

The Media, the Economy and the Vote ^{*}

Mark Andreas Kayser[†]
Michael Peress[‡]

September 2015

Abstract

What role does the media play in forming voters' perception of the economy and their choice of which party to vote for? Previous studies have suggested a variety of (sometimes conflicting) answers based on limited data—often no more than two sources and a single country—usually the United States. This paper introduces a large project designed to test the basic relationship between the media, partisanship, and voting cross-nationally. Having collected over 2 million articles related to the economy from 32 newspapers in 16 developed countries, we present preliminary results on three relationships: (1) How well does newspaper sentiment reflect the economy? (2) Does media partisanship bias reporting on the economy? and (3) Does newspaper reporting mediate the economic vote?

^{*}Prepared for presentation at the 2015 meeting of the European Political Science Association, Vienna. We gratefully acknowledge peerless research assistance from Giuliana Pardelli, CJ Yetman, Jeff Arnold, Arndt Leininger, William Pollock, and Johanna Willmann. We thank Paul Kellstedt for sharing his English language coding dictionary with us. We thank participants of seminars at the University of Bremen, UC Berkeley, MPSA 2015, APSA 2014, EPSA 2014 and the Workshop in Political Economy & Political Science (Santiago, 2014).

[†]Hertie School of Governance, Berlin, kayser@hertie-school.org

[‡]SUNY-Stony Brook, michael.peress@stonybrook.edu

1 Introduction

The economy influences election outcomes across a broad swath of countries, periods, institutions and contexts (Nadeau, Lewis-Beck and Bélanger, 2013). While other predictors of individual vote choice matter at least as much,¹ economic performance understandably captures a disproportionate share of scholarly and popular attention. Voters do not change party identifications quickly, nor do parties often shift major policy positions during a campaign; the economy, however, can and does continually change. Politicians and campaign strategists have been quick to grasp this (Gelman and King, 1993; Sides and Vavreck, 2014; Erikson and Wlezien, 2012) but scholarship tells us little about how voters form economic perceptions.

The study of economic effects on voting constitutes one of the largest and most established literatures in political science.² The preponderance of studies obtain their results by regressing election outcomes on economic aggregates—such as growth, unemployment, and inflation— or, usually more fruitfully (Kayser, 2014), economic perceptions from surveys. How such perceptions arise, however, is less clear. Few voters learn about the state of the economy directly through these aggregates. In this article, we focus on one possible mechanism: the media’s role in forming voters’ perception of the economy and the consequent effect on party choice. Previous studies have suggested a variety of (sometimes conflicting) answers to these questions based on limited data—often no more than two sources and a single country—usually the United States. This paper introduces a large project designed to test the basic relationship between the media, partisanship, and voting cross-nationally. How does the media influence economy perceptions and how does this matter for election outcomes?

We focus here on answering three questions. The most basic question is whether newspaper sentiment reflects the economy. We consider whether the adjectives newspapers use to describe growth, the labor market, and inflation, correspond with growth rates, unemployment rates, and inflation rates. We compare this relationship to the relationship between the vote share of incumbent parties and the economy. We also consider whether the emphasis newspapers place—that is, the relative share of economic news they devote to covering growth, the labor market, and inflation—is related to the actual performance of the economy.

We next study media bias in reporting on the economy, a question that has yielded contradictory results (cf. Lott Jr and Hassett, 2014; Puglisi and Snyder, 2014) but has been largely constrained to the United States. Do left-wing and right-wing newspapers

¹cf., party identification and proximity in policy preferences between parties and voters

²A search of Google Scholar for the term “economic vote” in April 2015, revealed 1300 hits, despite the fact that most publications on this subject probably do not explicitly use the term.

report on the economy in a way which is different and beneficial to left-wing and right-wing governments, respectively? We study both whether there is a bias in sentiment on the economy and whether there is a bias in the share of coverage devoted to growth, the labor market, and inflation.

Finally, we examine whether newspapers mediate the economic vote. That is, are incumbent governments punished when the economy performs poorly, or when the media report that the economy is performing poorly. Research has explored whether vote intentions condition (Wlezien, Franklin and Twiggs, 1997; Evans and Anderson, 2006) or do not condition (Duch and Stevenson, 2008; Lewis-Beck, Martini and Kiewiet, 2013) economic perceptions building on work showing that partisans filter (Zaller, 2004; Stanig, 2013) or neglect economic (Kayser and Wlezien, 2011) economic and other (Eggers, 2015) news to fit their partisan preferences. Scholars have also estimated the time-lead at which economic outcomes most influence vote choice (Achen and Bartels, 2004; Healy and Lenz, 2014; Wlezien, 2015). The actual source of economic perceptions, however, has remained relatively neglected.

Our analysis with respect to the media's role in the economic vote is relevant for a number of reasons—first, it says something about the level of sophistication of voters. Second, to the extent that the economic vote is mediated by newspaper coverage, it must be sociotropic. Third, media bias can only have an effect on elections if newspaper coverage exerts a causal effect on voting behavior (i.e. newspaper coverage mediates the economic vote). We further note that the voluminous literature on the economic vote, by neglecting the media, implicitly assumes either (a) that voters form economic perceptions exclusively through direct experience of the economy or (b) that the media does not distort perceptions of the economy. Both assumptions are directly at odds with findings that indirectly suggest a role for the media in the formation of economic perceptions and voting. Della Vigna and Kaplan (2006) demonstrate that towns in the U.S. with access to Fox News vote slightly more Republican than otherwise expected (although this effect is not necessarily attributable to coverage of the economy).³ The superior performance of real-time (i.e., initial) economic estimates over vintage (i.e., updated) in forecasting elections (Kayser and Leininger, 2015), for example, is most easily explained by the fact that initial estimates are more widely reported.

To answer these questions, we collected over 2 million articles related to the economy from 32 newspapers in 16 developed countries. From each country in our sample, we collected a left-wing and a right-wing newspaper. We apply sentiment analysis to generate a monthly time series for each of the 32 newspapers.

Our results indicate that newspaper sentiment on growth, unemployment, and infla-

³Arceneaux et al. (2015) show a similar effect of Fox News on the policy positions of elites.

tion track growth, changes in unemployment, and changes in inflation. We also find that the share of coverage devoted to growth, unemployment, and inflation is higher when the relevant aggregates suggest the economy is performing poorly. We find that newspapers are not biased in their tone—newspapers whose ideology matches the current government do not cover growth, unemployment, and inflation more favorably. We do however find some evidence of bias in the tone of economic coverage—newspapers whose ideology matches the current government focus somewhat less on unemployment when unemployment is high than newspapers whose ideology does not match the current government. Finally, we find that voters respond directly to unemployment, but only respond to media reports of growth.

2 The Data

2.1 Newspaper Articles

As our motivation was to study the mechanisms behind the economic vote, we began with a list of twenty five OECD countries which are typically included in studies of the economic vote. Our goal was to obtain a time-series of newspaper sentiment as long as possible, for both a left-wing and right-wing paper, for as many OECD countries as possible. One limitation is the many languages that are spoken among these countries. We focused our analysis on three languages which were spoken in many of the OECD countries—English, French, and German. We were able to include Spanish, Portuguese, and Italian newspapers as well because our French-speaking research assistant happened to speak these languages as well. Using these six languages combined, we were able to include 16 of the OECD countries in our sample.⁴

In each country, we attempted to identify a relatively left-wing and relatively right-wing newspaper for which we could obtain electronic copies of articles. Our preference was for papers that had large circulation, were mainstream rather than ideologically extreme or tabloid, and had a long time series of articles available. When a mainstream left-wing or right-wing paper was not available, we collected an extreme left-wing or right-wing newspaper. If either a left-wing or right-wing paper was not available, we collected a centrist paper. We coded the ideology of newspapers on a -2 to 2 scale, with -2 being extreme-left, -1 being left, 0 being centrist, 1 being right, and 2 being extreme right, based on a number of web-sources.

