On the Sources of Public Responsiveness to Policy[1]

Stuart Soroka
University of Michigan
ssoroka@umich.edu

Christopher Wlezien
University of Texas at Austin
Wlezien@austin.texas.edu

Abstract: There is a sizeable literature finding evidence of thermostatic responsiveness across a range of salient policy domains. We thus far have only a partial sense for what drives responsiveness to policy change at the individual level, however. One possibility is that individuals learn what they need to know from the mass media, but there is very little work exploring the prevalence of relevant policy cues in media content. This paper is a first effort at doing exactly this. We use automated content-analytic techniques to look at 35 years of reporting on both defense and welfare in the *New York Times* and the *Washington Post*. Cues about the direction in which spending is moving are surprisingly frequent, and in line with the direction of spending, especially in defense.

Democracy depends on an attentive and informed public. There is a growing body of literature suggesting that publics regularly respond thermostatically to policy change; at the same time, there is work suggesting that the average citizen has, at best, only a very partial understanding of politics and policy.[2] How can both of these be true? One prominent account focuses on the advantages of aggregation – although individuals make mistakes, those errors cancel out in the aggregate. We want to focus here on another possibility, namely, that thermostatic responsiveness does not actually require very detailed information about policy, and much of what is required is readily evident in media content. After all, some people must be receiving fairly accurate information about policy change even for the miracle of aggregation to work,[3] i.e., there is a responsive signal left after the noise cancels out.

To be clear: we do not argue that citizens learn about spending by reading budgets or even by reading news stories about budgets. We suggest, rather, that broad shifts in policy may be captured in news content, that citizens may thus be able to learn relatively easily about the general direction of policy change, and that this is all that is required for effective thermostatic responsiveness. (Indeed, the information requirements necessary for the public to respond thermostatic to policy actually are not that great. See, e.g. Soroka and Wlezien 2010: 41-42.)

This paper offers a first test of the possibility that media helps drive thermostatic responsiveness, focusing on a large-scale automated content analysis of thirty-five years

[1] This paper relies on on Lexicoder software development by Mark Daku and Lori Young; a good deal of database work from Dan Hiaeshutter-Rice (PhD candidate, Michigan); contributions on a related project from Fabian Neuner (PhD candidate, Michigan); and coding by undergraduates students Heidi Payter, Amanda Hampton, Thomas Donovan (Michigan), and Sean Riordan and Kayla Torres (UT Austin).

[2] As regards thermostatic responsiveness, see, e.g., Wlezien (1995; 1996), Erikson, et al (2002), Soroka and Wlezien (2005), Jennings (2009), Soroka and Wlezien (2010), Wlezien and Soroka (2012). As regards the public's inattentiveness and information, see, e.g., Berelson, et al (1954), Converse (1964; 1970), Bennett (1988), Page and Shapiro (1992), Delli Carpini and Keeter (1996), Popkin and Dimock (1999).

[3] Though see Althaus (2003) for a more complicated and less complimentary view of aggregation.

of news stories in the *New York Times* and *Washington Post*, and exploring the availability of cues about increases or decreases in spending commitments. We do not examine citizens' interest in or understanding of these cues at present, but take a first step, and simply look at the extent to which cues about policy direction are available in media content. We then compare the "signal" provided by these cues to trends in government spending; and consider this a preliminary examination of the possibility that thermostatic responsiveness is facilitated by media content.

Results suggest that in each of defense and welfare cues about the direction in which government policy are readily available. Indeed, they are so readily available that even a cursory engagement with news content seems likely to give citizens the information they require to make feasible thermostatic responsiveness, and representative democracy more generally. This is not universally true – our analyses show a "media signal" in defense that is impressively similar to trends in actual defense spending; but results in welfare are weaker. That said, the work that follows is very preliminary. We anticipate working with much more detailed data, and content analyses, in the future. We thus consider this an initial exploration into the feasibility of much broader study into micro-level explanations for thermostatic responsiveness – and one that points to the (perhaps surprising) availability of policy cues in media content, as well as the potential for this line of analysis on a much broader scale.

Mass Media & Thermostatic Responsiveness

If there is thermostatic public responsiveness, people will adjust their preferences for "more" or "less" policy in response to policy change, favoring less (more) policy in the wake of policy increases (decreases), other things being equal. Formally, the public's relative preference (R) would represent the difference between the publics preferred level of policy (P^*) and policy (P) itself:

$$R_t = P_t^* - P_t, \tag{1}$$

Where the subscripted t indicates time.[4] R can change because either P or P^* changes. Because we often do not have measures of P^* and also because the three variables are not measured using the same metric, it is necessary to rewrite the equation as follows:

$$R_t = B_1 U_t + B_2 P_t + e_t, \tag{2}$$

where U is a set of additional exogenous predictors of P^*.

If the coefficient for policy (B_2) in equation 2 is less than 0, we have evidence of thermostatic responsiveness. This often is true but it sometimes is not (Wlezien 1995). Moreover, even when the coefficient is less than 0, its size varies, owing partly to characteristics of issues and partly to political institutions (Soroka and Wlezien 2010; Wlezien and Soroka 2012). First, the public is more responsive in highly salient

[4] The equation can apply across space and time, though the latter is more common and also the focus of the analysis in this paper.

domains. This is highly intuitive. Second, political context matters. In particular, federalism dampens public responsiveness, seemingly because it lowers the clarity of information about what policymakers are doing at different levels of government.

Thermostatic public responsiveness clearly is important. It provides the basis for holding policymakers accountable. It also provides informed signals that policymakers could effectively represent. Establishing that there is thermostatic public responsiveness, and understanding how it varies across issues and political context, is thus of some significance. There is, as we have noted, a growing body of work that does exactly this. (See citations in note 2.) But there is relatively little work that seeks to understand the mechanism(s) driving the relationship. We know that most people pay little attention to politics and have little knowledge about it. How do people get the information? How is it that the patterns are so strikingly similar across various subgroups, including those with low levels of education and political interest?

What do we know? To begin, we know that thermostatic responsiveness does not require a high level of information. This must be true – otherwise we cannot account for the patterns that we (and others) observe. We have posited elsewhere (Soroka and Wlezien 2010) that people only need to have a sense for the direction of policy change – whether policy has gone up or down – and perhaps also the magnitude – whether it has gone up by a little or a lot. These are very basic cues. But even these cues need to come from somewhere.

