



Campaigns and the prediction of election outcomes: Can historical and campaign-period prediction models be combined?

Éric Bélanger*, Stuart Soroka

Department of Political Science, McGill University, 855 Sherbrooke Street West, Montreal, Quebec H3A 2T7, Canada

ARTICLE INFO

Article history:

Received 21 February 2012

Received in revised form 5 July 2012

Accepted 6 July 2012

Keywords:

Election campaigns

Electoral forecasting

Campaign effects

ABSTRACT

There are two approaches to predicting election outcomes: (1) a historical approach, which uses past election results alongside macroeconomic and political variables to forecast election results up to a year in advance, and (2) a campaign-oriented approach, which uses current campaign trends to forecast vote shares at the end of the campaign. They are in some way at odds—one approach says the campaign doesn't matter, the other focuses entirely on the campaign. This article considers whether the two approaches might be usefully combined; it considers whether the prediction errors in historical models may be related to trends during the campaign. That possibility is tested here using 17 elections in the US, UK and Canada, combining historical predictions and automated content analyses of campaign-period media content. Results suggest that campaigns do not account for errors in the historical predictions; but there may be other ways in which campaigns matter in conjunction with historical models.

© 2012 Elsevier Ltd. All rights reserved.

Political scientists have long been interested in trying to predict the outcome of future democratic elections. This exercise has value for at least three reasons. First, these forecasting models satisfy—with some degree of success—the natural curiosity that arises each time a new electoral contest is about to get under way. Second, and more importantly, these models put to the test the predictive power of theories about which factors influence the behavior of voters. If indeed we are able to identify variables that explain past election outcomes, then we should expect these same variables to be able to predict future elections with some level of accuracy. Third, prediction models address more or less directly the core question of whether or not election campaigns matter to the outcome.

There are, broadly speaking, two approaches to predicting election outcomes, and each holds different

assumptions about whether and how campaigns matter. The historical approach focuses on long-term macroeconomic factors combined with other political variables capturing, for instance, the popularity of the incumbent. The specifics of these forecast models vary—indeed, the relative strength of marginally different prediction models is regularly the focus of pre-election conference symposia and journal issues (for the most recent examples in the US and UK contexts, see [Campbell, 2008a](#); [Gibson and Lewis-Beck, 2011](#)). But the models are similar in their focus on the historical effects of long-term dynamics (i.e., the “fundamentals”) on election outcomes. Their success confirms the importance of macroeconomics and other long-term political factors in vote outcomes. Successfully predicting outcomes using pre-campaign information also speaks to the possible insignificance of election campaigns.

That said, the success of historical forecast models varies, both across countries and across elections. And there is a growing body of work showing that campaigns can matter—that is, that sometimes a campaign can lead to vote outcomes at the end that are markedly different than

* Corresponding author. Tel.: +1 514 398 4400; fax: +1 514 398 1770.

E-mail addresses: eric.belanger3@mcgill.ca (É. Bélanger), stuart.soroka@mcgill.ca (S. Soroka).

they were at the beginning (for recent reviews, see Brady and Johnston, 2006; Wlezien, 2010). This work does not often present itself as being about “prediction,” per se. Capturing campaign dynamics has tended to take some time, and so time-series analysis of vote intentions and campaign information has tended to take place after the campaign, in the form of explanatory rather than prediction models. That said, the models themselves are essentially about prediction. And recently developed automated content analytic techniques have led to a small, but growing, body of work that focuses on prediction, using campaign-period information, during the campaign (e.g., Soroka et al., 2009).

Both the historical and campaign approaches have seen some success; they are however clearly at odds with one another. A politico-economic approach suggests that campaigns barely matter; the campaign approach suggests otherwise, and recognizes but does not directly capture long-term effects. The truth most likely lies somewhere in between. This article represents a first attempt, then, at exploring whether the two approaches can be combined. That is, we begin with historical forecast models of campaign outcomes, and then verify whether using media content-analytic information from the campaign helps to account for the historical models' prediction errors. We do so for a significant sample of recent national elections in each of the US, UK and Canada. Our historical models draw directly on existing work in the field; our campaign-period content analyses rely on over 63,000 media stories collected in these three countries, analyzed for tone using Lexicoder.

Our basic aim is not to improve the specification of the historical models per se, since campaign information is not available far enough ahead of election day. Including campaign information into long-term models would thus mean sacrificing their lead time—something that would render their predictions more or less trivial. Rather, we wish to explore, in an ex post fashion, the extent to which the prediction errors of historical forecast models can be accounted for by campaigns in a meaningful manner. Put differently, we want to try to answer the following question: To what extent are the errors in the historical prediction models related to campaign dynamics?

1. The historical approach to prediction

What we call the “historical approach” to election forecasting refers to the various vote function models developed to predict election outcomes on the basis of political and macroeconomic variables. These politico-economic forecast models typically suggest that the vote for the incumbent candidate (or party) is a function of macroeconomic conditions, government approval, and the “cost of ruling”—variables that are sometimes referred to as “the fundamentals” (for reviews, see Campbell and Garand, 2000; Lewis-Beck, 2005; Lewis-Beck and Tien, 2011). Such models have been developed for several of the major Western democracies, with the US, the UK, and France having received the most attention to date (Lewis-Beck and Tien, 2011).

There are two other important features of these models that ought to be highlighted here. The first is that these are

long-term time-series models, meaning that their statistical estimation relies on a set of data covering several decades of election outcomes, typically going as far back as the 1950s. It is for this reason that these models can be considered as historical in nature: their predictions of future election outcomes are based on patterns of voter behavior observed over the past decades. These models' second key feature is their lead time. So as not to provide a trivial forecast, historical models use data collected a few months before the start of the election campaign to make their prediction—typically between three to six months in advance, sometimes a full year. In other words, the predictions are based on incumbent approval and macroeconomic data measured well before the onset of campaigns.

It is the latter feature—lead time—which has led to the charge that historical forecast models do not take campaigns into account in their predictions (e.g., Greene, 1993; Tetlock, 2005; van der Eijk, 2005; Silver, 2012). If political and macroeconomic variables measured several months in advance of an election campaign can predict its outcome correctly most of the time, then surely these models mean that campaigns do not matter. Election forecasters' response to this criticism has been to say that campaigns do matter, just not for prediction. As Lewis-Beck and Tien (2008) have pointed out, campaigns are “constant” in these models: campaign-free democratic elections do not exist, so we have no way of knowing if these models would work as well (or better) were no candidates or parties campaigning. In addition, candidates or parties all “play to win,” thus producing a campaign equilibrium where their competing efforts tend to cancel out, creating the aggregate-level illusion that campaigns have no net effects. The basis for these claims also comes from work having shown that the campaign process usually tends to bring the election's outcome in line with that predicted by the fundamental variables, because it leads voters to pay attention to the core issues of macroeconomic and political performance (see Gelman and King, 1993; Holbrook, 1996; Arceneaux, 2006; Campbell, 2008b; Erikson and Wlezien, 2012). In other words, campaigns “deliver” the fundamentals and they affect the election outcome through them. Thus, campaigns matter—but not independently of the variables already included in the historical models.