⁴These countries include Australia, Austria, Canada, France, Germany, Ireland, Israel, Italy, Japan, Luxembourg, New Zealand, Portugal, Spain, Switzerland, the United Kingdom, and the United States. The OECD countries excluded from our study include Belgium, Denmark, Finland, Greece, Iceland, Netherlands, Norway, Sweden, and Turkey.

Our dataset consists of over 2 million articles from 32 newspapers. Our sample represents a large increase in coverage over previous studies, in the number of newspapers, the number of countries, and the number of articles. Most previous studies have relied on human labor to categorize articles, which necessarily limited them to small samples (one or two newspapers) usually from a single country. We employ automated collection and content analysis which enabled the analysis of text on a scale not possible with human-coded text analyses. Prominent hand-coded studies, for comparison, such as [De Boef and Kellstedt \(2004\)](#) and [Soroka \(2006\)](#) were able to categorize 4000 and 5000 articles from a single newspaper, respectively.

Automated content analysis also enabled a smaller unit of analysis. Rather than classifying individual economic articles as positive or negative, we used text fragments as the basis of our sentiment analysis. This approach enabled us to capture more nuance than is possible with the discrete categorizing of economic articles as positive or negative. The actual unit of analysis was aggregated up to the month—the proportion of positive (or negative) economic text fragments in a given month—in order to match the economic data.

Why newspapers? Indeed, news content is fragmented over a variety of media in addition to newspapers, for example, television, twitter, social networks, online news portals, blogs, and other sources. Two reasons guide our decision for focus on newspapers. First, newspapers offer the longest coverage available to researchers. Television transcripts start later and social media, such as twitter, even later.⁵ Second, newspaper reporting tends to lead other news media (see, for example, [Roberts and McCombs, 1994](#)).

2.2 Economic Data

Our two sources of economic data were the OECD and the IMF. We used the highest frequency data that was available. If monthly data was available (as was sometimes the case for unemployment and inflation) we used monthly data. In only quarterly data was available, we converted the quarterly data to monthly data as follows. For growth we assumed a constant level of growth throughout the time period. For unemployment, we assumed a constant unemployment rate throughout the time period. For inflation, we assumed a constant rate of inflation throughout the time period. When quarterly data was not available, we imputed the quarterly data based on yearly data and we then imputed the monthly data based on the quarterly data. We used the highest level available preferentially, we used harmonized data (for unemployment and inflation) preferentially over unharmonized data, and we used the OECD data preferentially over the IMF data.

⁵See [Burckhardt, Duch and Matsuo \(2013\)](#) for an interesting use of twitter data to study regional perceptions of the economy in the UK.

We deleted a few clearly implausible values of the economic data that would greatly distort our results (i.e. a growth rate of a million percent one quarter followed by a growth rate of negative a million percent in the next quarter).

Once the data were converted to monthly values, we could then aggregate them to various other time periods. For example, consider an election held in March of 2003. In the election voters may respond to unemployment in the current month, the current quarter, the current year, etc. Quarterly unemployment in march would be computed as the average unemployment in the previous 3 months. Yearly unemployment would be computed as the average unemployment in the previous 12 months. However, suppose that only yearly unemployment was available. It was seem redundant to impute the monthly data based on yearly data only to covert the monthly data back to yearly data. Instead, the yearly data for March would be based on both the 2012 and 2013 unemployment rates, with the 2012 rate receive a weight of 9/12 and the 2013 data receiving a weight of 3/12, so that the coverting we do allows us to interpolate appropriate economic aggregates for elections that did not coincide with the repoting of economic aggregates.

2.3 Election Results

We assembled a dataset on democratic election results for 25 OCED countries for from a variety of sources. Our main sources were “The International Almanac of Electoral History” by Thomas T. Mackie and Richard Rose, “Elections in Europe: A Data Handbook”, by Dieter Nohlen and Phillip Staver, “Elections in the Americas: A Data Handbook: Volume 1: North America, Central America, and the Carribbean” by Dieter Nohlen, “Elections in Asia and the Pacific: A Data Handbook: Volume 1: Middle East, Central Asia, and South Asia”, by Dieter Nohlen, Florian Grotz, and Christof Hartmann, “Elections in Asia and the Pacific: A Data Handbook: Volume 2: South East Asia, East Asia, and the Pacific” by Dieter Nohlen, Florian Grotz, and Christof Hartmann, Parlgov (<http://www.parlgov.org/>), the Constituency-level Election Archive (<http://www.electiondataarchive.org/>), Adam Carr’s Election Archive (<http://psephos.adam-carr.net/>), and as a last resort, various countries’ Wikipedia websites. The ideological orientation of the prime minister’s party was coded based on the Comparative Manifesto Project.

3 Methodology

Our starting point was studies of the economic vote, which by and large focus on three aspects of the economy—growth, unemployment, and inflation [Powell and Whitten \(1993\)](#).

Our goal was to code sentiment for corresponding categories in newspaper coverage.

In line with our theory, we sought to consider the impressions that an average voter would receive about the economy upon reading the average newspaper article covering the economy. In our theory, voters form their impressions of the economy from newspapers as well as other sources. Voters as a group determine the electoral fortunes of incumbent parties. What matters then is the impression that voters as a group would obtain from reading newspaper coverage of the economy. This definition acknowledges that there may be some measurement error if a single human coder would rate media sentiment because our definition of sentiment is the impression that an average voter would have. It also acknowledges that different voters may read different articles, and potentially different newspapers.

We first paired down the set of articles down to a more reasonable size. We used keyword searches to identify articles that were related to the economy. This involved analyzing approximately 5% of the articles from each newspaper. From these articles, we sought to code sentiment as positive, negative, or neutral along the three possible dimensions—growth, unemployment, or inflation.

3.1 Approaches to Sentiment Analysis

We considered four different approaches for coding economic sentiment. The first is direct human coding. We rejected this possibility because it was too time intensive to be practical. This left 3 computer-aided approaches for coding economic sentiment—dictionary-based methods, supervised learning, or unsupervised learning. Unsupervised learning attempts to determine clusters of related words—or words that tend to appear in the same document. Often, it ignores the order of words and simply works with the raw frequency of each word, although it is possible to apply unsupervised learning to the frequency of two word or three word combinations. The algorithm will produce a number of clusters as well as words that are predictive of those clusters. There is no guarantee that those clusters will capture what we would like to capture. In the best case scenario, the algorithm would recover from among the clusters, clusters including growth-positive, growth-negative, unemployment-positive, unemployment-negative, inflation-positive, inflation-negative. Using the raw frequency of words, it seems exceedingly unlikely that the algorithm will recover this. Even using combinations of multiple words, it is unlikely that unsupervised learning that will recover such clusters.

A second approach is supervised learning. Here, a set of documents could be coded by humans. Supervised learning would then try to replicate the human coding using a statistical model. The main limitation of supervised learning is that determining sentiment on growth, unemployment, and inflation requires relating terms denoting positive

and negative sentiment with terms denoting growth, unemployment, and inflation. These words are likely to be separated by a few connecting words. Hence, employing supervised learning would require considering strings involving at least three or four words. Once we consider strings involving three or four words, the number of possible combinations of “features” to base the supervised learning estimators on becomes very large.

The third approach—and the one we use—is a dictionary-based approach. Consider the following simplified version of a dictionary based approach. We identify a number of words which denote growth. We also identify a number of words that denote positive and negative sentiment. We would then code sentiment based on the relative frequency of positive and negative words near growth words (where we could use 5 words away as our definition as “near”). This approach would require us to pre-identify a list of growth, positive, and negative words. The approach therefore requires a higher degree of language-specific knowledge than supervised and unsupervised learning, but solve the problem of having to deal with a very large feature space.