We know that the information is mediated in some way. Again, thus must be true. People do not directly observe what policymakers do in most domains. They also don't have copies of the federal budget (or other policies) on their coffee tables. Information about what policymakers do is thus conveyed to the public by some other means. But how exactly? Answering this question obviously is important unto itself, as it also helps us understand how the public learns about policy, which is critical to making representative democracies work.

The most likely suspect is mass media. Although the structure and form has changed substantially over time, the public still relies primarily on large-scale media organizations for its information about politics.[5] That said, the prospect that mass media provide the information that drives the thermostatic model may seem at odds with common critiques of media content. There are considerable bodies of work detailing a range of biases in media content,[6] and a good deal of work lamenting a lack of policy content specifically (e.g., Lawrence, 2000). There is a considerable body of work, in short, suggesting that media content offers a barely perceptible and systematically biased view of public policy. That said, just as we argue that citizens may not need much information to form general opinions, we argue that media can be inefficient and biased in many different ways, but still provide the basic information citizens need to assess the direction of policy change, at least in very salient policy areas that attract a lot of public, political and media

[5] See, for instance, http://www.people-press.org/2012/02/07/cable-leads-the-pack-as-campaign-news-source/. Of course, citizens may also rely on "new" media for policy information; and these too could facilitate thermostatic responsiveness, particularly if people's various new media sources also track policy over time.

[6] For instance, consider work on sensationalism and/or negativity in news content, e.g., Altheide 1997; Davie and Lee 1995; Lichter and Noyes 1995; Meyrowitz 1985; Patterson 1994; Sabato 1991; Soroka 2014.

attention. Sceptics abound, we suspect. But there is already work suggesting that people can and do learn about policy when there is sufficient media coverage (see Barabas and Jerit 2009 for both a review of the field, and a new study); and, as these authors note, existing work also identifies some areas in which the public actually has relatively high policy-specific knowledge (e.g., Delli Carpini and Keeter 1996).

If mass media actually are driving thermostatic public responsiveness, we will observe two patterns. First, obviously, media coverage will reflect what happens to policy. At minimum, we will observe a positive correlation between policy changes and the "signal" in the mass media. In the extreme, there will be perfect congruence between the two. But there is likely to be some degree of mismatch between policy and the media coverage; and the magnitude of the difference is important, particularly to the extent the public responds to media coverage. It also is something we can directly assess. Second, and quite relatedly, if mass media are driving thermostatic responsiveness, public opinion will reflect the media's policy signal. This also is something we can directly assess, though it is tricky, particularly to the extent that there is a close match between policy and the media signal over time. How can we determine that the public is following the media signal per se? To do this, we actually need to observe disjunctures between media and policy, then assess whether the public follows the media signal or policy itself.[7]

Data

We focus here on news content in the *New York Times* and *Washington Post*, from 1980 to the present. Media data are drawn from Factiva full-text indices. For now, we focus on two policy domains: defense and welfare. We select these because they are highly salient domains in which there is ample evidence of thermostatic public responsiveness (e.g., Wlezien 1995, 2004; Soroka and Wlezien 2010). These are domains in which public responsiveness is readily evident, and thus requires some explanation. We rely on Factiva subject keywords to identify articles: GWELF, and GDEF. These topics were selected based on preliminary searches; they tend to be policy-focused (rather than just focused on wars in the case of GDEF, for instance).

The use of Factiva keywords by no means perfectly captures only relevant articles. We invariably miss some relevant articles; and our analyses identify a good volume of irrelevant material captured in our searches as well. The keyword search is nevertheless more efficient than a full-text keyword search.[8] Of course, Factiva's assignment of topics is most likely a function of their own dictionary-based word search, but our assumption is that their search is more developed than ours would be. We suspect that our use of their keywords means that we err on the side of Type I rather than Type II errors. That said, we have to sort through irrelevant material as well, and one aim of the analyses that follow is to improve our ability to identify relevant material.

---

[7] Even if the public follows the former, we cannot easily rule out the other main possibility – that new/alternative media are the mechanisms of responsiveness and the aggregate signal from these sources parallels the traditional media signal. Still, as noted above, while this is possible, it is far less likely, i.e., are priors are *much* lower; and this clearly cannot account for responsiveness before about 2000.

[8] Consider what we might get when searching for articles using the word "welfare," for instance – a set of articles dealing not just with social welfare, but also child welfare, animal welfare, etc.

A more critical issue with subject codes in Factiva is that they have not been applied to all years – there are in particular years in the late 1990s when no subject codes exist for either of the newspapers that we focus on, and we thus have missing data in these years. For this preliminary work, we simply work around these missing data. That said, we do not lack for media content. Our database includes roughly 16,600 welfare articles and 52,800 defense articles. The analyses that follow are thus based on nearly 70,000 articles in total. The basic breakdown of articles is shown in Table 1, which reports the total number of articles downloaded by topic, newspaper and decade.

Table 1. Sample of News Articles, by Topic, Decade and Newspaper

	Defense		Welfare	
	NYT	*WPost*	*NYT*	*WPost*
1980s	15312	0	7415	0
1990s	3382	14	3885	0
2000s	10667	10346	2237	2063
2010s	5572	5807	580	424

The table makes clear the gaps that result from the intermittent non-existence of subject codes in the Factiva database. There are many more defense articles than welfare articles (at least based on these Factiva keywords); and we have no *Washington Post* articles before 2000. (There are also articles missing from Table 1 due to miscodes for date and/or news source, but these cases are relatively rare.) For the time being, however, we work with these data, with the belief that they will give us a sufficient sense for whether policy cues are available, and retrievable, in media content.

Content Analysis

Identifying Relevant Mentions

We analyze these data using two packages designed for automated content analysis in R. One is relatively common, the tm package. The other, Lexicoder 3.1, was released several years ago as a Java-based utility for basic large-scale content analyses (see, .e.g, Young and Soroka 2012; Soroka 2006, 2012). We rely here on the newest implementation, still in Java, but built to be run from the command line and thus also readily accessible within R. (The new version of Lexicoder was publicly released in late-August, 2015, and is readily available to academic researchers at www.lexicoder.com.)