A different viewpoint holds that the prediction errors of long-term models are not randomly distributed and that deviations from the forecasts are driven by independent, or direct, campaign effects on top of those associated with the fundamental variables (e.g., Vavreck, 2009). Early attempts at adding explicit campaign variables to US forecast models have yielded no statistically significant effects; these variables were the incumbent candidate's strength in the first primary election, the change in his/her polling advantage from June to Labor Day, and the “quality” of his/her campaign (Lewis-Beck and Tien, 2002, 2008). A more recent attempt, using Vavreck's (2009) proposed typology of “good” and “bad” campaigns, proved more successful (Nadeau and Lewis-Beck, 2012). According to Vavreck, forecast errors can be partly explained by whether the incumbent candidate (or the challenger) runs a “clarifying”

campaign on the economy when appropriate. That is to say, when the economy is doing well, the incumbent should talk about it and clarify his (or his party's) role in helping it, thus reducing the prediction error. When it is not doing well, it is the challenger who should talk about it and clarify the failure of the incumbent in preventing the economic decline, reducing the forecast error as well. If the incumbent or the challenger runs a bad campaign (that is, does not run a clarifying campaign on the economy when appropriate), the forecast error is increased. For Vavreck (2009: 107–109), then, it is the content of the candidates' campaign message on the economy that matters, and that ought to account for at least some of the prediction errors (see also Nadeau and Lewis-Beck, 2012).

2. Reconsidering the impact of campaigns

In contrast to work on historical approaches to election prediction, a relatively recent body of work on election forecasting is focused on how campaigns can influence election outcomes. It follows from fine-grained studies of public opinion showing that what goes on during the campaign period can have a significant impact on vote intentions and on the final result (e.g., Johnston et al., 1992, 2004; Just et al., 1996; Norris et al., 1999). The sound and fury of campaigns—including, but not restricted to, what the candidates or parties say—is found to be consequential. This work on the campaign period does not present itself as being about “prediction”, but its core conclusion that campaigns matter implies that if we have a good idea about the overall content of a campaign (especially the valence of that content), we will be able to predict its impact on the outcome.

A number of recent studies have started to take this next step—prediction—by focusing on how well the campaign tracks voting intentions. These studies have relied on media data as a measure of what goes on in the campaign. There are very good reasons for doing so. The modern election campaign is nearly by definition a media campaign. Recent work also shows that voters learn from media about election issues and candidates (Weaver, 1996; Brians and Wattenberg, 1996; Mendelsohn, 1996; Druckman, 2004). Media may thus play a leading role in campaign dynamics; or, alternatively, media may simply be a good reflection of what is going on in the campaign. Articles are written by journalists, after all, highly attentive to campaign events, to party strategies and proposals, and to shifts in public opinion as well. Regardless of whether media are affecting or reflecting opinion, they may be useful tools for campaign-period predictions.

What is it exactly in media content that is connected to vote shares in the short term? In some cases, there appears to be a connection between issue coverage and party support—where parties dominate on one issue or another, coverage of those issues tracks party support in the polls (see, e.g., Cutler et al., 2004; Walgrave and de Swert, 2004; Sheaffer and Weimann, 2005; Nadeau et al., 2008, 2010). A more general (and more widely applicable) approach is to look for a connection between vote intentions and the “tone” of media content. Indeed, within the context of election and referendum campaigns, media tone has recently been shown to track vote intentions very closely

(e.g., Fan, 1988; Mendelsohn and Nadeau, 1999; Dobrzynska et al., 2003; de Vreese and Semetko, 2004; Soroka et al., 2009; Hopmann et al., 2010).

Do campaigns matter to election prediction, or not? The campaign-period prediction models seem to suggest that they do, and that historical models may be missing a part of the story. Yet existing attempts at incorporating campaigns into politico-economic models have yielded mixed results, with only one successful attempt. Probably both approaches to election prediction are right to some extent, or at least campaigns account for part (if not all) of what the long-term macroeconomic and political variables do not.

We address this core question using media measures of the campaign. To help guide our analysis, we can lay out two extreme but plausible possibilities. The first is that campaigns account for all of the error in politico-economic predictions; that is, errors from politico-economic forecasts are entirely about change during the campaign. If this is the case, then we expect that the relationship between prediction error and tone is negative. That is, if media tone during campaigns matters, an under-estimation of incumbent performance by the historical models should be associated with a positive tone (on average) in the media's coverage of the incumbent's campaign. Inversely, an over-estimation of incumbent vote share should be associated with a negative media tone (on average) toward that incumbent during the campaign period.

The second possibility is that campaigns account for none of the error in historical predictions; that is, errors from politico-economic forecasts are just errors, unaccounted for by the campaign—leaving open the further possibility that the campaign's influence is indirectly felt through the fundamental variables' impact on the vote. In this case, of course, there should be no relationship between media tone and prediction error.

To see whether forecast error and tone are related to each other, we look at 17 recent electoral contests held in three major Western democracies. The countries selected are the US (seven elections, 1984–2008), the UK (four elections, 1997–2010) and Canada (six elections, 1993–2008). Despite their obvious institutional differences, these three countries share some important features. Of importance to our purpose, American, British, and Canadian electoral campaigns show a lot of similarities in their dynamics, and especially in the kind of campaign effects that are usually observed, be they bounces, bumps, or some combination of both (Wlezien, 2010: 113–114). In addition, there currently exist comparable long-term forecasting models for each of them, which we can use to generate prediction errors.

Individual elections have been selected here as a function of the availability of media data that are amenable to automated content analysis. Our starting point, however, is to take each of these three historical models and estimate their forecast errors for the elections being examined, a task to which we now turn.

3. Prediction errors in three countries

The historical models we use in this article are those by Abramowitz (2008), Bélanger et al. (2005), and Bélanger

Table 1
Prediction models.

	US (1948–2008)	UK (1955–2010)	CA (1953–2008)
Approval	.108*** (.021)	.433** (.099)	.481** (.153)
Economy	.547*** (.114)	-.654* (.267)	-2.168*** (.489)
Time	-4.429*** (1.017)	-1.777 (1.231)	-2.229 (1.211)
Constant	51.669*** (.831)	31.215*** (5.744)	42.099*** (8.336)
N	16	15	19
Rsqr	.917	.773	.701
Rsqr Adj	.896	.711	.641

* $p < .05$; ** $p < .01$; *** $p < .001$. Cells contain OLS regression coefficients with standard errors in parentheses. Economy is captured using GDP in the US, inflation in the UK, and unemployment in Canada. Time is captured using the number of terms a President/Government is in power in the US and UK, and the number of months a Government is in power in Canada.

and Godbout (2010). They forecast US presidential elections, UK general elections, and Canadian federal elections, respectively. These three models share the exact same structure, which facilitates comparisons. They use the same dependent variable: incumbent vote share in percentage (recalculated in the US as the two-party vote share). They include the same three long-term independent variables: government approval, the state of the national economy, and a “time-for-change” variable (government longevity). Each independent variable is measured several months in advance of election day. These models’ theoretical expectations are thus similar: the higher the approval rate and the stronger the economy, the more votes the incumbent candidate (or party) should receive; and the longer the incumbent (or his/her party) has been in power, the less votes he/she/it should receive due to voter fatigue with that incumbent.