3.2 Dictionary Coding of Sentiment

The actual approach we used is somewhat more involved than what we described above. We have a separate dictionary of negations and a nearby negation alters the meaning of a positive or negative word. We have a dictionary of words indicating increasing and decreasing where increasing words near growth contribute to positive sentiment and decreasing words near growth contribute to negative sentiment. We have a separate list of words indicating recession, which are coded as negative sentiment. All together, we calculate the number of positive growth instances divided by positive plus negative growth instances in a given month to generate our measure of sentiment for that newspaper in that month. Similar rules are used to generate sentiment for unemployment and inflation.

To developing our English dictionaries, we made small modifications to existing dictionaries. We developed our French, German, Italian, Portuguese, and Spanish dictionaries from scratch, using the existing English dictionaries as reference. The dictionaries were all developed by fluent speakers.

4 Does Newspaper Sentiment Reflect the Economy?

Does newspaper sentiment reflect the economy? We first consider the correlation between our measure of newspaper sentiment on growth, unemployment, and inflation, and our economic aggregates. We measure sentiment as the proportion of relevant text fragments that positively depict the economy. Growth, unemployment, and inflation are measured

monthly. However, newspaper sentiment may not necessarily reflect only the performance of the economy over the last month. For this reason, we compare monthly newspaper sentiment to economic performance over the most current month, quarter, semi-year, year, two-year period, and four-year period. In addition, we consider both levels and changes in unemployment and inflation. The correlations are presented in Table 1.

We see that same-month economic statistics associate less strongly with monthly newspaper sentiment than do longer-period economic measures. Growth in the year up to a given month is most strongly related to growth sentiment in that month. The change in unemployment over the prior six months correlates most strongly with unemployment sentiment in a given month and changes in the inflation rate over the prior 12-months associate most strongly with inflation sentiment. These results suggest that economic outcomes lead newspaper sentiment. This is far from claiming that news reporting faithfully, objectively and myopically reflects the real economy. It is still clearly possible that although economic news often precedes sentiment (Blood and Phillips, 1995), economic sentiment among the population may influence how and what is reported (Gentzkow and Shapiro, 2006; Hopkins, 2010). Nor are economic news the only determinant of economic sentiment (De Boef and Kellstedt, 2004). Given these caveats, it is nevertheless reassuring to see preliminary results that prior and longer-term economic variation matters for media sentiment.

	Month	Quarter	Semi-year	Year	2 Years	4 Years	N
Growth Sentiment							
Growth	0.31	0.36	0.41	0.42	0.40	0.34	6650
Unemployment Sentiment							
Unemployment	0.01	0.02	0.03	0.06	0.11	0.14	6652
Change in Unemployment	-0.18	-0.30	-0.32	-0.26	-0.12	-0.03	6605
Inflation Sentiment							
Inflation	-0.05	-0.07	-0.04	0.02	0.08	0.11	6606
Change in Inflation	0.01	-0.06	-0.11	-0.17	-0.09	-0.01	6606

Table 1: Correlations between Newspaper Economic Sentiment and the Macro-economy.

In Table 2, we consider the ability of our economic variables—growth, changes in unemployment, and changes in inflation—all measured yearly, to predict economic sentiment. We see that there is quite a bit of variation across newspapers. The R-squared from bivariate OLS regressions is as high as 0.68 for the Austrian newspaper, *Die Presse*, and as low as 0.03 for the Luxembourgian newspaper, *Le Quotidien*. There is a relatively poor fit for both Spanish newspapers, one Portuguese newspaper, one German newspaper, both Japanese newspapers, and one Italian newspaper. Before taking these as evidence that these newspaper provide an inaccurate assessment of the economy, we must consider

Country	Newspaper	Relative Ideology	Language	R^2	N
Australia	<i>The Age</i>	L	English	0.34	273
Australia	<i>Herald Sun</i>	R	English	0.27	245
Austria	<i>Der Standard</i>	L	German	0.43	69
Austria	<i>Die Presse</i>	R	German	0.68	113
Canada	<i>Toronto Star</i>	L	English	0.40	336
Canada	<i>The Globe and Mail</i>	R	English	0.31	429
France	<i>Le Monde</i>	L	French	0.33	276
France	<i>Le Figaro</i>	R	French	0.58	200
Germany	<i>Die Zeit</i>	L	German	0.37	62
Germany	<i>Frankfurter Allgemeine</i>	R	German	0.11	45
Ireland	<i>Irish Times</i>	L	English	0.43	247
Ireland	<i>Irish Independent</i>	R	English	0.59	83
Israel	<i>Globes</i>	L	English	0.22	204
Israel	<i>Jerusalem Post</i>	R	English	0.29	295
Italy	<i>La Stampa</i>	L	Italian	0.20	252
Italy	<i>Corriere della Sera</i>	R	Italian	0.42	56
Japan	<i>Nikkei Weekly</i>	L	English	0.12	399
Japan	<i>Daily Yomiuri</i>	R	English	0.14	283
Luxembourg	<i>Le Quotidien</i>	L	French	0.03	60
Luxembourg	<i>Le FAX d'Agefi</i>	R	French	0.34	25
New Zealand	<i>The Press</i>	L	English	0.38	207
New Zealand	<i>New Zealand Herald</i>	R	English	0.53	178
Portugal	<i>Correio da Manha</i>	L	Portuguese	0.66	15
Portugal	<i>Jornal de Noticias</i>	R	Portuguese	0.20	191
Spain	<i>El Pais</i>	L	Spanish	0.15	200
Spain	<i>El Mundo</i>	R	Spanish	0.05	134
Switzerland	<i>Tages-Anzeiger</i>	L	German	0.27	163
Switzerland	<i>Neue Zurcher Zeitung</i>	R	German	0.34	232
United Kingdom	<i>The Guardian</i>	L	English	0.54	343
United Kingdom	<i>London Times</i>	R	English	0.43	330
United States	<i>New York Times</i>	L	English	0.48	289
United States	<i>The Wall Street</i>	R	English	0.48	415

Table 2: Newspaper Sample and Fit Statistics – 32 newspapers, 16 countries, 6 languages, and 6663 months. *Nikkei Weekly* and *Die Zeit* are weekly papers.

whether these poor fits are due to coding issues specific to these newspapers.

One possible explanation is that the newspapers we collected that did not report in the native language would provide worse fits because they were not placing as much coverage on the local economy. This could indicate that these newspapers do not provide good measures of the media in these countries. These countries are Israel and Japan. There is some evidence that this is the case—these countries have some of the worst performance in our sample. Note also that we have used English language newspapers for these countries—news sources that might have a more international readership and orientation.

Country	Correlation	N
Australia	0.86	212
Austria	0.55	69
Canada	0.88	335
France	0.75	192
Germany	0.36	45
Ireland	0.87	75
Israel	0.68	203
Italy	0.46	48
Japan	0.64	283
Luxembourg	0.41	28
New Zealand	0.71	178
Portugal	0.55	13
Spain	0.45	126
Switzerland	0.78	150
United Kingdom	0.82	326
United States	0.85	289

Table 3: Correlation between Left-wing and Right-wing media sentiment across the 16 countries in our sample.

Our second concern is that some dictionaries will exhibit poorer performances than others. Our maintained hypothesis is that a poorly coded dictionary will lead to measurement error, which will depress correlations. We thus argue that observing either a strong correlation between right and left wing newspapers or observing a strong correlation between sentiment and the economy is evidence against a poor dictionary.⁶ The correlation between the left and right wing papers only falls below 50% in Germany, Italy, Luxembourg, and Spain. Since the German language dictionary performs very well in Austria and Switzerland and the French dictionary performs very well in France, we believe that we can dismiss problems with these dictionaries. This leaves us with the Italian and Spanish dictionaries. Even in these case, the left-right correlations are reasonable strong. Moreover, the dictionaries were constructed by the same research assistant who compiled the English, French and Portuguese dictionaries that performed well for other newspa-

⁶Table 3 presents the correlations between sentiment in the left-wing and right-wing newspapers.

pers, so we cautiously proceed assuming that the low correlations between the economy and sentiment in Italy and Spain reflect the fact that newspapers in these countries cover the economy in a way that is only loosely related to actual economic conditions.