We begin with a simple text-cleaning function in Lexicoder that removes punctuation and other strange characters, and changes all words to lower case. Subsequent analyses rely on both word-count and dictionary-count functions. We begin these with a simple dictionary search for terms related to spending. That dictionary is as follows:

```
allocate*, appropriation*, budget*, cost*, earmark*, expend*, fund*, grant*, outlay*,
resourc*, spend*
```

The first rows of Table 2 show the number of articles that include (or do not include) at least one of these spending keywords. The defense search clearly casts a wider net than does the welfare search. Even so, we have roughly 20,000 defense articles that include one spending keyword, and 11,000 similar articles in welfare. These form the body of data on which we will focus our analyses below.

Table 2. Frequency of Spending Mentions

	Defense	Welfare
Articles with no spending mention	31534	5216
Articles with at least one spending mention	21271	11388
Number of sentences mentioning spending	*68873*	*55003*

The final row of Table 2 shows a different quantity: the number of sentences, rather than articles, mentioning spending. These sentences are identified using the Lexicoder *kwic* (keyword in context) function, which both identifies and then extracts (into a separate database) all instances in which a given set of keywords is used. Note that the *kwic* function captures not just the keyword, but the surrounding *N* words, up to the beginning or end of the sentence. We set a wide, 30-word window here, and in so doing capture the entire sentence in which a set of spending-related keywords are used. This results in a separate database of 123,876 such instances. Below, we rely on a combination of the article-level and sentence-level datasets.

Note that our data reflect a potentially large understatement of the frequency of policy information in media content. There will be many articles that deal with policy but do not mention spending explicitly, for instance; and (relatedly) there are many aspects of policy that are not about spending. That said, our corpuses focus on a subset of policy-relevant information that passes a relatively high bar information-wise. For those who suspect that media do not provide any information about the direction of spending, the size alone of the *kwic* dataset should be striking. Even ignoring decades of missing data, our corpus amounts to just under 10 spending-related sentences on these two topics every day for the past 35 years.

Table 3. Frequency of Spending Keywords

	Keyword Frequency	
	Defense	Welfare
allocate*	239	317
appropriation*	1859	438
budget*	17614	10889
cost*	18354	14121
earmark*	188	167
expenditure*	813	430
fund*	7633	13605
grant*	2008	4474
outlays*	420	310
resource*	5081	3649
spend*	14662	6602

Table 3 provides some basic diagnostics about the frequency of spending keywords across issues. There are some common trends, i.e., the words more frequent in the welfare domain tend to also be more frequent in defense. We take this as evidence that these are the words most relevant to spending beyond these domains as well; though we do not yet explore that possibility here.

Whether these spending mentions (a) actually deal with government spending and/or (b) contain any useful / accurate information about policy change is another matter. We explore this by manually coding a subset of *kwic* retrievals as either (a) clearly related to

spending in welfare or defense, or (b) irrelevant to spending in these domains.[9] Our coding began with a random sample of 350 retrievals in each topic. We then supplemented that sample with additional random draws of articles using specific spending keywords, to ensure that there is no single keyword for which we have fewer than 50 coded retrievals per topic. (This was necessary because some spending keywords are used relatively infrequently, and thus do not appear often in our initial random sample.)

Table 4. Relevance of Spending Keywords

	Alone		With direction		With Relevance Keywords	
	N	Relevant	N	Relevant	N	Relevant
allocate*	100	49.0%	39	48.7%	64	67.2%
appropriation*	101	49.5%	37	62.2%	87	54.0%
budget*	179	81.6%	101	88.1%	179	81.6%
cost*	177	49.7%	88	53.4%	94	71.3%
earmark*	100	70.0%	42	64.3%	73	83.6%
expend*	99	67.7%	53	75.5%	74	77.0%
fund*	132	40.2%	51	35.3%	67	64.2%
grant*	100	25.0%	48	29.2%	57	35.1%
outlay*	101	82.2%	57	78.9%	85	83.5%
resourc*	100	23.0%	47	23.4%	43	46.5%
spend*	126	72.2%	78	78.2%	96	82.3%

The first two columns of Table 4 show the total number of coded (relevant and irrelevant) *kwic* retrievals by spending keyword, followed by the percent of relevant retrievals. Some keywords are very reliably related to spending, such as budget*, outlay*, earmark*, and spend*. Of the 179 coded *kwic* retrievals containing the word budget*, for instance, 82% are clearly related to federal government spending in either defense or welfare. Others, like resource*, are much less reliably related. Indeed, the non-reliability of resource* and grant*, both of which are relevant roughly 25% of the time, points towards the advantages of simply dropping these keywords from subsequent analyses.

Note that our analyses do not focus on *any* mention of spending, however – our media signal relies on there being information about the direction in which policy is moving. The clearest versions of this are instances in which a spending keyword co-occurs with a direction keyword. For instance, spend* occurring alongside "more" or "less" offers very clear information about the direction of fiscal policy. We accordingly add to our analyses the following dictionaries for direction, constructed from our own reading of *kwic* retrievals, augmented by thesaurus searches:

```
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*
```

[9] Coding is currently based on single expert coder.

Not all of these words will be used in relation to spending, of course; just as not all spending mentions will co-occur with direction keywords. But we start by casting a relatively wide net in both instances; and expect that the concurrence of these two dictionaries will identify sentences that are related to policy direction in each domain. That said, the middle two columns of Table 4 show only a marginal improvement in relevance scores when we focus only on the *kwic* retrievals that contain both spending and directional keyword. Directional keywords are obviously important – these are the means by which readers will learn about policy change. But they do not easily solve the issue of relevance in automated keyword searches.

The final column of Table 4 adds one more test. Here, we focus on relevance scores for *kwic* retrievals that contain one of the following 'relevance' keywords:

```
senate, congress*, president, governor, bush, reagan, clinton, obama, republican,
democrat*, federal, propos*, request*, approv*, keep, get, future, next, estimat*,
domestic, national, economy, annual, program*, tax, bill, budget*, revenue, fiscal,
deficit, debate, number, percent, billion*, million*, dollar*, pentagon, military,
aircraft, medicare, social
```

This set of words is determining as follows: we build a "term document matrix" (using the tm package in R) of all words used in all human-coded *kwic* retrievals. We then look at the likelihood that a *kwic* retrieval is coded as relevant based on the occurrence of each word. The above list includes all words for which 80% of coded retrievals are coded as relevant (excluding words already in the spending or direction dictionaries). The list is derived statistically, then, but it makes good sense: it includes a set of institutional words (and presidents' names), words related to budgetary policy and decision-making, words related to numbers, and finally some words specific to the defense and welfare domains.

