspapers are concerned.

Not all of this content is focused exclusively on defense spending. The analyses that follow are thus not based on articles in their entirety, but rather all *sentences* in this corpus that focus on spending. To be clear: our working database in the analyses that follow is at the sentence level, where sentences are extracted from the larger database using a simple keyword-in-context (KWIC) search identifying all sentences with a keyword related to government spending. We do this using a SPEND dictionary, which includes the following words:

> SPEND: allocat*, appropriation*, budget*, cost*, earmark*, expend*, fund*, grant*, outlay*, resourc*, spend*

This dictionary search (and all subsequent dictionary searches) is implemented in the *quanteda* package (Benoit et al. 2018) in *R*. Note that the dictionary has been subjected to testing in Soroka and Wlezien (2019), and was constructed from our own reading of keyword-in-context (kwic) retrievals, augmented by thesaurus searches.

Our prior work indicates that we can reduce the number of false positives, i.e., sentences that are captured in our search but are in fact not directly related to defense spending, by running a second dictionary search over the set of spending change sentences to confirm that each includes at least one of the following DEFENSE words:

> DEFENSE: army, navy, naval, air force, marines, defense, military, soldier, war, cia, homeland, weapon, terror, security, pentagon, submarine, warship, battleship, destroyer, airplane, aircraft, helicopter, bomb, missile, plane, service men, base, corps, iraq, afghanistan, nato, naval, cruiser, intelligence

Applying this second dictionary to our "spending" corpus somewhat narrows the number of sentences included. In the end, we have 1,775,008 sentences in our newspaper corpus, each of which includes at least one word in the SPEND dictionary, and at least one word in the DEFENSE dictionary.

The Television Corpus

Our corpus of television news broadcasts is also extracted from Lexis-Nexis, again using the Web Services Kit (WSK). Transcripts are held in Lexis-Nexis in a somewhat different way than newspaper articles. In some cases, content is stored at the story level, like newspapers. In other cases, content is stored at the show level, i.e., there is a single transcript for the entire half-hour program. The makes a story-level search relatively ineffective – we end up downloading an entire program, only a small part of which may be on defense spending. Thankfully, our sentence-based analysis for newspapers is still straightforward to implement for television coverage: we download all available content, and then focus on the sentences that match the SPEND dictionary above.

For the three major broadcasters, ABC, CBS and NBC, we download all available content from 1990 onwards in the feature 6pm evening news broadcasts. This is straightforward. The cable news networks, CNN, MSNBC and Fox, do not have a feature news program – indeed, MSNBC and Fox do not clearly separate news from half-hour commentary programs. So we cannot quite get comparable programming from the cable news networks, but we can download all available content, drop infrequent programs, and keep the major programs, e.g., Rachel Maddox, Rush Limbaugh, etc. We do this, and then extract from all transcripts the sentences including at least one work in the

SPEND dictionary and at least one word in the DEFENSE dictionary. This identifies 105,494 sentences over the 1990-2018 period.