Dependent Variable:	Growth Sentiment	Unem. Sentiment	Inflation Sentiment	Economic Sentiment
Independent Variables:				
Growth (yearly) (SD = 3.072)	0.023*** (0.003)	0.011*** (0.002)	-0.003* (0.001)	0.013*** (0.002)
Change in Unem. (yearly) (SD = 0.955)	-0.002 (0.006)	-0.013** (0.003)	0.006 (0.003)	0.001 (0.003)
Change in Inf. (yearly) (SD = 12.570)	-0.003+ (0.001)	0.002 (0.001)	-0.008*** (0.002)	-0.003** (0.001)
Number of Months	6650	6644	6595	6649
Number of Newspapers	32	32	32	32
Number of Countries	16	16	16	16
R-Squared	0.601	0.354	0.282	0.643

Table 4: The Effect of the Economy on Newspaper Sentiment – Dependent variables are positive sentiment for growth, unemployment, inflation, and overall economic performance. Newspaper fixed effects were included in each regression, but omitted from the table. Standard errors were clustered by country.

In Table 4, we use four types of monthly positive economic sentiment as our dependent variables and we use our yearly measure of growth, change in unemployment, and change in inflation as independent variables. We include newspaper fixed effects to account for difference in the way language is used by different newspapers. We clustered the standard errors by country to account for correlations in the unobserved shocks between newspaper-months in the same country. We find that growth, changes in unemployment, and changes in inflation, affect sentiment across growth, unemployment, and inflation, respectively. We also find that growth has a positive effect on sentiment for unemployment and inflation. Growth exhibits a positive and inflation a negative effect on overall economic sentiment, but unemployment does not have a statistically significant effect on overall sentiment. Overall, these results suggest at least partial measurement reliability and validity in our economic sentiment measures.

4.1 Comparison with the Economic Vote

It is instructive to compare our results relating media sentiment to the economy in Table 1 to similar results for the economic vote. Table 5 presents the correlations between the vote share for the prime minister’s party and various macro-economic measures. The economic vote, in contrast to economic sentiment, responds more closely to level rather than changes in unemployment and inflation. We also find that the economic vote responds

most closely to growth over a two-year period, monthly unemployment, and quarterly inflation. Thus, growth over a fairly long period of time affects the economic vote, while the economic vote responds to relatively recent levels of unemployment and inflation. A possible explanation for this finding—which we will further consider—is that voters learn about growth from the media, but experience unemployment and inflation more directly. This would also concord with other research that shows that real-time measures of economic growth outperform updated (vintage) measures in predicting incumbent vote share, presumably because initial (real-time) estimates are most reported in the press (Kayser and Leininger, 2015). The reaction of the economic vote to growth is relatively consistent with the reaction of newspaper sentiment to growth. The reaction of the economic vote to unemployment and inflation is not consistent with reaction of newspaper sentiment to unemployment and inflation—both because the economic vote responds to levels rather than changes and because the economic vote reacts only to recent levels of unemployment and inflation.

	Month	Quarter	Semi-year	Year	2 Years	4 Years	N
Vote Share of Prime Minister’s Party							
Growth	0.14	0.21	0.22	0.17	0.28	0.26	348
Unemployment	-0.18	-0.17	-0.17	-0.17	-0.15	-0.14	274
Change in Unemployment	0.04	-0.04	-0.02	-0.10	-0.10	-0.02	243
Inflation	-0.07	-0.08	-0.07	-0.06	-0.07	-0.06	394
Change in Inflation	0.06	-0.08	-0.02	0.04	0.00	0.02	369

Table 5: Correlations between Vote Shares for the PM’s Party and the Macro-economy.

4.2 The Relative Share of Coverage for Growth, Unemployment, and Inflation

We next investigate which aspect of economic performance—growth, unemployment, or inflation—receives the most newspaper coverage. In our data, the average share of economy-related sentences (averaged over months) devoted to growth, unemployment, and inflation, are 46.5%, 18.3%, and 35.3%, respectively. It is interesting to investigate the variation in these. We start by comparing the correlations of the share of coverage of each to our measures of economic performance. As before, we consider different windows for our calculation of economic performance. The results can be seen in Table 6.

We find that newspapers are more likely to cover growth, unemployment, and inflation, when economic performance according to these measures is poor. Moreover, we find that coverage responds more strongly to levels rather than changes in unemployment and inflation (an interesting difference to our finding for sentiment). In terms of the window,

	Month	Quarter	Semi-year	Year	2 Years	4 Years	N
Growth Share of Coverage							
Growth	-0.06	-0.09	-0.12	-0.14	-0.15	-0.12	6649
Unemployment Share of Coverage							
Unemployment	0.56	0.56	0.55	0.54	0.52	0.47	6656
Change in Unemployment	0.09	0.16	0.23	0.27	0.29	0.24	6609
Inflation Share of Coverage							
Inflation	0.16	0.22	0.27	0.25	0.22	0.15	6660
Change in Inflation	0.01	0.02	0.09	0.11	0.06	0.01	6660

Table 6: Good News Gets Reported, Bad News... Less So – Correlations between newspaper coverage and the macro-economy.

Dependent Variable:	Growth Share of Coverage	Unemployment Share of Coverage	Inflation Share of Coverage
Independent Variables:			
Growth (yearly)	-0.013*** (0.003)		
Unemployment (yearly)		0.017*** (0.003)	
Inflation (yearly)			0.009* (0.004)
Number of Months	6649	6656	6660
Number of Newspapers	32	32	32
Number of Countries	16	16	16
R-Squared	0.779	0.734	0.644

Table 7: Newspaper fixed effects were included in each equation, but omitted from the table. Standard errors were clustered by country.

the results are somewhat different here. For example, unemployment coverage responds most strongly to contemporaneous unemployment, but the difference in fit is relatively small. For this reason, it is reasonable to model the share of coverage for growth, unemployment, and inflation, as a function of yearly levels of growth, unemployment, and inflation.

We consider additional results in Table 7. We include fixed effects for each newspaper in each equation to account for differential patterns across newspaper. We cluster the standard errors by country to account for correlations in the unobserved shocks between newspaper-months in the same country. Consistent with the previous results, we find that for each economic measure, the media pay more attention when the economy is performing poorly according to that measure.⁷ Overall, the results provide strong evidence that newspaper pay more attention to measures on which the economy is performing poorly.

⁷These results are broadly consistent with those of Soroka (2006).

5 Media Bias in Reporting on the Economy

In the previous section, we found that the tone of media coverage reflected the economy—newspaper sentiment on growth, unemployment, and inflation, were to a large degree explained by growth, unemployment, and inflation. The fact that sentiment does not perfectly track these economic aggregates leaves open the possibility the newspapers differ in their reporting of the economy. In this section, we focus on partisan differences. Specifically, we focus on whether right-wing (left-wing) newspapers report more positive sentiment when there is a right-wing (left-wing) government. Our main analysis focuses on the relative left/right coding of newspapers we previously reported. We coded the left/right ideology of the incumbent prime-minister’s party based on the Comparative Manifesto Project. We then coded the variable *Ideological Match* as 1 for observations where the newspaper and the prime-minister had the same ideological orientation.