If we focus on relevance when these words are included, as in the last columns of Table 4, we start to identify a more reliably relevant set of policy cues. There are now four spending keywords with relevancy scores over 80%, for instance. We take this as a good sign that we can further improve the identification of relevant policy cues in future work. Doing so is an important step in identifying just how many policy-related cues are available in media content, and in estimating the "policy signal" provided to media consumers.

Extracting the "Policy Signal" in Media Content

Focusing on co-occurrences of spending and direction keywords leaves us with roughly 63,000 sentences in our database that suggest (albeit with varying degrees of relevance) the direction of fiscal policy.[10] We aim here to convert these sentences into a "policy signal," and to compare that signal with actual budgetary data. Doing so offers a relatively simple test of the extent to which media provide an accurate signal about what governments are doing (/spending).

Our measure of budgetary policy come from the Policy Agendas' database of appropriations (available at policyagendas.org). These data have the advantage of using functional definitions that are more temporally consistent than the standard OMB

[10] More specifically, of the 68,873 sentences on defense spending, 35,207 also include direction keywords; of the 55,003 sentences on welfare spending, 28,350 including spending keywords.

Historical Tables.[11] We distinguish spending on each of welfare and defense using the definitions in Wlezien (2004) and Soroka and Wlezien (2010); in each case, we rely on constant dollars (controlling for inflation).

Converting our dictionary counts into a media signal is not straightforward. There is a simple route, however, which we take here: we use co-occurring spending and directional keywords to attribute one of two codes to every retrieval, (a) increasing spending (+1), or (b) decreasing spending (-1); and then plot the sum of all mentions, aggregated by fiscal year, over the 35-year time period. We focus for the time being on all co-occurrences of spending and direction keywords; with the exception of grant* and resource*, which we drop based on relevancy results above.

This approach is far from perfect. Perhaps most importantly, it takes the frequency of co-occurrences as an indication of magnitude. A month in which there are many co-occurrences in a positive direction will show a strongly positive signal, while a month in which there are only a few co-occurrences in a positive direction will show a weakly positive signal, for instance. There is however no reason to expect that the magnitude of spending change will be systematically matched to the number of mentions of upward or downward movement. A highly salient but fiscally small change in policy is a problem for our current measure, then. We thus do not expect a perfect correspondence between our current media signal and fiscal policy. But we regard this work as an initial test only. Is there *any* correspondence? Are there any hints in these data that media content captures over-time trends in policy?

We try to answer these questions using an approach that explores two possibilities: (1) the direction of spending change is reflected in the direction of the media signal, or (2) both the direction and magnitude of spending change is reflected in the media signal. The first is an easier test – it asks only whether the directions in spending and media match, regardless of the magnitude of either measure. The second is much tougher, and as we have noted above, may be too much to ask of our relatively simple count-based 'media signal.'

Figure 1 allows us to easily see relationships between the two series in terms of both direction and magnitude. In every panel of Figure 1, the x-axis shows the upward or downward change in appropriations while the y-axis shows upward or downward change suggested by the aggregation of media content. Each dot shows one fiscal year, and the dotted lines show the zero-point for both axes. Consider first a test focused on direction: to the extent that media content points in the same direction as actual spending, we expect dots to appear in either the bottom left quadrant (where spending is decreasing, and media content suggests a decrease), or in the upper right quadrant (where spending is increasing, and media content suggests an increase). Dots in other quadrants signal years in which the media signal is in conflict with the direction of spending. Of course, the figure allows for a test focused on magnitude as well: to the extent that media content mirrors spending in both direction and magnitude, we expect a strong (diagonal) relationship between the two series.

[11] OMB classification actually changes over time, though probably less often than in other countries, where reliable spending data can be even more difficult to identify (Soroka, Wlezien and McLean 2006).