Figure 1. Defense Spending Sentences, by Fiscal Year

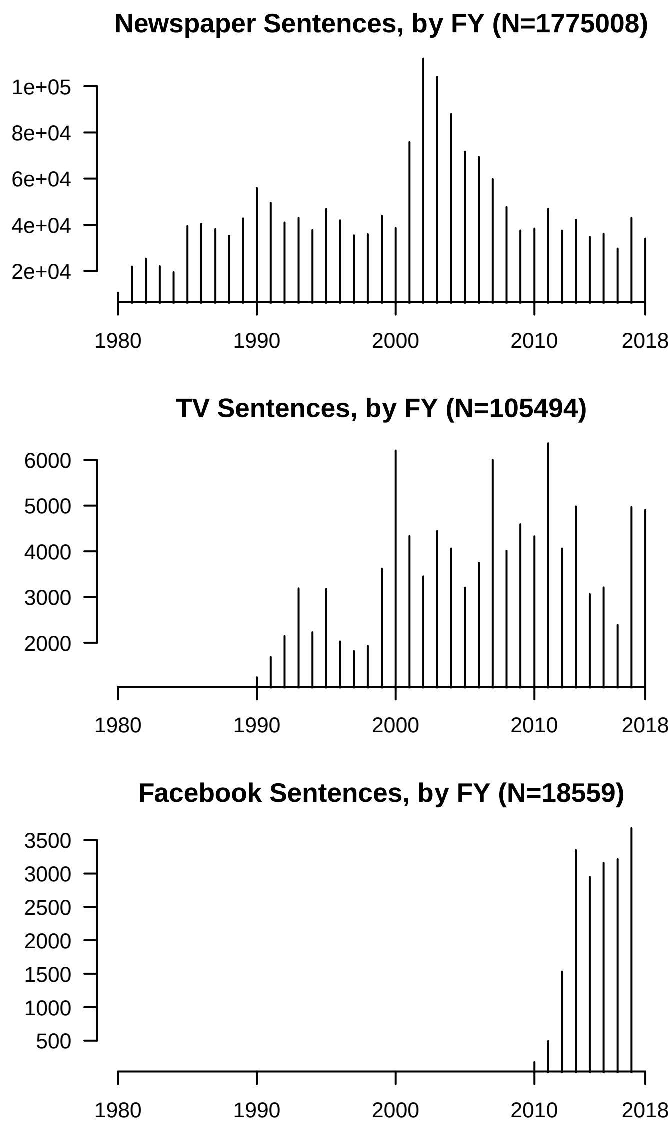

The Facebook Corpus

It has become very difficult to gather data from Facebook. Thankfully, Dan Hiaeshutter-Rice was gathering data on public Facebook pages leading up to, and through, the 2016 election. His database ends in 2017, but goes back to the near-beginning of Facebook, and has been used in Hiaeshutter-Rice and Weeks (N.d.) and Hiaeshutter-Rice (N.d.). Both papers include much more detail about the data and use it much more intensively than is our objective here. Even so, we have kindly been given access to the data and so are able to extend our analyses to Facebook.

The corpus itself includes all posts from the 500 top public-affairs-oriented public pages on Facebook. This includes most major news agencies, alongside a good number of "fringe" news outlets. We analyze them all together here, by extracting sentences in the same way as we do above – we keep all sentences that include at least one word from each of the SPEND and DEFENSE dictionaries. There are 18,559 sentences in total over the 2010-2017 period.

The resulting number of sentences in all three corpora is shown, over time, in Figure 1. The analysis that follow, tracing the accuracy of news coverage of defense spending, are based on this body of data.

The "Media Policy Signal"

Having identified roughly comparable sentences on defense spending across our three corpora, the next task is to estimate a media policy signal, i.e., a measure of the extent to which content suggests increases, decreases, or no change in defense spending. We do so using another set of dictionaries, built and implemented using the same process described above. The dictionaries are as follows:

> UP: accelerat*, accession, accru*, accumulat*, arise*, arose, ascen*, augment*, boom*, boost, climb*, elevat*, exceed*, expand*, expansion, extend*, gain*, grow*, heighten*, higher, increas*, increment*, jump*, leap*, more, multiply*, peak*, rais*, resurg*, rise*, rising, rose, skyrocket*, soar*, surg*, escalat*, up, upraise, upsurge, upward
>
> DOWN: collaps*, contract*, cut*, decay*, declin*, decompos*, decreas*, deflat*, deplet*, depreciat*, descend*, diminish*, dip*, drop*, dwindl*, fall*, fell, fewer, less, lose, losing, loss, lost, lower*, minimiz*, plung*, reced*, reduc*, sank, sink*, scarcit*, shrank, shrink*, shrivel*, shrunk, slash*, slid*, slip*, slow*, slump*, sunk*, toppl*, trim*, tumbl*, wane, waning, wither*

Applying these UP and DOWN dictionaries allows us to assign a direction to the sentences in our database.[1]

(In)accuracies in Media Coverage of Defense Spending

Capturing (in)accuracies in media coverage of defense spending requires several additional steps. First, we must produce a measure of the media policy signal for each corpus. We code sentences in which there are more UP words than DOWN words as "1" and sentences in which there are more DOWN words than UP words are coded as "-1;" other sentences are coded "0." We then sum these values across sentences, by fiscal year (October-September). The resulting measure, while admittedly coarse, nevertheless captures both the direction and magnitude of media coverage (also see footnote 1).