Although these models’ independent variables are the same, there are variations in the measures being employed that should be kept in mind. For the US, net approval of the President in June, GDP growth during the second quarter of the election year, and a dummy variable indicating whether or not the incumbent party is seeking a third consecutive mandate (or more) are used as measures. For the UK model, the indicators being used are the approval of the government’s record and the inflation rate, both measured six months before the election, as well as the number of consecutive terms the incumbent party has been in power. For the Canadian case, the measures are the popularity of the incumbent party in terms of vote intentions three months before the election, the unemployment rate for the second quarter before the election, and the logged number of consecutive months the incumbent party has been in power.

We run the historical models on their own, for each country, in Table 1. As expected given their respective track record, the models perform very well. The parameter estimates are almost always statistically significant.¹ That said, there are some important cross-country differences in model fits. The US forecast model stands out with its very high R-squared values (in the .90s) while the other two models display more moderate ones (in the .60s–.70s

¹ Only the term variables in the UK and Canada fail to reach statistical significance in the Table, although they are all significant if one adopts a much less stringent test (cut-off value of .10, one-tailed).

Table 2
Actual and predicted incumbent vote shares, by election.

Country	Year	Actual	Predicted	Error
US	1984	59.2	57.247	-1.953
	1988	53.9	51.451	-2.449
	1992	46.5	47.964	1.464
	1996	54.7	57.158	2.458
	2000	50.3	54.318	4.018
	2004	51.2	53.526	2.286
	2008	46.3	44.435	-1.865
UK	1997	30.7	29.658	-1.042
	2001	40.7	44.321	3.621
	2005	35.3	38.389	3.089
	2010	29.0	36.045	7.045
CA	1993	16.0	27.248	11.248
	1997	38.5	37.648	-.852
	2000	40.8	45.546	4.746
	2004	36.7	33.183	-3.517
	2006	30.2	33.312	3.112
	2008	37.6	37.479	-.121

range). That is to be expected, given that the UK and Canada have multiparty systems that complicate voters’ choice, compared to the clear-cut opposition between Republicans and Democrats found in the US.²

Table 2 presents the error from each prediction. The sample of national elections is restricted here to only those for which we have campaign data available (a total of 17). The forecast error is simply the difference (in percentage points) between the predicted vote share and the actual share received by the incumbent. For each election, our strategy for generating the predicted share of the vote is to run each model using all the data except for the year being forecast (i.e., they are “out-of-sample” forecasts). As can be seen, the errors are relatively small in general—a mean absolute error of 2.36 for the US, 3.70 for the UK, and 3.93 for Canada. That said, there is some variation that could be accounted for by the campaigns. The 2000 US presidential election, the 2010 UK general election, and the 1993 and 2000 Canadian federal elections are cases that particularly stand out in this respect.

Can we account for these prediction errors using media measures of the campaign? It is this central question that we now examine.

4. Tone of media coverage and the prediction error

We capture the content of an election campaign here using the tone of media coverage of major parties and leaders in campaign-period newspaper articles. For Canada and the UK, we include all stories related to the campaign from the day the election was called until the day of the election. For the US, we use all election stories from September 1st to election day.

In each case, campaign-related stories are identified in Nexis by full-text searching for the word “election” or any major party or leader name. Stories unrelated to the election are excluded manually. In the US, we use the New York

² Of course, it may not be the two-party system itself so much as a host of related factors, including the powerful role of party identification in the US.

Times and the Washington Post; in the UK, the Times, Guardian, Independent, Sun, and Daily Telegraph; in Canada, the Calgary Herald, Globe and Mail, Montreal Gazette, National Post, Toronto Star, and Vancouver Sun. The aim is to select the major national broadsheets in every case. In the US, due to the volume of data, we use just the two major newspapers of record; in the UK, we use five. The Canadian data rely on English-language newspapers only, due to limitations in automated coding (discussed below). So media data are limited in different ways in each country. Even so, we expect that content in these newspapers will accurately capture trends in coverage over the campaign (and across campaigns as well). In the end, we have roughly 25,000 stories across seven US elections, 12,000 stories across four UK elections, and 27,000 stories across six Canadian elections.

The tone of campaign stories is identified using the Lexicoder Sentiment Dictionary (LSD) in Lexicoder, a multi-platform Java-based automated content software available at lexicoder.com. The software applies a simple dictionary-based approach to content analysis—in this case, it counts the number of positive and negative words. The reliability of this approach is of course in large part a function of the dictionary; the LSD is discussed and tested in great detail in Young and Soroka (2012).

We attach tone to individual parties and leaders here using a relatively simple proximity analysis. That is, we count the number of positive and negative words that co-occur in the same sentences as party and leader names. We then use those counts to produce measures of net tone; or, rather, coefficients of imbalance (CI). The CI is drawn from work by Janis and Fadner (1943), and is intended to capture a combination of the volume and tone of content. The CI is calculated as follows:

$$C_f = (f^2 - fu)/rt, \text{ where } f > u, \quad (1)$$

$$C_u = (uf - f^2)/rt, \text{ where } f < u, \text{ and} \quad (2)$$

$$CI = C_f - C_u, \quad (3)$$

where f is the number of favorable items (e.g., words, phrases, articles), u is the number of unfavorable items, r is the number of relevant items and t is the total number of items (so r includes just the items categorized as positive or negative, while t includes all items, including those categorized as neutral). C_f is the coefficient of favorable imbalance, and C_u is the coefficient of unfavorable imbalance. Each component is discussed in detail in Janis and Fadner (1943); here, we just note that it provides the same basic trend as a regular measure of positive minus negative content, though the CI does a somewhat better job of capturing not just shifts in tone but in volume.³

³ In this case, using the CI actually produces only a minor improvement over the more standard “net tone” measure (% positive words – % negative words). That is, we get the same general results using a regular “net tone” measure.

We calculate the CI, as above, for both the incumbent and the challenger, and use them to generate the Incumbent CI Advantage—the gap between the CI for the incumbent and the challenger. In the American two-party system, the calculation of incumbent and challenger CI is relatively simple—one party is the incumbent, and the other is the challenger. In the UK and Canada, things are slightly more complicated. In the former, we include positive and negative words for the Liberal Democrats in the estimation of challenger CI, alongside Labour or the Conservatives of course. In Canada, we include the positive and negative words for all parties except the incumbent and the Bloc Québécois.⁴

The reliability of the LSD has been tested, against nine other content-analytic dictionaries, elsewhere (Young and Soroka, 2012). We nevertheless present some broad evidence here that the LSD identifies the “tone” of campaigns. First, Figs. 1–3 show weekly data for Incumbent CI Advantage across campaigns. The data show, for instance, the relatively positive coverage of Reagan in 1984, the precipitous decline of H.W. Bush in 1992, and of W. Bush in 2008. They show the steady erosion of positive coverage for the British Labour party from 2001 to 2010. They also show the negative coverage of the Martin Liberals in Canada in 2006, as compared to the previous two elections.