Figure 1. Changes in Defense Appropriations and Media Cues

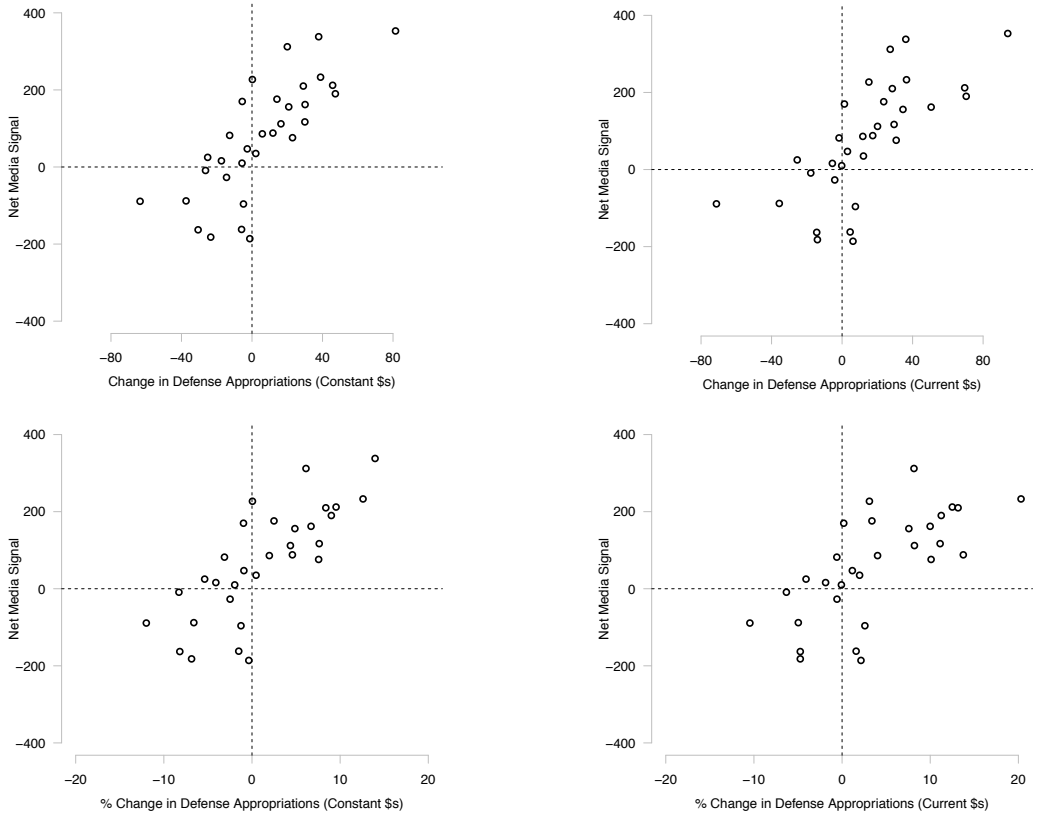

Figure 2. Trends in Defense Appropriations and Media Cues

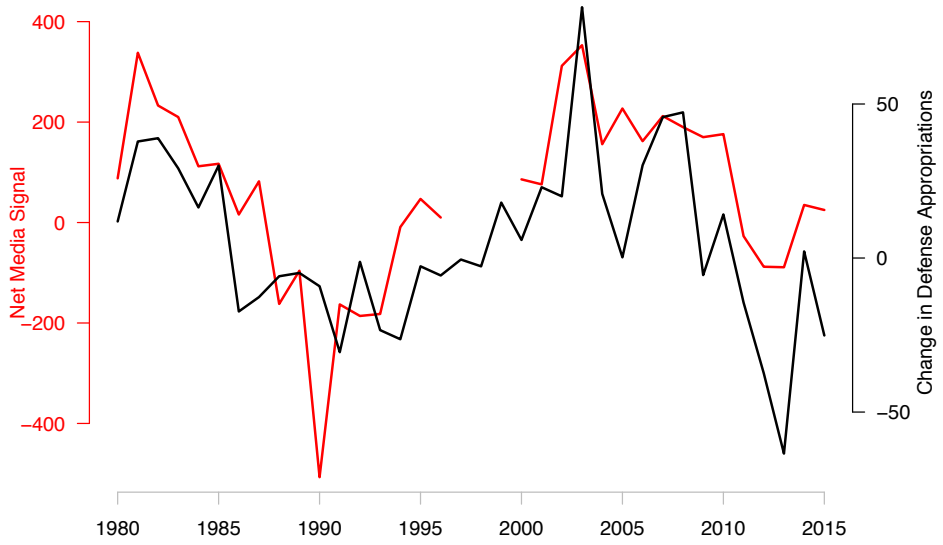

There is a rather strong relationship between the series in terms of magnitude as well: the figure shows what we regard as rather striking correlations between the direction of spending change and the direction of our media signal, with correlations of .68 and .62 in the top row, and .71 and .67 in the bottom row. There is no strong signal here that media content follows current versus constant dollars, or percentage versus absolute change.

But it is clear that media cues are pretty consistent with the direction of spending change, both in terms of direction and magnitude.

The correlation is clearer still when we look at the series plotted over time. The black line in Figure 2 shows changes in defense appropriations, in billions of constant dollars, from 1980 to the present. (These are the spending data used in the top left panel of Figure 1.) The red line shows our basic media "policy signal" over the same time period (with missing data in quarters for which Factiva stories are missing subject codes). The year 1990 appears to be a real outlier where media content is concerned – we cannot yet explain why – but the figure certainly confirms what the scatterplot in Figure 1 suggests: media content includes a good number of relatively reliable cues about the direction of government spending. Note that our results hold when we focus on actual expenditures instead of appropriations as well (see Appendix Figure A1); though there are one or two more years in which the direction of media content differs from the direction in outlays.[12] In sum: media consumers will be exposed to cues that give them, on average, a good sense for the direction (and magnitude) of spending change.

Figure 3. Changes in Welfare Appropriations and Media Cues

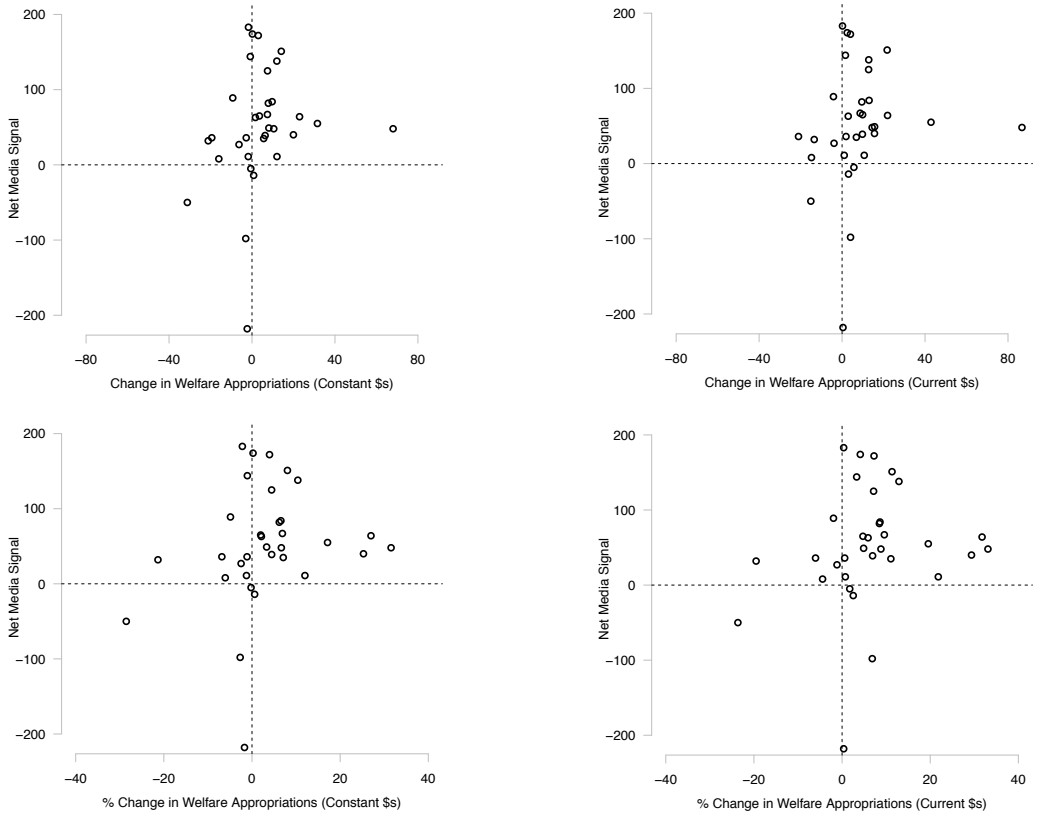

[12] This is as we might expect, given that our past work suggests that public preferences respond more to appropriations than to outlays (see Wlezien and Soroka 2003) and very early in the fiscal year, before much spending has occurred (Soroka and Wlezien 2010). To be clear: if media play a role in informing the public about policy change, it appears here as though that signal is more attuned to appropriation than to outlays.