Comparing this measure directly with spending change requires that we place the two series on a similar metric. We use spending in 10,000s FY2000 USD. We then rescale our media policy index so that it has the same mean and range as the spending measure. This is of course just one of many possible approaches to putting the two variables on a similar scale; and different approaches may yield somewhat different results about the match between the two series. For instance, the approach we use here is especially useful in examining the accuracy of year-to-year changes in the media policy signal. If media coverage is perfectly accurate, the signal will be identical to actual spending change.

[1] Note that we may still have irrelevant sentences in our database, but layering dictionaries on top of each other will, we think, produce an increasingly reliable measure. Indeed, our own recent work suggests that this is the case (Soroka and Wlezien 2019). Our approach here is identical to the use of hierarchical dictionary counts as implemented in Young and Soroka (2012) and Belanger and Soroka (2012). We also regard this application of dictionaries as very similar to the "learning" inherent in supervised learning methods used for large-N content analysis (e.g., Jurka et al. 2012; see Grimmer and Stewart 2013 for an especially helpful review.) Our previous research details how the application of successive dictionaries works (Soroka and Wlezien 2019); another paper further highlights the strong connection between this dictionary-based approach and human coding (Neuner et al. 2019); and a still another paper compares the approach with a version augmented by supervised learning guided by thousands of human-coded sentences (Dun et al. 2019). We will not outline the details of these findings here, and ask that readers either consult that work or else accept that our approach yields an accurate estimate of the extent to which sentences – and aggregations of sentences – reflect an increase or decrease in spending on defense.

But this approach ignores the possibility that media coverage is systematically biased (upwards or downwards) on defense spending, because the zero-point is lost in the rescaling described above. Another approach would be to retain the zero-point in the media policy signal, allowing it to be systematically below or above spending change, reflecting a systematic over-reporting (or under-reporting) of either increases or decreases. We do not explore this possibility here, leaving it for future work.

The measures of both spending change and the rescaled media policy signal are illustrated in Figure 2, separately for newspapers, television, and Facebook.[2] The blue line in each panel shows the estimated media policy signal. The light gray line shows actual changes in spending. And the shaded area, the space between the media policy signal and actual spending change, shows the error; put differently, the shaded area represents misinformation.

Note that this gap varies over time. There are years in which media coverage very accurately captures policy change. In these years, were citizens to get a representative sample of media coverage (since presumably no citizen will read *all* defense news), they would have a highly accurate view of the state of defense policy change. There would be, in effect, no misinformation. There are other years in which the story is rather different. In most years there is some gap between media coverage and policy, and in some years that gap is sizeable. In these years, even if citizens received a representative sample of media coverage, they would be likely to have an inaccurate view of the state of defense policy change. News coverage would in this instance be a source of misinformation about the state of public affairs. That said, we suspect that not of all the gaps in Figure 3 are in fact misinformation. It also appears to be the case that media coverage of defense spending is partly forward-looking – in some instances it appears as though the media signal anticipates future changes in spending (see Neuner et al. 2019).[3] It thus may be that part of what we identify as misinformation in Figure 3 is the consequence of media coverage that focuses both on current and future spending. This will be a focus of future work.

In the meantime, note that in many years the scope of misinformation in this domain is relatively small. There are strong correlations between spending and the newspaper signal (r=0.58), the television signal (r=0.53) and the Facebook signal (r=0.80;). The Facebook signal is based only on six years of data, during a strong upward swing in spending, so we do not regard it as directly comparable with the others. Even so, there is evidence here that media report on the direction of spending change on defense. This helps explain past evidence that the US public responds to changes in defense spending (e.g., Wlezien 1995; 1996; Soroka and Wlezien 2010). It also serves as a warning against taking too seriously arguments focused entirely on the failure of mass media. In some domains, at least some of the time, it seems that media coverage can facilitate both electoral accountability and dynamic responsiveness to policy change.[4]

---

[2] Note that we use a recently- updated media database, and rescaled measures, so the values in Figure 2 are slightly different from those in Soroka and Wlezien (2019) and Neuner et al. (2019); and although they will be identical to Dun et al. (2019), our rescaling will slightly alter the illustration of the two series.