Fig. 4 takes a slightly more detailed look at shifts in campaign-period media tone, in the 2004 and 2008 US elections. The figure plots (smoothed) daily Incumbent CI Advantage against vote intentions.⁵ Media content is clearly not just about vote intentions, though the two do clearly track each other over time, as we should expect. We investigate the relationship between media and opinion more closely in Table 3, which shows results from a simple Granger causality test where each of vote intentions and media content is regressed on lagged values of itself as well as the other variable. Results suggest that media lead opinion, rather than the other way around. We do not interpret this as meaning that opinion simply follows media, of course. As has been discussed elsewhere (Soroka et al., 2009), it may simply be that media content, written by journalists highly attentive to the campaign, captures a trend before it emerges in opinion polls. And there are some differences between trends in media and trends in opinion—differences which we regard not as error in the media measure, but as actual difference between the content of media and the content of vote intentions. Regardless, we take the evidence in Fig. 4 and Table 3 as a further, albeit broad, indication that our automated media content measure captures an important part of what is going on during the campaign vis-à-vis party tone.

Having established, albeit in a rather general way, the apparent reliability of the media measures, we can proceed

⁴ Preliminary tests suggest that our results do not change fundamentally if we use only the largest challenger party, though choosing the largest challenger party is not as simple in Canada as in the UK. And note that we exclude the Bloc Québécois in the Canadian case because (a) the party runs only in Quebec, and (b) we use only English-language newspapers, in which coverage of the Bloc Québécois is rather slight.

⁵ Vote intentions data come from Wlezien and Erikson (2002, 2005) and Erikson and Wlezien (2012).

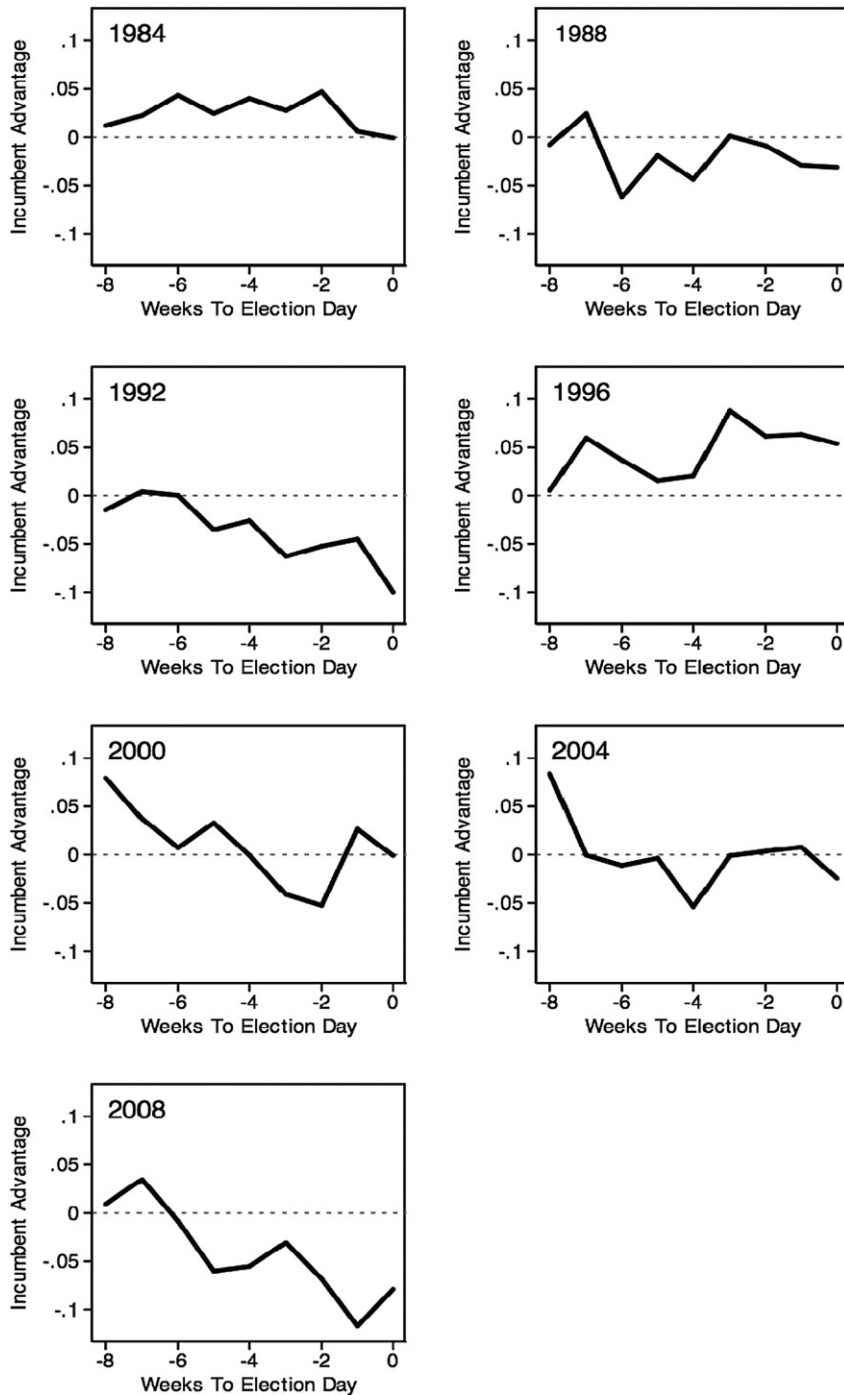


Fig. 1. Incumbent CI Advantage, US Election Campaigns.

to the main event: What is the relationship between campaign-period media tone and prediction errors from the historical models? Fig. 5 shows a scatterplot of prediction errors (y-axis) and mean Incumbent CI Advantage (x-axis) for each of the three countries. Appendix Table A1 presents a somewhat more formal analysis: it shows results from simple OLS models regressing prediction error on various campaign-period measures, including the mean

Incumbent CI Advantage (ninth row). These are pooled results, using all 17 elections, and including binary variables for the UK and Canada (with the US as the residual category).