This is true in the defense domain, at least. Is this also the case for welfare? Figures 3 and 4 present the results. These are not at all as clear-cut as for defense – the connection between our estimated media signal and spending change is without question weaker in welfare. That said, the number of years in which the media signal points in the opposite direction as spending change is relatively limited. The panels in Figure 3 point to six or seven years in which the direction of the media signal is at odds with the direction of spending. This is actually very close to what we saw in defense. Indeed, viewed in this way, the media signal in the welfare domain is relatively accurate.

Figure 4. Trends in Welfare Appropriations and Media Cues

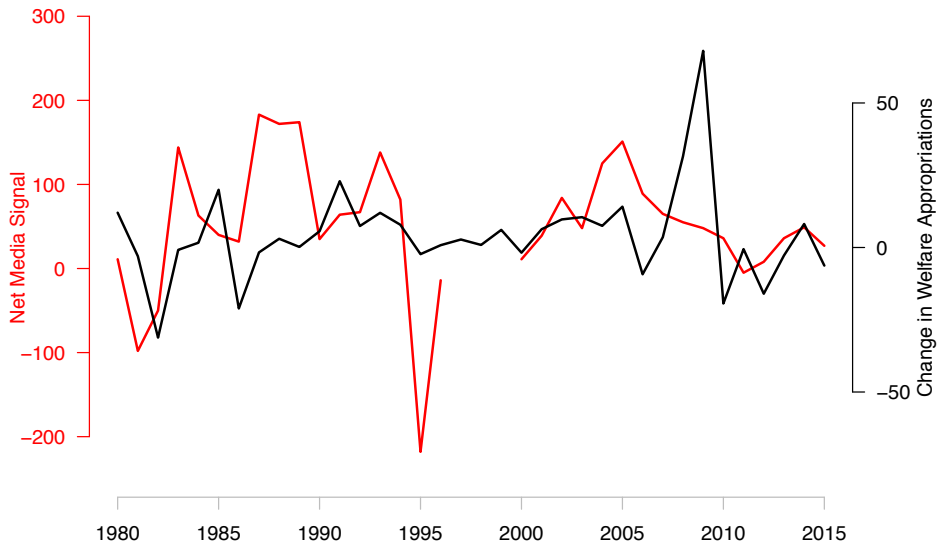

Taking the magnitude of change into account makes for a rather different story. In comparison with correlations hovering near .7 in defense, the correlations between spending and the estimated media signal are .19 and .14 in the top row of Figure 3; .22 and .19 in the bottom row. The corresponding over-time trends for welfare are shown in Figure 4, using the same spending series as in the top left panel of Figure 3. The relatively weak relationship between the two series is clear here – there are some common trends, but the over-time relationship between the two series is noisy at best. (Again, much the same is true using expenditures; see Appendix Figure A2.) Assuming our identification of cues is equally strong in the two domains, it appears as through the media signal is much clearer for defense than for welfare. But, of course, we cannot be sure that our content analysis is as strong in welfare as it is in defense. We take up this issue, and others, below.

Discussion

There are large and growing bodies of literature focused on the failures of both representative democratic institutions and mass media; and it is worth emphasizing that we do not view our findings as being fundamentally at odds with all of that work. There

clearly are situations in which thermostatic link between opinion and policy is weak, and even non-existent (Wlezien 2004; Soroka and Wlezien 2010). There are situations in which policy decisions are at odds with public opinion (e.g., Page and Shapiro 1983; Druckman and Jacobs 2015); and levels of spending may not match preferences as well (although this can be difficult to distinguish; see Wlezien and Soroka N.d.). There also is no question that media content is regularly incorrect, or incomplete, or simply silent on a wide range of topics of real relevance for democratic citizens. It is nevertheless the case that there are – conditional on, for instance, issue salience and institutional arrangements – domains in which policy and opinion are powerfully inter-related; and this is true in spite of what appear to be widespread failures on the part of journalists, politicians, and citizens. Clearly, there is some kind of signal getting through. We have argued that even a very general signal may be enough to make representative democracy work. This paper represents a first step towards identifying that signal.

Our aim for this paper has been relatively prosaic: we want to know whether media content includes the kinds of cues that would be required in order for citizens to learn about policy. This, we believe, is one basic requirement for thermostatic responsiveness to policy change. In order for citizens to adjust their preferences based on the level of policy, they need to know how policy actually has changed.

The starting point for our investigation has been a body of media content from the last 35 years, drawn from the *New York Times* and the *Washington Post*. The results are encouraging: there are a good number of cues about policy direction in these papers, and these cues are more often than not in line with the direction of policy change. This is particularly true for defense, where there is a very strong correlation between the "policy signal" evident in media content and policy itself. The welfare domain does not exhibit similarly strong correlations, but even there we see evidence of a media policy signal.

Why do we find a difference between defense and welfare? There are many possible reasons, and we cannot easily adjudicate between them here. One possibility is that our content analysis simply does not capture what it needs to in welfare. It could be that our spending and/or direction keywords are not ideal in this domain. It could be that media content on welfare is less focused on "spending," and more focused on increases or decreases in other types of policy action. It could also be that media content provides an amalgam of federal and state-level policy in this domain, and we do not adequately separate out the federal material here. Each of these possibilities is testable, relying on a larger body of human-coded data. This is one objective of future work.

Of course, it could also be the case that our estimates above are accurate: the media policy signal may simply be less clear on welfare than on defense. If this is the case, then why does past work find such strong public responsiveness in welfare? We should not forget that the role that media in thermostatic responsiveness likely varies widely across domains. There will be domains in which media content is very clear and abundant; and other domains in which media content is both unclear and sparse.