5.1 Media Bias in Tone

Table 3 has already shown that economic sentiment in left and right papers correlates at above .5 in all countries of our sample except four, suggesting that left and right wing newspapers report similar sentiment, but leaving some room for partisan differences. We expand on this in Figure 1. The results here suggest that left and right newspapers track each other closely and the differences are not well explained by partisan differences.

The first three columns of Table 8 regresses the proportion of positive sentiment for growth, unemployment, and inflation on a dummy variable indicating whether a newspaper in each month matches the ideology of the prime minister’s party. The results are striking: co-partisanship has no effect whatever on the media sentiment. The tone of reporting in a paper does not change when the party in power is replaced by one with a different ideological orientation. Newspaper reporting on the economy shows the same proportion of positive stories regardless of the ideology of the governing party.

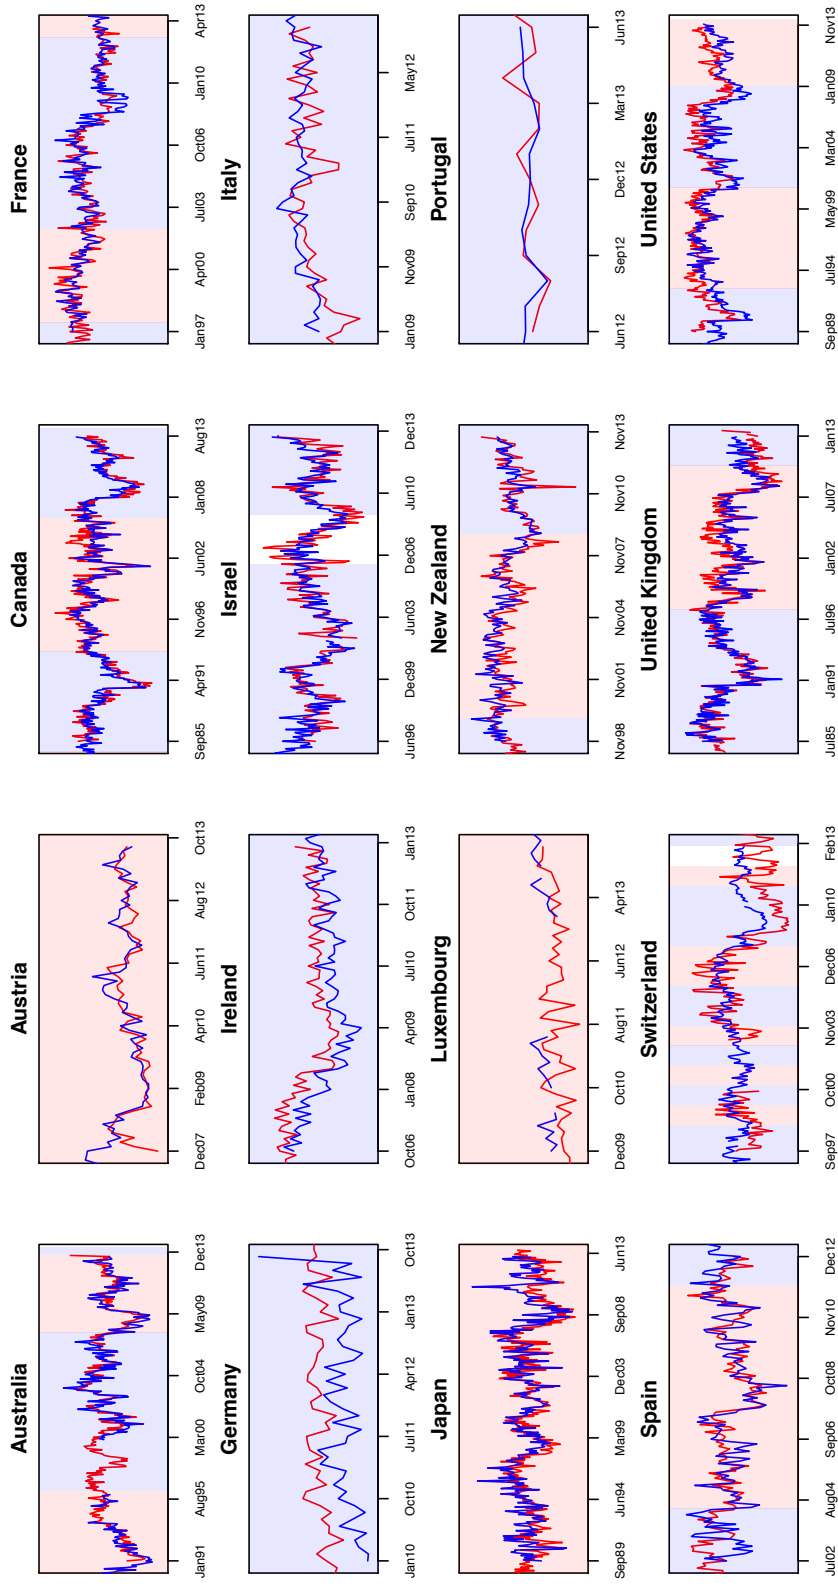


Figure 1: Proportion of positive reports on the economy in left (red) and right (blue) newspapers over time.

Dependent Variable:	Growth		Unem.		Inf.	
	Sent.	Unem. Sent.	Sent.	Unem. Sent.	Sent.	Inf. Sent.
Independent Variables:						
Ideological Match	-0.013 (0.019)	-0.008 (0.010)	0.011 (0.010)	0.000 (0.008)	-0.001 (0.008)	
Growth (yearly)			0.023*** (0.003)	0.012*** (0.001)	-0.007*** (0.002)	
Change in Unem. (yearly)				-0.019*** (0.004)		
Change in Inf. (yearly)					-0.011*** (0.002)	
Growth * Ideo. Match			0.000 (0.004)	-0.001 (0.002)	0.004+ (0.002)	
Change in Unem. * Ideo. Match				0.009 (0.005)		
Change in Inf. * Ideo. Match					0.005+ (0.002)	
Number of Months	6664	6658	6608	6644	6595	
Number of Newspapers	32	32	32	32	32	
Number of Countries	16	16	16	16	16	
R-Squared	0.467	0.250	0.249	0.601	0.285	

Table 8: Media Bias in Tone

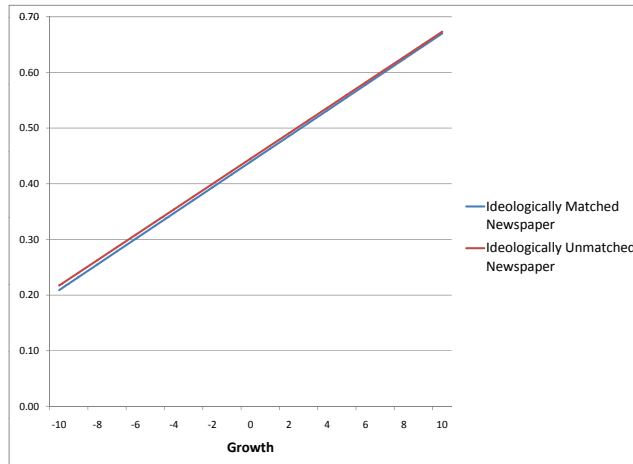


Figure 2: Growth Sentiment vs. Growth for Matched and Unmatched Newspapers

In the next three columns of Table 8, we include the economic aggregates and well as their interactions with the ideological match variable. We find similar results to the first three columns. We find marginally significant results in columns 6 however, but these results are not consistent with partisan bias. This can be better seen in Figures 2, 3, and 4. Each figure plots the predicted level of sentiment as a function of the ideological match and the underlying economic aggregate. For growth, we find that matched and unmatched newspapers produce nearly identical tone. For unemployment, we find that unmatched newspapers report slightly more positively when unemployment has been improving and that matched newspapers report slight more positive when unemployment has been worsening. These effects are only a few percentage points and are vnot statistically significant. For inflation, we find that unmatched newspapers report slightly more positively when inflation has been improving and that matched newspapers report slight more positively when inflation has been worsening. These differences are marginally statistically significant and are not consistent with bias since bias suggests that the ideologically matched newspaers should uniformly report more favorably.