Neither the figure nor the pooled regression models show any relationship between media tone and prediction errors in this sample of elections. These findings hold with many measures of media tone. In the figure, we focus on

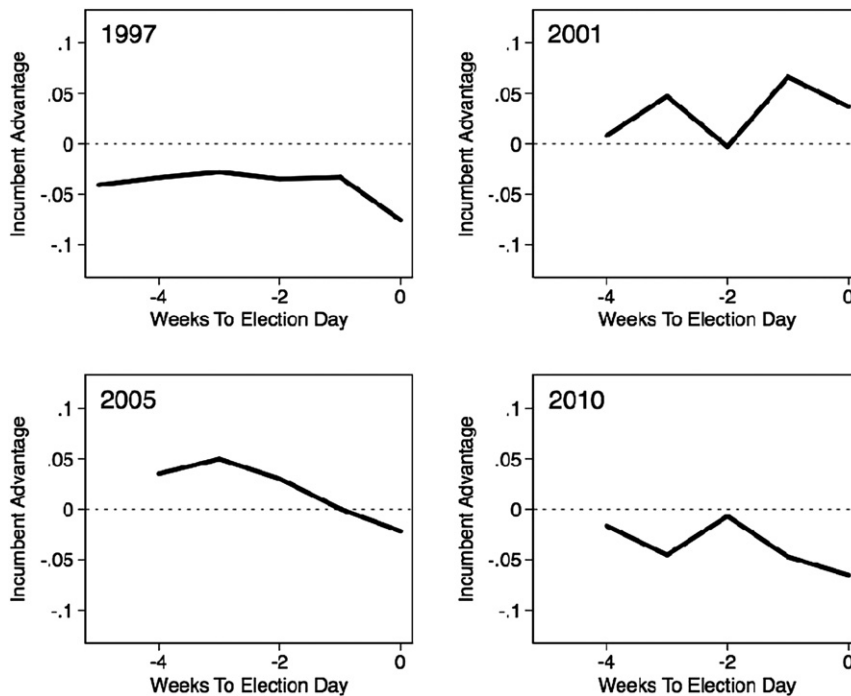


Fig. 2. Incumbent CI Advantage, UK Election Campaigns.

a simple average of Incumbent CI Advantage over the campaign, but we have tested a number of other possibilities, including: Incumbent CI, Incumbent C_f , Incumbent C_u , combinations of Incumbent and Challenger C_f or C_u , and campaign-period differences, where we captured the shift in Incumbent CI and other measures from the first two weeks of the campaign to the final two weeks of the campaign. A summary of results from these alternative tests is presented in [Appendix Table A1](#). As can be seen, in no case do the pooled results reveal any relationship between errors from the historical models and campaign-period media tone.⁶

In the meantime, we note that an even more direct test of the connection between media tone and prediction error would be to add our measure of media tone to the historical models of [Table 1](#). Doing so would test the possibility that media matter above and beyond the impact of the other independent variables. It is not possible, however: because media data are available only for recent elections, we lose most of the cases in the historical models. In the UK, there are not enough degrees of freedom to estimate the model at all; in Canada and the US, the models can be estimated but yield

no significant results.⁷ (And note that a pooled model is not possible since many of the independent variables differ across countries.) We are left, then, with analyses of four to seven elections in each country. Even these limited cases reveal an interesting possibility, however, explored below.

5. The possibility of indirect campaign effects

Based on the evidence presented thus far, it seems that campaigns cannot account for the error in historical models, at least in the small sample of elections that can be examined in this study. Are there campaign effects at all? One possibility is that campaigns influence the election outcome indirectly—that rather than add (or subtract) from the impact of long-term political and macroeconomic variables, campaigns capture or engage the impact of those variables. This is what some electoral forecasters have suggested; and it would explain why campaigns do not account for errors in long-term predictions. (Indeed, this would explain why we see signs of a positive relationship between error and media tone in the preceding results.)

If this is the case, then we should expect the content of a campaign to be determined in good part by pre-existing

⁶ These pooled models require that the relationship between error and Incumbent CI Advantage is similar across countries—not just in direction, but in magnitude as well. This may be an unrealistic assumption. Running country-specific models reveals a possibly interesting difference across countries; namely, a much stronger positive connection between error and media tone in the US than in the other two countries. But to be clear, no coefficient here is significant, sample sizes are tiny, and the relationship is in the opposite direction from what we would expect if campaigns account for the error in the historical models.

⁷ Not only is the regression coefficient for the Incumbent CI Advantage variable not significant, it is wrongly signed (negative)—a result to which we briefly come back in the next section (see footnote 11). Here are the forecasting equation estimates for the two countries including the media tone variable (none of the variables are statistically significant at the .05 level): US: $\text{vote} = 52.61 + .18 (\text{approval}) + .40 (\text{growth}) - 6.39 (\text{time}) - 69.97 (\text{media tone})$ CA: $\text{vote} = 55.74 + .52 (\text{approval}) - 3.31 (\text{unemployment}) - 4.44 (\text{time}) - 50.95 (\text{media tone})$.

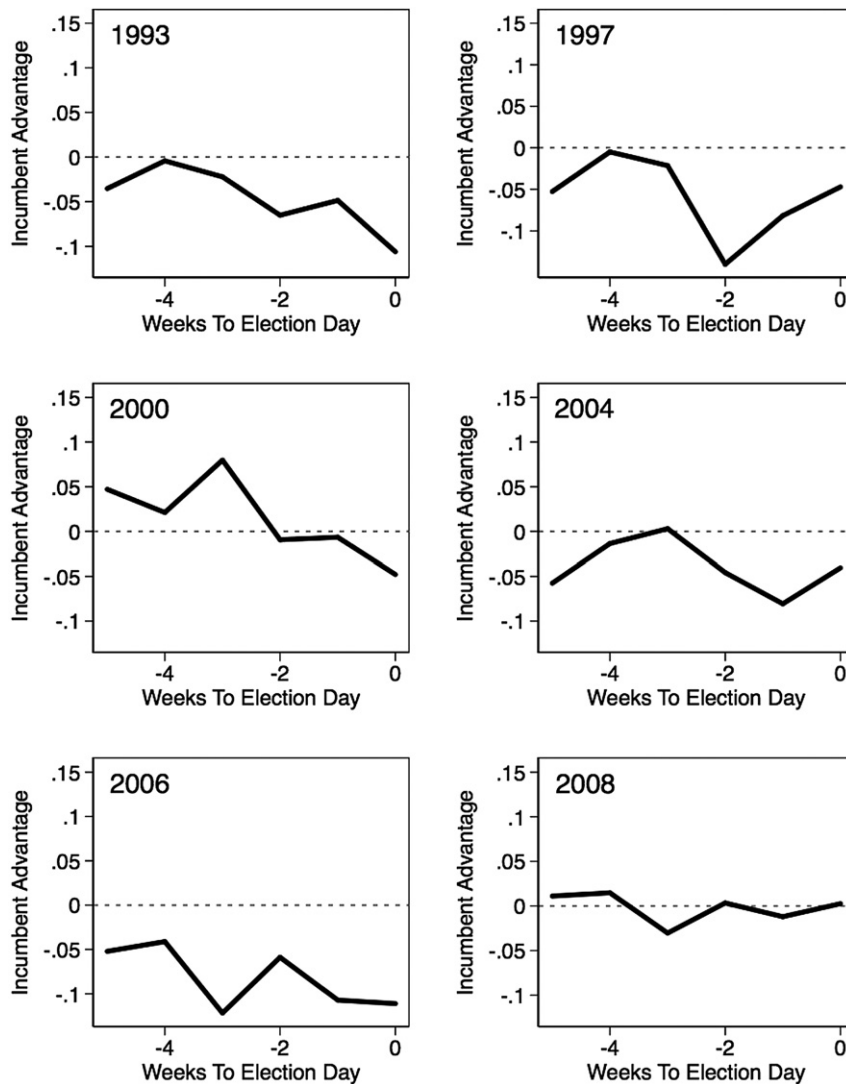


Fig. 3. Incumbent CI Advantage, CA Election Campaigns.

national political and economic conditions—those same conditions that are already captured by the historical prediction models. One way to explore this possibility is to examine directly whether our measure of campaign-period media tone actually reflects macroeconomic conditions, government approval, and/or government longevity, as measured a few months in advance of campaigns. We do so via the simple bivariate correlations presented in Table 4.⁸

The first column shows correlations between Incumbent CI Advantage and each of the independent variables in the prediction models for the US, UK and Canada. Results are clearly tentative from a statistical point of view—with only one exception, no correlations are statistically significant at the .05 level (though there are several significant at

the .10 level). We are dealing with very small sample sizes, however. And taking the correlations at face value, results do point toward the ‘indirect influence’ argument. Eight out of nine correlation coefficients display the expected sign. For the US and Canada cases, all relationships are in the expected direction; for the UK, only one relationship is wrongly signed (inflation).