[3] Note that this is line with work indicating that economic news coverage is forward-looking; see, e.g, Soroka et al. 2015.

[4] Note that what we show here is likely an underestimate of the link between media coverage and spending change, given that Dun et al. (2019) suggests that the addition of supervised learning methods, on top of dictionaries, increases slightly the correlation between the media policy signal and spending change.

Figure 2. The Media Policy Signal and Defense Spending

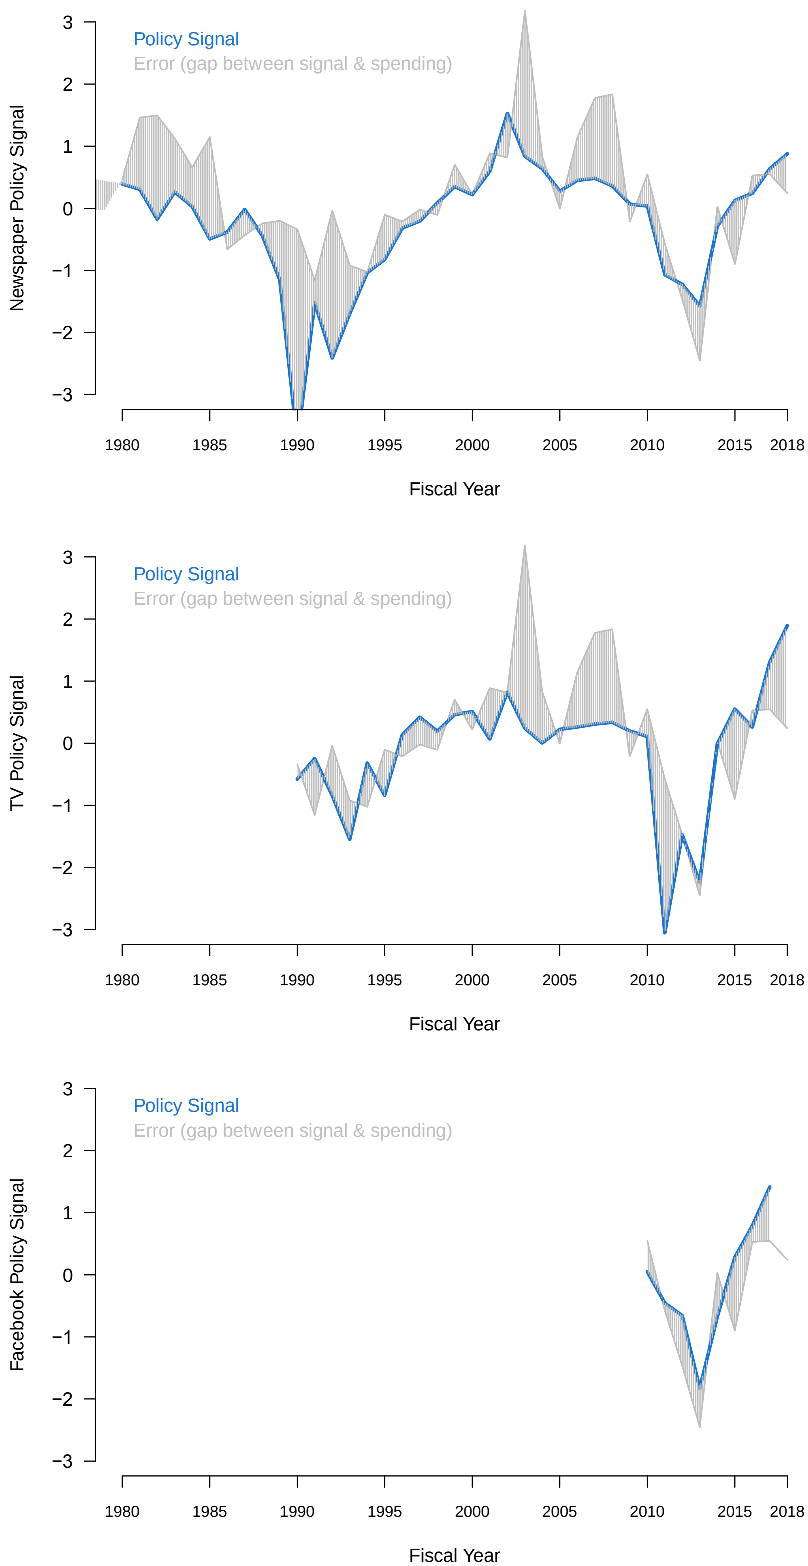

Figure 2 shows results across all 17 newspapers, and 6 television networks, and Facebook posts, combined. There will of course be differences across specific outlets, e.g., the different newspapers, as well, and one strength of our approach is that it facilitates an estimation not just of accuracy for media generally, but for the particular outlets. Table 1 offers a first attempt at characterizing their accuracy.

Table 1. Measuring the Accuracy of Media Outlets

Source	Volume	Direction	Correlation
Atlanta Journal-Constitution	2298.04	0.79	0.64
Arkansas Gazetter	1872.32	0.65	0.39
Arizona Republic	1800.70	0.70	0.75
Boston Globe	3095.97	0.84	0.55
Chicago Tribune	3883.50	0.76	0.44
Denver Post	1485.60	0.72	0.54
Houston Chronicle	2858.71	0.79	0.62
LA Times	5937.06	0.82	0.40
Minnesota Star-Tribune	1055.18	0.68	0.38
New York Times	7642.74	0.82	0.56
Orange County Register	1136.84	0.84	0.44
Philadelphia Inquirer	1978.52	0.72	0.57
St. Louis Post-Dispatch	3109.47	0.90	0.50
Seattle Times	2104.86	0.86	0.70
Tampa Bay Times	2558.84	0.69	0.38
USA Today	1912.13	0.80	0.69
Washington Post	9295.46	0.82	0.58
ABC	332.38	0.69	0.38
CBS	307.36	0.75	0.53
CNN	1705.86	0.76	0.52
FOX	900.24	0.57	0.45
MSNBC	655.79	0.53	0.46
NBC	291.55	0.64	0.47