5.2 Media Bias in Coverage

While newspapers are not very biased in thier tone, they may be biased in the coverage they devote to the economy. In the first three columns of Table 9, we regress the share of coverage devoted to growth, unemployment, and inflation, on the economic variables, whether the newspaper is ideologically matched with the prime minister, and an interaction between ideological match and the economic variables. Our results suggests that




Figure 3: Unemployment Sentiment vs. Changes in Unemployment for Matched and Unmatched Newspapers

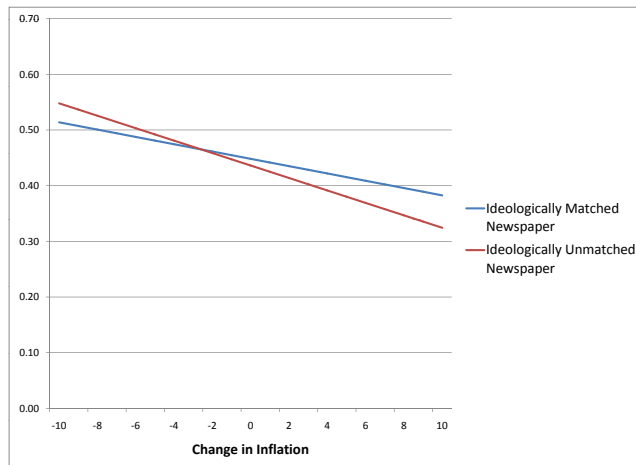


Figure 4: Inflation Sentiment vs. Changes in Inflation for Matched and Unmatched Newspapers

Dependent Variable:	Growth Coverage	Unemployment Coverage	Inflation Coverage
Independent Variables:			
Ideological Match	0.002 (0.014)	0.047* (0.023)	0.011 (0.017)
Growth (yearly)	-0.012** (0.004)		
Unem. (yearly)		0.022*** (0.002)	
Inf. (yearly)			0.013*** (0.003)
Growth * Ideo. Match	-0.002 (0.003)		
Unem. * Ideo. Match		-0.008** (0.002)	
Inf. * Ideo. Match			-0.007 (0.004)
Number of Months	6649	6656	6660
Number of Newspapers	32	32	32
Number of Countries	16	16	16
R-Squared	0.7788	0.7431	0.6461

Table 9: Media Bias in Coverage

the ideological orientation of newspapers does not color the share of coverage they devote to growth and inflation—while we find that newspapers focus on bad news, unmatched newspapers are not more likely to do so for growth and inflation. For unemployment, we do find that ideology colors newspapers coverage.

These patterns are expanded on in Figures 5, 6, and 7. These figures report the predicted share of coverage, as a function of the economy and whether there is an ideological match between the newspaper and the current prime minister. We find little difference between matched and unmatched papers for growth and inflation. For unemployment, we find that while both match and unmatched newspapers increase their coverage of unemployment as unemployment increases, for low levels of unemployment, the share of coverage by matched papers is higher than the share of coverage by unmatched newspapers, and high levels of unemployment, the reverse is true. The inflection point happens very close to the average unemployment rate in the sample, suggesting that matched papers report relatively more on unemployment when unemployment is below average and unmatched newspapers report relatively more on unemployment when unemployment is above average.

Taken together, we find very little bias in the tone of media coverage, but we find some bias in the amount of media coverage for unemployment, but not for growth or inflation. Interestingly, our results on the share of coverage comport with [Larcinese, Puglisi and Snyder \(2011\)](#), who also find that in the United States, unmatched papers focus more on unemployment when unemployment is high.

5.3 Possible Concerns and Robustness Checks

One concern is that our finding of no media bias in tone occurs because newspapers we collected are not sufficiently ideologically differentiated. Readers familiar with U.S.

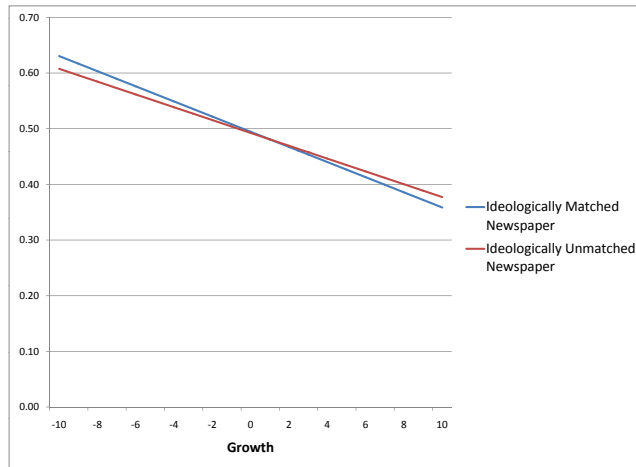


Figure 5: Growth Coverage vs. Growth for Matched and Unmatched Newspapers




Figure 6: Unemployment Coverage vs. Unemployment for Matched and Unmatched Newspapers

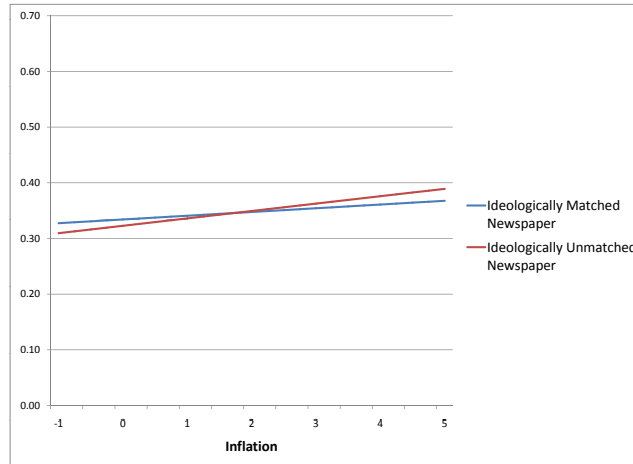


Figure 7: Inflation Coverage vs. Inflation for Matched and Unmatched Newspapers

newspapers may view the New York Times and the Wall Street Journal as relatively moderate (although [Gentzkow and Shapiro \(2006\)](#) estimate their ideological orientation to be quite different). Here, our study of share of media coverage serves as a type of manipulation check—if our finding of no difference in media tone is driven by selecting pair of papers that are ideologically similar, then we should not be able to detect differences in coverage between left and right wing newspapers.

A second concern is our binary measures of newspaper and prime minister ideology. To address this problem, we coded the thirty two newspapers in our sample on a five point ideology scale based on various online sources.⁸ We also used a continuous measure of prime minister ideology based on the comparative manifesto project. The scales made the analyses somewhat more difficult to interpret. We specified the regression by interacting newspaper ideology and prime minister ideology and their interaction with the economic variables, focusing on the triple interaction between the economy, newspaper ideology, and prime minister ideology (this is essentially the same approach as [Larcinese, Puglisi and Snyder \(2011\)](#)). We found essentially similar results.

We considered a number of other robustness checks. We coded government ideology rather than prime minister ideology.⁹ Our choice of whether to use levels or changes in unemployment and inflation was based on our findings in Section 4, but we considered alternative specifications. The share of coverage can be coded in multiple different ways—the share of words, sentences, or articles (using various cutoff for how much an

⁸The five point scale included the categories very liberal, somewhat liberal, moderate, somewhat conservative, and very conservative.

⁹Government ideology was computed as the ministry-weighted mean ideological position of the parties in government.

article must discuss growth, unemployment, and inflation to be considered a growth, unemployment, and inflation articles). These different robustness checks did not effect the results substantially.