There will also be domains in which, regardless of the nature of content, public preferences are only tangentially related to media coverage. Media dependency theory (Ball-Rokeach and DeFleur 1976) makes clear that media effects are most prevalent in areas with which citizens do not have direct experience; we often do not need media to

tell us about the state of roads in our area, for instance. Other work also makes clear that policy preferences are driven by more than just domain-specific information; for instance, welfare preferences are deeply affected by macro-economic trends. In short, not only will media content sometimes not reflect policy, but public preferences in some domains will have other, potentially more important, inputs.

This is not to say that media do not matter, of course. Media content plays a significant role across a wide range of policy domains, and understanding the varying magnitude of media impact on public preferences is of real significance if we are to understand how exactly thermostatic responsiveness – and representative democracy – works. We take these results as evidence that a large-scale analysis of media cues is feasible and potentially illuminating. This will thus be one objective of our future work.

References

Althaus, Scott. 2003. *Collective Preferences in Democratic Politics*. Cambridge: Cambridge University Press.

Altheide, David L. 1997. "The News Media, the Problem Frame, and the Production of Fear." *Sociological Quarterly* 38(4): 647–68.

Ball-Rokeach, Sandra J. and M.L. DeFleur. 1976. "A dependency model of mass-media effects." *Communication Research* 3 (1): 3–21.

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.

Bennett, Stephen. 1988. "'Know-Nothings' Revisited: The Meaning of Political Ignorance Today. *Social Science Quarterly* 69:476-490.

Berelson, Bernard R., Paul F. Lazarsfeld, and William N. McPhee. 1954. *Voting*. Chicago: University of Chicago Press.

Converse, Philip E. 1970. "Attitudes and Non-Attitudes: Continuation of a Dialogue. In Edward R. Tufte (ed.), *The Quantitative Analysis of Social Problems*. Reading, Mass: Addison-Wesley.

-----. 1964. "The Nature of Belief Systems in Mass Publics." In David Apter (ed.), *Ideology and Discontent*. New York: Free Press.

Dahl, Robert. 1971. *Polyarchy: Participation and Opposition*. New Haven: Yale University Press.

Davie, William R., and Jung Sook Lee. 1995. "Sex, Violence, and Consonance/Differentiation: An Analysis of Local TV News Values." *Journalism and Mass Communication Quarterly* 72(1): 128–38.

Delli Carpini, Michael and Scott Keeter. 1996. *What American Know about Politics and Why it Matters*. New Haven: Yale University Press.

Druckman, James N. and Lawrence R. Jacobs. 2015. *Who Governs? Presidents, Public Opinion and Manipulation*. Chicago: University of Chicago Press.

Enns, Peter and Christopher Wlezien (eds.). 2011. *Who Gets Represented?* New York: Russell Sage Foundation.

Eichenberg, Richard, and Richard Stoll. 2003. "Representing Defence: Democratic Control of the Defence Budget in the United States and Western Europe." *Journal of Conflict Resolution* 47: 399-423.

Erikson, Robert S., Michael B. MacKuen and James A. Stimson. 2002. *The Macro Polity*. Cambridge: Cambridge University Press.

Franklin, Mark and Christopher Wlezien. 1997. "The Responsive Public: Issue Salience, Policy Change, and Preferences for European Unification." *Journal of Theoretical Politics* 9:347-63.

Geer, John Gray. 1996. *From Tea Leaves to Opinion Polls: A Theory of Democratic Leadership*. New York: Columbia University Press.

Jennings, Will. 2009. "The Public Thermostat, Political Responsiveness and Error Correction: Border Control and Asylum in Britain, 1994-2007." *British Journal of Political Science* 39:847-870.

Lawrence, Regina G. "Game-Framing the Issues: Tracking the Strategy Frame in Public Policy News." *Political Communication* 17(2): 93-114.

Lichter, S. Robert and Richard Noyes. 1995. *Good Intentions Make Bad News*. Lanham

MD: Rowman and Littlefield.
Meyrowitz, Joshua. 1985. *No Sense of Place: The Impact of Electronic Media on Social Behavior*. New York: Oxford University Press.
Page, Benjamin I. and Robert Y. Shapiro. 1992. *The Rational Public: Fifty Years of Trends in Americans' Policy Preferences*. Chicago: University of Chicago Press.
———. 1983. "Effects of public opinion on policy." *American Political Science Review* 77: 175–90.
Popkin, Samuel and Michael Dimock. 1999. Political Knowledge and Citizen Competence. In Stephen Elkin and Karol Salton (eds), *Citizen Competence and Democratic Institutions*. University Park: Pennsylvania State University.
Patterson, Thomas. 1994. *Out of Order*. New York: Vintage.
Sabato, Larry, 1991. *Feeding Frenzy*. New York: Free Press.
Soroka, Stuart. 2014. *Negativity in Democratic Politics: Causes and Consequences*. Cambridge Studies in Public Opinion and Political Psychology, Cambridge University Press.
Soroka, Stuart N. and Christopher Wlezien. 2010. *Degrees of Democracy: Politics, Public Opinion and Policy.* Cambridge University Press.
Stuart Soroka, Christopher Wlezien and Iain MacLean. 2006. "Public Expenditure in the UK: How Measures Matter," *Journal of the Royal Statistical Society (Series A)*, volume 169.
Turcotte, Jason. 2014. "Debates and the Disincentives of Democracy: How News Norms and Values Influence Moderator Agendas." *Journalism Practice 8(6): 772-788*
Wlezien, Christopher. 2004. "Patterns of Representation: Dynamics of Public Preferences and Policy." *Journal of Politics* 66:1-24.
———. 1996. "Dynamics of Representation: The Case of US Spending on Defense." *British Journal of Political Science* 26:81-103.
———. 1995. "The Public as Thermostat: Dynamics of Preferences for Spending." *American Journal of Political Science* 39:981-1000.
Wlezien, Christopher and Stuart Soroka. N.d. "Measuring and Modeling Public Preferences for Policy: The Case of Spending on Welfare in the US."
———. 2012. "Political Institutions and the Opinion-Policy Link." *West European Politics* 35(6): 1407-1432.
———. 2003. "Measures and Models of Budgetary Policy." *Policy Studies Journal* 31(2): 273-286.