We can explore the extent to which the existing economic situation is captured in media content a little more directly, however. To do so, we separate media content into articles with some discussion of the economy, versus articles with little to no discussion of the economy. This is done with a simple keyword search, where we count the number of economic keywords in each article.⁹ Articles with a count of 2 or more are categorized as (at least partly)

⁸ Again, performing multivariate tests is difficult here given the small number of observations that we have in each country and given the measurement differences across countries.

⁹ Keywords are: economy, economic, *employment, inflation, stock market, stock exchange, fiscal, monetary, budget, debt, and interest rate.




Fig. 4. Incumbent CI Advantage and Vote Intentions, Daily, US Election Campaigns.

economic; all other articles are categorized as non-economic. (The proportion of economic articles, by election, is shown in Appendix Table A2.) We then recalculate the Incumbent CI Advantage, for each election, for (a) economic and (b) non-economic articles.

The resulting bivariate correlations are shown in the second and third columns of Table 4. The link between the tone of economic coverage and the long-term variables in the US is as we might expect—the link between tone and both approval and the economy strengthens (while the link between tone and terms in office declines somewhat). Correlations weaken considerably when we look at non-economic news only. We take these results as possible

Table 3

Incumbent CI Advantage and Vote Intentions, Daily, US Election Campaigns.

	2004		2008	
	DV: Opinion _t	DV: Media _t	DV: Opinion _t	DV: Media _t
Opinion _{t-1}	.963*** (.039)	-.002 (.002)	.932*** (.032)	.002 (.002)
Media _{t-1}	2.395** (.833)	.787*** (.044)	2.819* (1.128)	.844*** (.074)
Constant	3.253 (2.043)	.088 (.108)	3.313* (1.576)	-.118 (.103)
N	61	61	62	62
Rsqr	.907	.839	.958	.785

*p < .05; **p < .01; ***p < .001. Cells contain regression coefficients from a vector autoregression with standard errors in parentheses.

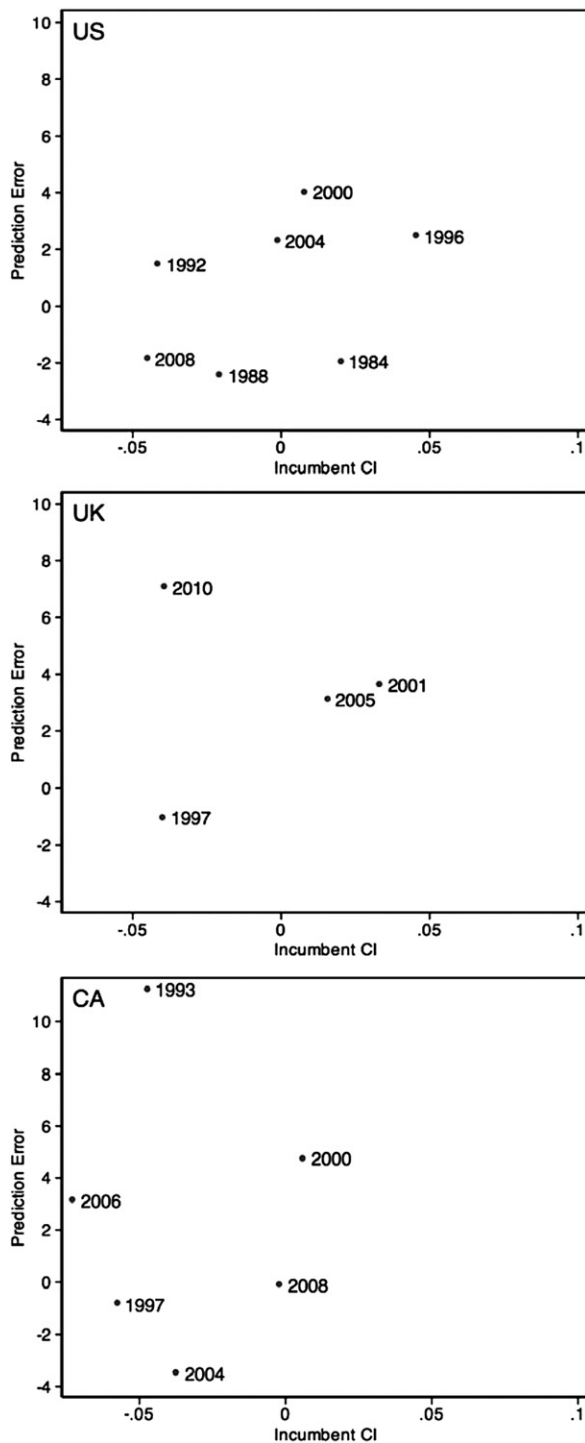


Fig. 5. Incumbent CI Advantage & Prediction Error.

evidence that the tone of the campaign is linked to the components already included in the long-term models.

Evidence in Canada points in the same direction. Recall that the economy in Canada is captured using the unemployment rate. We thus expect the coefficient to be negative, as it is, and the coefficient strengthens when we focus on economic news alone. The link between tone and approval is also stronger in

Table 4
Incumbent CI Advantage & the fundamentals, bivariate correlations.

	Bivariate correlations with Incumbent CI Advantage		
	All news	Economic news	Non-economic news
US (N = 7)			
Approval	.735 (.059)	.882 (.009)	.421 (.347)
Economy	.734 (.060)	.818 (.025)	.473 (.284)
Time	-.747 (.053)	-.664 (.102)	-.629 (.130)
UK (N = 4)			
Approval	.983 (.017)	.869 (.131)	.988 (.012)
Economy	.681 (.319)	.712 (.288)	.591 (.409)
Time	-.940 (.059)	-.843 (.157)	-.929 (.071)
CA (N = 6)			
Approval	.307 (.554)	.502 (.311)	.171 (.746)
Economy	-.426 (.400)	-.487 (.327)	-.228 (.664)
Time	-.427 (.398)	-.361 (.483)	-.427 (.399)

Cells contain Pearson's R coefficients with *p*-value in parentheses.

Canada when we look at economic news only. In the UK, results are more mixed. Inflation is not clearly linked to the tone of the campaign. Approval clearly is, however; though note that the link is strongest here for non-economic articles.