6 Do Newspaper Mediate the Economic Vote?

Do voters vote based on the economy, or based on media reports of the economy? To the extent we observe a correlation between media reports of the economy and the vote share of the incumbent prime ministers, is this because voters are affected by the newspaper coverage, or because the newspaper coverage simply provides a proxy for the actual state of the economy. To investigate this, we perform a mediation analysis, following the approach originally developed by [Baron and Kenny \(1986\)](#).

We estimated the following structural equations,

$$V_n = \beta_0 + \beta_1 Growth_n + \beta_2 * Unem_n + \beta_3 Inf_n \quad (1)$$

$$+ \beta_4 GrSent_n + \beta_6 UnSent_n + \beta_7 InfSent_n + \epsilon_n$$

$$GrSent_n = \gamma_0^G + \gamma_1^G Growth_n + \eta_n^G \quad (2)$$

$$UnSent_n = \gamma_0^U + \gamma_1^U Unem_n + \eta_n^U \quad (3)$$

$$InfSent_n = \gamma_0^I + \gamma_1^I Inf_n + \eta_n^I \quad (4)$$

We can plug (2), (3), and (4) into (1) to obtain,

$$V_n = \beta_0 + (\beta_1 + \beta_4\gamma_1^G)Growth_n + (\beta_2 + \beta_5\gamma_1^U)Unem_n + (\beta_3 + \beta_6\gamma_1^I)Inf_n \quad (5)$$

$$+ \epsilon_n + \beta_4\eta_n^G + \beta_5\eta_n^U + \beta_6\eta_n^I$$

Here, we have that $\beta_1 + \beta_4\gamma_1^G$, $\beta_2 + \beta_5\gamma_1^U$, and $\beta_3 + \beta_6\gamma_1^I$ represent the total effects of Growth, Unemployment, and Inflation; β_1 , β_2 , and β_3 represent the direct effects; $\beta_4\gamma_1^G$, $\beta_5\gamma_1^U$, $\beta_6\gamma_1^I$ represent the indirect effects; and $\epsilon_n + \beta_4\eta_n^G + \beta_5\eta_n^U + \beta_6\eta_n^I$ is an error term. To simplify the setup, we assume that ϵ_n and $(\eta_m^G, \eta_m^U, \eta_m^I)$ are independent for all n and m which allows us to compute the standard errors for the total, direct, and indirect effects using the delta method.

We report results for the regressions in (5), (1), (2), (3), and (4) in Table 10. In column (1), we see that growth, unemployment, and inflation are statistically significant and have the expected signs. In column (2), the sentiment variables are added to the analysis. We see that the sentiment on growth and actual unemployment are the only statistically significant variables. In columns (3), (4), and (5), we find that all three measures of economic performance are correlated with the respective components of sentiment. Combining these results, we have evidence for an effect of growth that is entirely mediated by newspaper sentiment and an effect of unemployment that is not mediated by newspaper sentiment.

DV:	PM Vote Sh.	PM Vote Sh.	Growth Sent.	Unem. Sent	Inf. Sent.
Constant	36.950*** (1.644)	27.814*** (7.837)	-	-	-
Growth	0.690** (0.226)	-0.181 (0.493)	0.023*** (0.000)		
Unem.	-0.485* (0.196)	-0.886* (0.397)		-0.001** (0.000)	
Inf.	-0.141* (0.065)	0.438 (0.444)			-0.004** (0.000)
Gr Sent.		19.752* (9.481)			
Un. Sent.		0.882 (9.107)			
Inf. Sent.		5.850 (9.005)			
R^2	0.075	0.119	0.600	0.251	0.253
N	293	89	6650	6652	6606
Newspaper FE?			X	X	X

Table 10: Structural equation model for the effect of economic performance on elections, as mediated by newspaper sentiment.

	Growth	Unemployment	Inflation
Total Effect	0.273 (0.421)	-0.887* (0.398)	0.413 (0.438)
Direct Effect	-0.181 (0.493)	-0.886* (0.397)	0.438 (0.444)
Indirect Effect	0.454* (0.218)	-0.001 (0.013)	-0.025 (0.038)

Table 11: Mediation Analysis—Calculated using the linear structural model estimated in Table 10. Standard errors are calculated using the delta method.

We expand on this analysis in Table 11 and compute the direct, indirect, and total effects, along with their associated standard errors. Starting with growth, we find that

the direct effect is negative, but not statistically significant, and that the indirect effect is positive and statistically significant. The total effect of growth is positive, but not statistically significant. For unemployment, we find that the direct effect is negative and statistically significant, the indirect effect is negative but not statistically significant, and the total effect is negative and statistically significant. For inflation, none of the effects are statistically significant.

Taking these results at face value, growth has an effect that is entirely transmitted through newspaper coverage of the economy, unemployment has an effect that is not at all transmitted through newspaper coverage, and inflation has no effect. Taking the results at face value again, voter punish incumbents based on media reports of low growth and through their direct experiences with unemployment.

The fact that the total effect for growth is statistically insignificant may initially seem troubling. We can compare the total effects to the results in column (1) of Table 10. According to the model we estimated, the coefficient on growth in column (1) should equal the total effect. This could fail if the assumptions of the linear structural model are incorrect or if the coefficients are not invariant to the sample.

To check this, we estimated the model in column (1) of Table 10 on the sample restricted to the 89 observations for which we observe media data (not shown). The results indicate that growth is positive and not statistically significant, unemployment is negative and statistically significant, and inflation is positive but not statistically significant. This suggests that (a) the total effects we calculated comport very well with the linear regression restricted to the sample for which we observe newspaper data, and (b) the lack of statistical significance for growth and inflation seems to be driven by the small sample size for which we observe both newspaper data and an election results.

7 Conclusion

This paper is only the first-cut at a large, original dataset of machine-coded newspaper sentiment based on over 2 million newspaper articles in 32 newspapers in 16 countries. The scale of our dataset not only allows us to generalize more securely but also to conduct analyses that were previously not possible. This paper tentatively investigated the validity and reliability of its economic sentiment measures. Economic sentiment tracks fairly closely to macro-economic aggregates themselves and newspapers based in the same country correlate fairly closely with each other. We also see one interesting null effect. The tone of newspaper sentiment does not depend on whether the newspaper and the government share an ideological orientation. Newspapers do not report more positively about the economy when their co-partisans hold the prime-ministership. Initial results,

however, do suggest that partisanship expresses itself via another means. Newspapers run more stories on the economy during boom times when a co-partisan is in office and even more stories during a recession when a party of an opposing ideology is in office. Finally, a structural equations model revealed that the effect of economic growth on the vote is fully mediated by reporting on the economy while, in contrast, inflation and unemployment directly influence the vote.