The data in Table 1 are preliminary -- we expect to improve these measures in the coming months. That said, these are to our knowledge the first measures of by-outlet accuracy of coverage, on defense or other policy areas, over an extended period. We offer three different measures.

Volume is a simple measure of the volume of coverage. Is it the number of sentences that reference defense spending (i.e., have at least one word from each of the SPEND and DEFENSE dictionaries), by fiscal year. We weight by fiscal year since the outlets are not all available over the same number of years, so this weighting make the data somewhat more comparable. Television outlets have markedly fewer sentences that match our criteria, as we should expect. Newspapers like the *Washington Post* and *New York Times* have a lot of content on defense spending in a fiscal year. Where television is concerned, CNN is the clear outlier.

Direction is the proportion of years in which the direction of the media signal matches the direction of spending change. The measure is agnostic about the magnitude of change, just the direction. So if spending is moving upwards (downwards), and the media signal for that fiscal year is above 0 (below 0), then this is a year in which the signal matches spending. For most outlets, this happens most of the time. There nevertheless are outlets for which accuracy, as captured by *Direction*, is nearly a coin flip. For both MSNBC and Fox, *Direction* is just above .5.

Correlation is straightforward – it is the Pearson correlation between the media signal and spending change. The measure thus takes into account both direction and magnitude. It varies from a high of .75 (for the *Arizona Republic*) to a low of .38 (for the *Tampa Bay Times* and ABC). Clearly, there are significant cross-source differences in the accuracy of news coverage of defense spending. That on its own is an important finding, we believe. We consider it below, in light of ongoing concerns about misinformation.