We cannot account easily for these differences across countries.¹⁰ Results nevertheless point in the direction of a positive link between the content of a campaign and the long-term variables. In short, the mean Incumbent CI Advantage is more positive (negative) when the long-term variables are more (less) favorable to the incumbent candidate or party—which is exactly what one would expect if campaigns and their media coverage really are about making voters aware of the incumbent's political and economic performance of the recent past (Gelman and King, 1993; Holbrook, 1996; Erikson and Wlezien, 2012). Viewed in this light, our finding that campaigns have no significant relationship to the historical models' prediction errors is not too surprising.¹¹ The implication is that these prediction errors may be just that—errors unaccounted for by the campaigns because the latter's influence is (or can be assumed to be) embedded within these models. For the time being, Table 4 points in this direction; a greater sample of observations might of course go a long way toward clarifying this possible indirect effect of campaigns on election outcomes.

6. Summary and conclusion

Can historical and campaign-period prediction models of election outcomes be combined? While this question remains difficult to answer in the absolute, the evidence provided here seems to suggest that they cannot. With data from 17 national elections held in the US, the UK, and Canada, we have shown that the prediction errors of long-

¹⁰ Though note that macroeconomic variables are different across countries (based on differences in the specification of existing prediction models), and it may be that different economic factors are reflected to greater or lesser degrees in media content.

¹¹ In fact, it is akin to a multicollinearity problem. Recall that in footnote 7 the Incumbent CI Advantage regression coefficient is negatively signed when included together with the other predictors of the vote—in spite of the fact that the same media tone variable correlates positively with the vote in a bivariate relationship (Pearson's *R* of .75 for the US, .96 for the UK, and .46 for Canada).

term politico-economic models are not related to campaign content as measured by media tone. In other words, it appears that deviations between predicted and actual incumbent vote share in these elections cannot be accounted for by what is going on in the campaigns.

Why is there no relationship between the prediction errors from the historical models and measures of media content during the election campaigns? One possible explanation is that the errors from the predictions are simply not the consequence of campaign events. This might be true in spite of the fact that campaign-period variations in vote intentions clearly move, and clearly do so alongside (if not also as a consequence of) media. So campaigns may barely matter, given the relative success of macroeconomic predictions. But they may also matter, given that there are sometimes significant shifts in the campaign; and given that what is going on in the media is reflected in vote intentions. That said, we think that the small sample sizes with which we had to work prevent us from outright accepting the null hypothesis. It might be more appropriate to simply say that we fail to reject the null in this study.

There are of course other possibilities. Our measures of media content may not adequately capture the campaign, for instance. The newspapers selected may not be a sufficiently representative sample (particularly in the US), or the way in which overall campaign tone is calculated may not be sufficiently accurate. As discussed above, however, we have tried a wide variety of summary measures of campaign tone; and the weekly and daily trends in media tone suggest that we are indeed capturing a good part of what is going on in campaigns.

Another possibility might be that what matters are the candidates' messages about the economy, as measured via their campaign advertisements and speeches—campaigns may matter in their potential to “clarify” the impact of long-term factors (Vavreck, 2009). Put differently, it is not media content on the campaign as a whole that accounts for prediction error, but rather particular components of media content. This may well be true, though our emphasis here has been to examine the possibility that the campaign as a whole accounts for prediction error—given that campaign-period prediction models show (in past work as in our own analyses) that campaign content relates closely to vote intentions.¹² As it turns out, while campaign-period prediction models clearly work during the campaign, they do not add much to historical prediction models.

Still another possible reason for why historical and campaign-period prediction models cannot be combined may be surprisingly simple: it might just be that they are tracking very different things. As Gelman and King (1993) have argued, responses to pollsters and vote decisions are two different phenomena—the former mostly are gut-feeling answers that merely reflect day-to-day campaign events, whereas the latter are the final product of preferences enlightened by the campaign process as a whole. In other words, campaign-period prediction models may be capturing the day-to-day

dynamics between campaign events and vote intentions, while historical models may be predicting what the outcome should be once the campaign has finished performing its “enlightening” function and the preferences of the public have come into sharper focus (see esp. Erikson and Wlezien, 2012).

The last section of our analysis has presented some evidence, albeit weak from a statistical standpoint, that campaigns on average are partially reflecting the state of national political and economic conditions. Campaigns may thus influence election outcomes, but the impact may be in line with, rather than orthogonal to, the various components of historical politico-economic predictions. This is a significant finding. It helps account for the success of two apparently fundamentally opposed projects—historical versus campaign-period prediction models. Put differently, it helps account for the success of historical forecast models, even as we know that campaigns can matter to election outcomes. At the same time, it is clear that there is much more to understand about the relationship between the long-term predictions and election campaigns.

Acknowledgements

A previous version of this article was presented at the annual meeting of the American Political Science Association held in Seattle in September 2011. We are grateful to participants for comments, and especially to Christopher Wlezien and Michael Lewis-Beck. We are additionally grateful to Wlezien and Robert Erikson for providing campaign-period trends in US vote intentions, to Alan Abramowitz for sharing his US forecast data, and to Cameron Stark for his research assistance. This research was supported in part by the Québec Fonds de recherche sur la société et la culture.

Appendix A

Results in Fig. 5 are based on pooled OLS models in which prediction errors are regressed on campaign tone, captured using Incumbent CI Advantage, alongside dummy variables for the UK and Canada. There are other possibilities where campaign tone measures are concerned, however. To confirm that our results are not a product of mis-specifying the campaign tone variable, then, Appendix Table A1 shows results from regression models using a variety of specifications for the campaign tone variables; specifically, (a) campaign-period means, (b) campaign-period sums (allowing for a slightly different cumulative effect of media over the campaign), and (c) campaign-period differences, capturing the change in each media variable from the first two weeks of the campaign to the final two weeks of the campaign.

Note that we show only pooled results, not separate by-country results, because the latter do not show any additional results. Note also that results used for Fig. 5 are found in the ninth row of the table. Overall, no other specification leads to improved results; that is, in no case does there appear to be a systematic relationship between prediction error and campaign tone.

As noted in the text, Appendix Table A2 shows the proportion of articles in each election dealing with the economy.

¹² As Vavreck (2009: 49) points out, testing as we do for whether the tone of campaign coverage affects election outcomes is “another interesting, but essentially different project” than her own.

Table A1
Regressing error on media tone, with various specifications for media.