A Appendix: Dates of Coverage

Country	Newspaper	Language	Relative Partisanship	Coverage
Austria	<i>Der Standard</i>	German	Left	Dec. 2007 – Aug. 2013
Austria	<i>Die Presse</i>	German	Right	Apr. 2004 – Aug. 2013
Australia	<i>The Age</i>	English	Left	Jan. 1991 – Sept. 2013
Australia	<i>The Herald Sun</i>	English	Right	Jan. 1987 – Aug. 2013
Canada	<i>The Globe and Mail</i>	English	Right	Nov. 1977 – July 2013
Canada	<i>Toronto Star</i>	English	Left	Sept. 1985 – Aug. 2013
France	<i>Le Monde</i>	French	Left	Jan. 1990 – Dec. 2012
France	<i>Le Figaro</i>	French	Right	Jan. 1997 – Aug. 2013
Germany	<i>Die Zeit</i>	German	Left	Nov. 2008 – Apr. 2014
Germany	<i>Frankfurter Allgemeine</i>	German	Right	Jan. 2010 – Sept. 2013
Ireland	<i>The Irish Independent</i>	English	Right	Oct. 2006 – Aug. 2013
Ireland	<i>The Irish Times</i>	English	Left	Jun. 1992 – Dec. 2012
Israel	<i>Globes</i>	English	Left	June 1996 – Sept. 2013
Israel	<i>The Jerusalem Post</i>	English	Right	Jan. 1989 – Aug. 2013
Italy	<i>Corriere della Serra</i>	Italian	Right	Jan. 2009 – Aug. 2013
Italy	<i>La Stampa</i>	Italian	Left	Jan. 1992 – Dec. 2012
Japan	<i>Daily Yomiuri</i>	English	Right	Sept. 1989 – Mar. 2013
Japan	<i>Nikkei Weekly</i>	English	Left	June 1980 – Sept. 2013
Luxembourg	<i>Le Quotidien</i>	French	Left	Apr. 2008 – Dec. 2013
Luxembourg	<i>Le Fax d'Agefi</i>	French	Right	Dec. 2009 – Apr. 2014
New Zealand	<i>New Zealand Herald</i>	English	Right	Nov. 1998 – Aug. 2013
New Zealand	<i>The Press</i>	English	Left	June 1996 – Aug. 2013
Portugal	<i>Correio da Manha</i>	Portuguese	Left	June 2012 – Aug. 2013
Portugal	<i>Jornal de Noticias</i>	Portuguese	Right	July 1997 – June 2013
Spain	<i>El Mundo</i>	Spanish	Right	July 2002 – Aug. 2013
Spain	<i>El Pais</i>	Spanish	Left	Apr. 1996 – Dec. 2012
Switzerland	<i>Neue Zürcher Zeitung</i>	German	Right	May 1993 – Dec. 2012
Switzerland	<i>Tages-Anzeiger</i>	German	Left	Sept. 1997 – Sept. 2013
United Kingdom	<i>The Guardian</i>	English	Left	July 1984 – July 2013
United Kingdom	<i>The Times (London)</i>	English	Right	Jul. 1985 – Dec. 2012
United States	<i>New York Times</i>	English	Left	Sept. 1989 – Sept. 2013
United States	<i>Wall Street Journal</i>	English	Right	June 1979 – Dec. 2013

Table 12: Dates of coverage for 32 newspapers.

References

- Achen, Christopher H. and Larry M. Bartels. 2004. “Musical Chairs: Pocketbook Voting and the Limits of Democratic Accountability.” *Working Paper, Princeton University* .
- Arceneaux, Kevin, Martin Johnson, René Lindstädt and Ryan J Vander Wielen. 2015. “The Influence of News Media on Political Elites: Investigating Strategic Responsiveness in Congress.” *American Journal of Political Science* .
- Baron, Reuben M and David A Kenny. 1986. “The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations.” *Journal of personality and social psychology* 51(6):1173.
- Blood, Deborah J and Peter CB Phillips. 1995. “Recession headline news, consumer sentiment, the state of the economy and presidential popularity: A time series analysis 1989–1993.” *International Journal of Public Opinion Research* 7(1):2–22.
- Burckhardt, Philipp, Raymond Duch and Akitaka Matsuo. 2013. “Locating Economic Perceptions: the Geographical Distribution of Responses to Economic News DRAFT.” .
- De Boef, Suzanna and Paul M Kellstedt. 2004. “The political (and economic) origins of consumer confidence.” *American Journal of Political Science* 48(4):633–649.
- Della Vigna, Stefano and Ethan Kaplan. 2006. “The Fox News Effect: Media Bias and Voting.” *NBER Working Paper* .
- Duch, Raymond M. and Randy Stevenson. 2008. *The Economic Vote: How Political and Economic Institutions Condition Election Results*. Cambridge: Cambridge University Press.
- Eggers, Andrew. 2015. “Partisanship and Electoral Accountability : Evidence from the UK Expenses Scandal.” *Quarterly Journal of Political Science* .
- Erikson, Robert S and Christopher Wlezien. 2012. *The timeline of presidential elections: How campaigns do (and do not) matter*. University of Chicago Press.
- Evans, Geoffrey and Robert Anderson. 2006. “The Political Conditioning of Economic Perceptions.” *Journal of Politics* 68(1):194–207.
- Gelman, Andrew and Gary King. 1993. “Why Are American Presidential Election Campaign Polls So Variable When Votes Are So Predictable?” *British Journal of Political Science* 23:409–451.

- Gentzkow, Matthew and Jesse M. Shapiro. 2006. "Media Bias and Reputation Jesse M. Shapiro." *Journal of Political Economy* 114(2):280–316.
- Healy, Andrew and Gabriel S. Lenz. 2014. "Substituting the End for the Whole: Why Voters Respond Primarily to the Election-Year Economy." *American Journal of Political Science* 58(1):31–47.
- Hopkins, Daniel J. 2010. "Does Newspaper Coverage Influence or Reflect Public Perceptions of the Economy?" *Working Paper, APSA 2010* .
- Kayser, Mark A. 2014. The Elusive Economic Vote. In *Comparing Democracies 4*, ed. Lawrence LeDuc, Richard G. Niemi and Pippa Norris. Thousand Oaks, CA: Sage.
- Kayser, Mark A. and Christopher Wlezien. 2011. "Performance Pressure: Patterns of Partisanship and the Economic Vote." *European Journal of Political Research* 50(3):365–394.
- Kayser, Mark Andreas and Arndt Leininger. 2015. "Vintage Errors: Do Real-Time Economic Data Improve Election Forecasts?" *Research & Politics* 2(3).
- Larcinese, Valentino, Riccardo Puglisi and James M Snyder. 2011. "Partisan Bias in Economic News : Evidence on the Agenda-Setting Behavior of US Newspapers." *Journal of Public Economics* 95(9).
- Lewis-Beck, Michael S., Nicholas F. Martini and D. Roderick Kiewiet. 2013. "The Nature of Economic Perceptions in Mass Publics." *Electoral Studies* .
- Lott Jr, John R and Kevin A Hassett. 2014. "Is newspaper coverage of economic events politically biased?" *Public Choice* 160(1-2):65–108.
- Nadeau, Richard, Michael S Lewis-Beck and Éric Bélanger. 2013. "Economics and elections revisited." *Comparative Political Studies* 46(5):551–573.
- Powell, Bingham G. and Guy D. Whitten. 1993. "A Cross-National Analysis of Economic Voting: Taking Account of the Political Context." *American Journal of Political Science* 37:391–414.
- Puglisi, Riccardo and James M Snyder. 2014. "The balanced US press." *Journal of the European Economic Association* .
- Roberts, Marilyn and Maxwell McCombs. 1994. "Agenda setting and political advertising: Origins of the news agenda." *Political Communication* 11(3):249–262.

- Sides, John and Lynn Vavreck. 2014. *The gamble: Choice and chance in the 2012 presidential election*. Princeton University Press.
- Soroka, Stuart N. 2006. "Good News and Bad News: Asymmetric Responses to Economic Information." *Journal of Politics* 68(2):372–385.
- Stanig, Piero. 2013. "Political Polarization in Retrospective Economic Evaluations During Recessions and Recoveries?" *Electoral Studies* 32(4):729–745.
- Wlezien, Christopher. 2015. "The Myopic Voter? The Economy and US Presidential Elections." *Electoral Studies* Forthcoming.
- Wlezien, Christopher, Mark Franklin and Daniel Twigg. 1997. "Economic Perceptions and Vote Choice: Disentangling the Endogeneity." *Political Behavior* 19(1):7–17.
- Zaller, John. 2004. Floating Voters in US Presidential Elections, 1948-2000. In *Studies in Public Opinion: Attitudes, Nonattitudes, Measurement Error, and Change*, ed. Paul Sniderman and Willem Saris. Princeton: Princeton University Press pp. 166–214.