Conclusion and Discussion

What can we draw from the preceding analyses related to current debates about misperceptions and misinformation? The primary general lesson is that the nature and magnitude of (mis)information is in some domains traceable over extended periods of time, across media platforms, and across sources as well. A more specific lesson is that the mass media has, at least in this one domain, provided as much information about policy change as it has misinformation, keeping in mind that some of the disjunctures that we observe between policy and coverage may reflect limitations of our measure of the media signal, not true differences. These results may help explain effective thermostatic public responsiveness to policy change, particularly in high salience domains, where information presumably is plentiful (Soroka and Wlezien 2010). But that responsiveness may be imperfect because the information people receive is itself imperfect. As such, it may be that the variation in public opinion may be best explained by some combination of actual policy change and the other factors that are reflected in the media policy signal (see Neuner, et al 2019).

There is reason to suspect that the proportion of the media policy signal that is due to "other factors" is increasing. Results in Table 1 do not include Facebook, as it is tough to come up with directly comparable metrics. The Facebook signal is, after all, an amalgam of 420 different sources, some but not all of which would show up in a user's feed. Indeed, given interests, "likes," and algorithms, the average Facebook feed won't include a random sample of these sources, but rather a curated and most likely partisan-leaning subsample. In future work we will be able to explore signals based on the partisanship of the source. For the time being, consider the following. If we estimate a Facebook media policy signal based on all 420 sources, the correlation with spending is 0.80. This is higher than any correlation we see using newspapers and television networks in Table 1, though note that Facebook data exist over a comparatively short time period, and the correlations are thus not directly comparable. Even so, we can compare that Facebook signal with another one where Facebook content is weighted by "likes." Calculating this is relatively straightforward: before averaging the direction of sentences by fiscal year, we weight each sentence by the number of likes the post received. Doing so produces a signal that may be more in line with Facebook feeds, prioritizing popular posts (although doing so in this instance across an unlikely feed of 420 news outlets). And this weighting reduces the correlation between the media policy signal and spending, from 0.80 (the highest in our sample) to 0.32 (the lowest in our sample).

User-curated information may thus be markedly less accurate than professionally-curated information, at least for defense policy, and so social media content may reduce thermostatic responsiveness. (Of course, even to the extent the weighted measure better reflects the average user's feed, it is not clear people consider it a reflection of what is happening in the world instead of what should – or should have – happened.) And, as we have seen, there is a fair bit of variation in accuracy even amongst professional news outlets as well. There thus are both glass-half-full and glass-half-empty interpretations of our results.

On the one hand, it has over the past four decades been possible to get a rather accurate sense of changes in defense spending from media outlets. Indeed, many outlets have offered an entirely reasonable account of changes in defense spending. This is good for democratic citizenship and accountability; and it helps explain past findings of thermostatic responsiveness in defense (e.g., Wlezien 1995; Soroka and Wlezien 2010).

On the other hand, there are good reasons for concern about misinformation in news coverage of defense policy – not just recently, but well before the advent of social media. There are professional news outlets that have produced a media policy signal that is only weakly connected to actual policy. Following the *Tampa Bay Times* or the evening news on ABC would produce somewhat different perceptions of year-to-year changes in spending than would following the *Washington Post* or CBS. We have no reason to expect that these differences in coverage are intentional on the part of major news agencies – they likely are barely perceptible to news outlets, which surely do not consider their coverage vis-à-vis spending over decades.[5] Even so, the failure to accurately reflect spending, for whatever reason, has consequences for the quality of public responsiveness and government accountability. Observing those consequences, over decades across policy domains, is not just a reminder that misinformation is not a new phenomenon – it is an opportunity to examine the potential implications of the new forms of misinformation with which citizens are currently confronted. This is one objective of our ongoing work.

[5] This is not to say that there are not systematic biases in coverage, there surely are; see citations above; but that these systematic biases may not differ across outlets, and thus can't account for the differences in coverage observed here.

References

Allcott, Hunt, and Matthew Gentzkow. 2017. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives* 31(2): 211–236.

Altheide, David L. 1997. "The News Media, the Problem Frame, and the Production of Fear." *Sociological Quarterly* 38: 647–668.