Appendix Figure A1. Changes in Defense Outlays and Media Cues

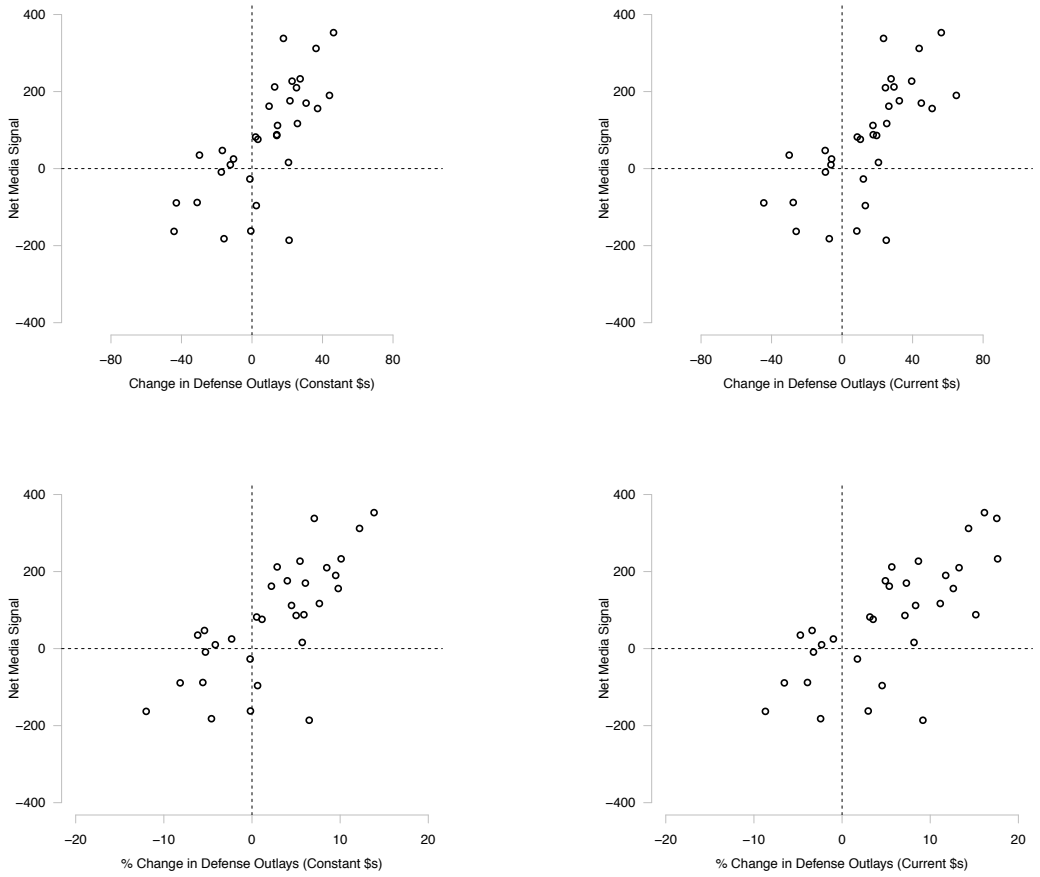
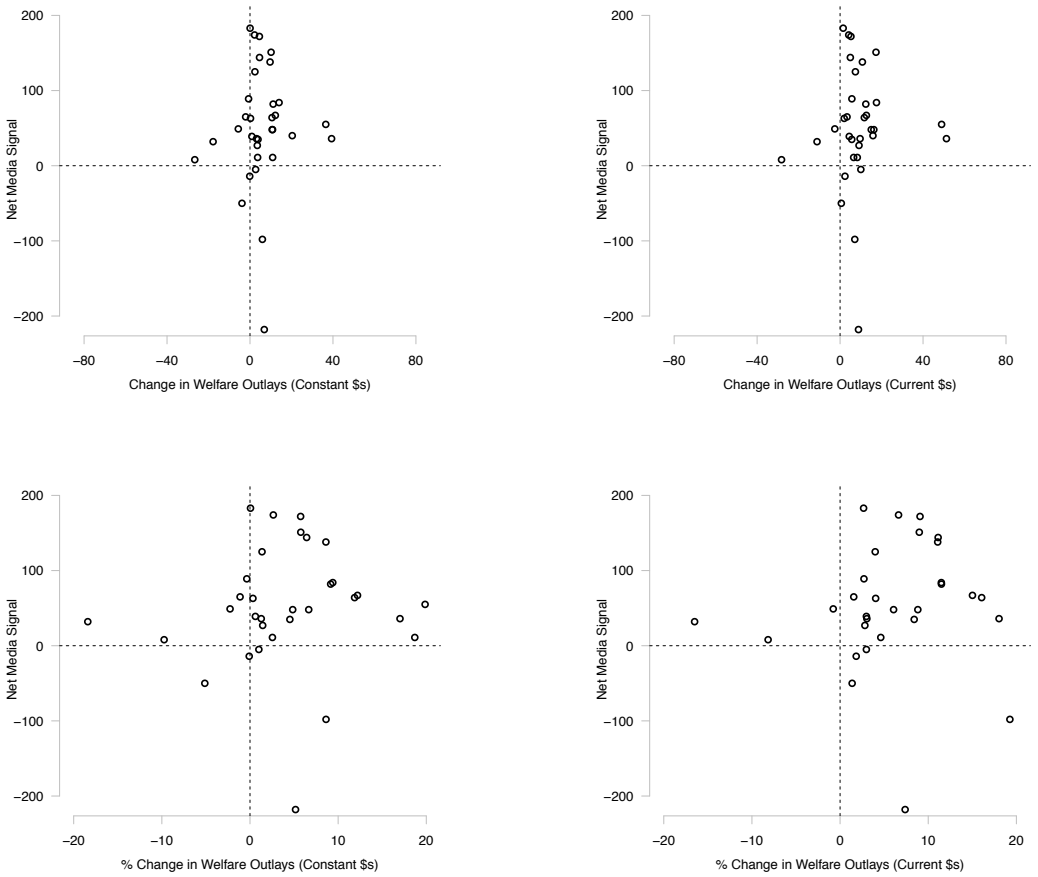
Appendix Figure A2. Changes in Welfare Outlays and Media Cues