Barabas, Jason, Jennifer Jerit, William Pollock, and Carlisle Rainey. 2014. "The Question(s) of Political Knowledge." *American Political Science Review* 108(4): 840–855.

Barabas, Jason, and Jennifer Jerit. 2009. "Estimating the Causal Effects of Media Coverage on Policy-Specific Knowledge." *American Journal of Political Science* 53(1): 73–89.

Baum, Matthew A, and Tim J Groeling. 2010. *War Stories: The Causes and Consequences of Public Views of War*. Princeton NJ: Princeton University Press.

Belanger, E., and S. Soroka. 2011. "Campaigns and the Forecasting of Election Outcomes: Combining Historical and Campaign-Period Prediction Models."

Bennett, W. Lance. 1990. "Toward a Theory of Press-State Relations in the United States." *Journal of Communication* 40(2): 103–127.

Bennett, W Lance, and Steven Livingston. 2018. "The disinformation order: Disruptive communication and the decline of democratic institutions." *European Journal of Communication* 33(2): 122–139.

Bode, Leticia, and Emily K. Vraga. 2015. "In Related News, That was Wrong: The Correction of Misinformation Through Related Stories Functionality in Social Media." *Journal of Communication* 65(4): 619–638.

Cappella, J N, and K H Jamieson. 1997. *Spiral of Cynicism: the press and the public good*. New York: Oxford University Press.

Druckman, James N. 2005. "Media Matter: How Newspapers and Television News Cover Campaigns and Influence Voters." *Political Communication* 22(4): 463–481.

Dun, Linsday, Stuart Soroka and Christopher Wlezien. N.d. "Dictionaries, Supervised Learning, and Media Coverage of Public Policy." *Political Communication*, forthcoming.

Dunaway, Johanna. 2011. "Institutional Effects on the Information Quality of Campaign News." *Journalism Studies* 12(1): 27–44.

Entman, Robert, ROBERT M. 2003. "Cascading Activation: Contesting the White House's Frame After 9/11." *Political Communication* 20(4): 415–432.

Eveland, William P. 2001. "The Cognitive Mediation Model of Learning From the News: Evidence From Nonelection, Off-Year Election, and Presidential Election Contexts." *Communication Research* 28(5): 571–601.

Flynn, D. J., Brendan Nyhan, and Jason Reifler. 2017. "The Nature and Origins of Misperceptions: Understanding False and Unsupported Beliefs About Politics." *Political Psychology* 38(S1): 127–150.

Franklin, Mark, Stuart Soroka, and Christopher Wlezien. 2014. "Elections." In *The Oxford Handbook Public Accountability*, eds. Mark Bovens, Robert E. Goodin, and Thomas Schillemans. OUP Oxford, p. 389–404.

Friedman, Sharon M., Sharon Dunwoody, Carol L. Rogers, Sharon Dunwoody, et al. 2012. *Communicating Uncertainty : Media Coverage of New and Controversial Science*. Routledge.

Garrett, R. Kelly, Brian E. Weeks, and Rachel L. Neo. 2016. "Driving a Wedge Between Evidence and Beliefs: How Online Ideological News Exposure Promotes Political Misperceptions." *Journal of Computer-Mediated Communication* 21(5): 331–348.

Grimmer, Justin, and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21(3): 267–297.

Hiaeshutter-Rice, D., Weeks, B. E. N.d. "The Nature of News on Facebook."

Hiaeshutter-Rice, D. N.d. "Political Platforms: Technology, User Affordances, and Campaign Communications."

Jerit, Jennifer, Jason Barabas, and Toby Bolsen. 2006. "Citizens, Knowledge, and the Information Environment." *American Journal of Political Science* 50(2): 266–282.

Jr, Edson C. Tandoc, Zheng Wei Lim, and Richard Ling. 2018. "Defining 'Fake News.'" *Digital Journalism* 6(2): 137–153.

Lawrence, Regina G. 2000. "Game-framing the issues: Tracking the strategy frame in public policy news."

Political Communication 17(2): 93–114.
Lawrence, Regina G. 2014. "Indexing." http://www.oxfordbibliographies.com/view/document/obo-9780199756841/obo-9780199756841-0090.xml (Accessed March 28, 2019).
Lazer, David M. J. et al. 2018. "The science of fake news." *Science* 359(6380): 1094–1096.
Lichter, S. Robert, and Richard E. Noyes. 1996. *Good Intentions Make Bad News: Why Americans Hate Campaign Journalism*. Lanham, MD: Rowman & Littlefield Publishers.
Neuner, Fabian, Stuart Soroka and Christopher Wlezien. 2019. "Mass Media and Public Responsiveness." *International Journal of Press/Politics* 24(3):269-292.
Nyhan, Brendan, and Jason Reifler. 2010. "When Corrections Fail: The Persistence of Political Misperceptions." *Political Behavior* 32(2): 303–330.
Pasek, Josh, Gaurav Sood, and Jon A. Krosnick. 2015. "Misinformed About the Affordable Care Act? Leveraging Certainty to Assess the Prevalence of Misperceptions." *Journal of Communication* 65(4): 660–673.
Patterson, T E. 1994. *Out of Order*. New York: Vintage Books.
Sabato, Larry. 1991. *Feeding Frenzy: How Attack Journalism Has Transformed American Politics*. New York: The Free Press.
Scheufele, Dietram A., and Nicole M. Krause. 2019. "Science audiences, misinformation, and fake news." *Proceedings of the National Academy of Sciences*: 201805871.
Schudson, Michael. 2002. "The News Media as Political Institutions." *Annual Review of Political Science* 5(1): 249–269.
Schultz, Julianne. 1998. *Reviving the Fourth Estate: Democracy, Accountability and the Media*. Cambridge University Press.
Soroka, Stuart N. 2014. *Negativity in Democratic Politics: Causes and Consequences*. New York, NY: Cambridge University Press.
Soroka, Stuart N., Dominik A. Stecula, and Christopher Wlezien. 2015. "It's (Change in) the (Future) Economy, Stupid: Economic Indicators, the Media, and Public Opinion." *American Journal of Political Science* 59(2): 457–474.
Soroka, Stuart and Christopher Wlezien. 2019. "Tracking the Coverage of Public Policy in Mass Media." *Policy Studies Journal* 47(2): 471-491.
Soroka, Stuart N., and Christopher Wlezien. 2010. *Degrees of Democracy: Politics, Public Opinion and Policy*. Cambridge MA: Cambridge University Press.
Southwell, Brian G., and Emily A. Thorson. 2015. "The Prevalence, Consequence, and Remedy of Misinformation in Mass Media Systems." *Journal of Communication* 65(4): 589–595.
Stocking, S. Holly, and Lisa W. Holstein. 2009. "Manufacturing doubt: journalists' roles and the construction of ignorance in a scientific controversy." *Public Understanding of Science* 18(1): 23–42.
Thorson, Emily. 2016. "Belief Echoes: The Persistent Effects of Corrected Misinformation." *Political Communication* 33(3): 460–480.
Vicario, Michela Del et al. 2016. "The spreading of misinformation online." *Proceedings of the National Academy of Sciences* 113(3): 554–559.
Walgrave, Stefaan, and Peter Van Aelst. 2006. "The Contingency of the Mass Media's Political Agenda Setting Power: Toward a Preliminary Theory." *Journal of Communication* 56(1): 88–109.
Weeks, Brian E. 2015. "Emotions, Partisanship, and Misperceptions: How Anger and Anxiety Moderate the Effect of Partisan Bias on Susceptibility to Political Misinformation." *Journal of Communication* 65(4): 699–719.
Wlezien, Christopher. 1996. "Dynamics of Representation: The Case of US Spending on Defense." *British Journal of Political Science* 26:81-103.
Wlezien, Christopher. 1995. "The Public as Thermostat: Dynamics of Preferences for Spending." *American Journal of Political Science* 39:981-1000.
Young, Lori, and Stuart N. Soroka. 2012. "Affective News: The Automated Coding of Sentiment in Political Texts." *Political Communication* 29: 205–231.