Media specification	Coefficients				R ²
	Media	UK (dummy)	CA (dummy)	Constant	
<i>Means</i>					
Incumbent C _f	−15.483 (45.936)	2.520 (2.522)	1.865 (2.225)	2.412 (5.682)	.098
Incumbent C _u	−37.514 (94.238)	2.052 (2.872)	.118 (4.933)	4.359 (9.647)	.101
Incumbent CI	−4.935 (39.872)	2.657 (2.542)	2.099 (2.900)	.655 (1.681)	.091
Challenger C _f	−53.200 (47.298)	2.038 (2.458)	2.926 (2.331)	6.432 (5.414)	.170
Challenger C _u	−41.334 (76.943)	1.690 (3.026)	−.471 (4.887)	4.176 (6.886)	.109
Challenger CI	−20.862 (37.043)	2.853 (2.524)	3.466 (3.592)	1.044 (1.724)	.111
Incumbent C _f ADV	17.822 (34.834)	2.527 (2.499)	2.231 (2.322)	.406 (1.535)	.108
Incumbent C _u ADV	44.967 (129.494)	2.281 (2.683)	1.425 (2.568)	−.054 (2.338)	.098
Incumbent CI ADV	12.087 (31.910)	2.644 (2.506)	2.234 (2.422)	.624 (1.518)	.100
<i>Sums</i>					
Incumbent C _f	.014 (.012)	2.755 (2.402)	−.631 (3.043)	−4.148 (4.344)	.174
Incumbent C _u	.009 (.010)	3.091 (2.492)	2.416 (2.243)	−2.353 (3.467)	.146
Incumbent CI	.000 (.013)	2.599 (2.573)	1.789 (3.875)	.554 (1.587)	.090
Challenger C _f	.004 (.008)	2.674 (2.496)	.847 (2.983)	−.672 (2.848)	.108
Challenger C _u	.006 (.011)	3.117 (2.641)	2.678 (2.627)	−1.162 (3.398)	.112
Challenger CI	.001 (.009)	2.532 (2.572)	1.369 (4.015)	.500 (1.578)	.091
Incumbent C _f ADV	.002 (.009)	2.603 (2.515)	2.027 (2.365)	.522 (1.531)	.092
Incumbent C _u ADV	.041 (.031)	1.469 (2.526)	−.915 (2.999)	−1.024 (1.879)	.195
Incumbent CI ADV	−.001 (.008)	2.590 (2.523)	1.715 (2.547)	.549 (1.523)	.091
<i>Differences</i>					
Incumbent C _f	−61.969 (44.236)	2.600 (2.347)	1.363 (2.115)	.579 (1.415)	.209
Incumbent C _u	105.342 (80.799)	2.891 (2.378)	1.663 (2.108)	.287 (1.444)	.195
Incumbent CI	−51.103 (31.971)	2.737 (2.303)	1.351 (2.069)	.441 (1.390)	.239
Challenger C _f	−11.162 (40.583)	2.217 (2.895)	1.609 (2.422)	.962 (2.090)	.095
Challenger C _u	161.315 (115.486)	1.444 (2.493)	.088 (2.444)	1.735 (1.645)	.208
Challenger CI	−21.159 (34.522)	1.708 (2.888)	1.142 (2.504)	1.470 (2.102)	.115
Incumbent C _f ADV	−26.315 (34.833)	3.542 (2.754)	2.270 (2.251)	−.363 (1.929)	.128
Incumbent C _u ADV	31.033 (90.523)	2.920 (2.662)	2.153 (2.372)	.258 (1.757)	.098
Incumbent CI ADV	−23.889 (30.495)	3.692 (2.820)	2.451 (2.306)	−.514 (2.025)	.131

Table A2
Proportion of economic articles, by election.

Country	Year	% Economic articles
US	1984	30.9%
	1988	26.6%
	1992	37.9%
	1996	26.3%
	2000	24.9%
	2004	17.9%
UK	2008	33.5%
	1997	20.2%
	2001	18.8%
	2005	17.2%
CA	2010	19.0%
	1993	34.5%
	1997	27.9%
	2000	20.3%
	2004	19.9%
	2006	15.0%
	2008	29.7%

References

- Abramowitz, A.I., 2008. Forecasting the 2008 presidential election with the time-for-change model. *PS: Political Science and Politics* 41, 691–695.
- Arceneaux, K., 2006. Do campaigns help voters learn? A cross-national analysis. *British Journal of Political Science* 36, 159–173.
- Bélanger, É., Godbout, J.-F., 2010. Forecasting Canadian federal elections. *PS: Political Science and Politics* 43, 691–699.
- Bélanger, É., Lewis-Beck, M.S., Nadeau, R., 2005. A political economy forecast for the 2005 British general election. *British Journal of Politics and International Relations* 7, 191–198.
- Brady, H.E., Johnston, R. (Eds.), 2006. *Capturing Campaign Effects*. University of Michigan Press, Ann Arbor, MI.

- Brians, C.L., Wattenberg, M.P., 1996. Campaign issue knowledge and salience: comparing reception from TV commercials, TV news, and newspapers. *American Journal of Political Science* 40, 172–193.
- Campbell, J.E., 2008a. Editor's introduction: forecasting the 2008 national elections. *PS: Political Science and Politics* 41, 679–682.
- Campbell, J.E., 2008b. *The American Campaign*, second ed. Texas A&M University Press, College Station, TX.
- Campbell, J.E., Garand, J.C. (Eds.), 2000. *Before the Vote: Forecasting American National Elections*. Sage, Thousand Oaks, CA.
- Cutler, F., Fournier, P., Soroka, S., Lyle, G., 2004. "Need for Cognition, Need to Evaluate, and Change in Vote Choice." Paper presented at the Annual Meeting of the American Political Science Association, Chicago.
- de Vreese, C.H., Semetko, H.A., 2004. News matters: influences on the vote in the Danish 2000 Euro referendum campaign. *European Journal of Political Research* 43, 699–722.
- Dobrzynska, A., Blais, A., Nadeau, R., 2003. Do the media have a direct impact on the vote? The case of the 1997 Canadian election. *International Journal of Public Opinion Research* 15, 27–43.
- Druckman, J.N., 2004. Priming the vote: campaign effects in a U.S. Senate election. *Political Psychology* 25, 577–594.
- Erikson, R.S., Wlezien, C., 2012. *The Timeline of Presidential Elections: How Campaigns Do (And Do Not) Matter*. University of Chicago Press, Chicago.
- Fan, D.P., 1988. *Predictions of Public Opinion from the Mass Media: Computer Content Analysis and Mathematical Modeling*. Greenwood, New York.
- Gelman, A., King, G., 1993. Why are American presidential election campaign polls so variable when votes are so predictable? *British Journal of Political Science* 23, 409–451.
- Gibson, R., Lewis-Beck, M.S., 2011. Methodologies of election forecasting: calling the 2010 UK 'hung parliament'. *Electoral Studies* 30, 247–249.
- Greene, J.P., September 21, 1993. Forecasting follies. *The American Prospect*.
- Holbrook, T.M., 1996. *Do Campaigns Matter?* Sage, Thousand Oaks, CA.
- Hopmann, D.N., Vliegenthart, R., de Vreese, C., Albaek, E., 2010. Effects of election news coverage: how visibility and tone influence party choice. *Political Communication* 27, 389–405.
- Janis, I.L., Fadner, R.H., 1943. A coefficient of imbalance for content analysis. *Psychometrika* 8, 105–119.

- Johnston, R., Blais, A., Brady, H.E., Crête, J., 1992. *Letting the People Decide: Dynamics of a Canadian Election*. McGill-Queen's University Press, Montreal and Kingston.
- Johnston, R., Hagen, M.G., Jamieson, K.H., 2004. *The 2000 Presidential Election and the Foundations of Party Politics*. Cambridge University Press, Cambridge.
- Just, M.R., Crigler, A.N., Alger, D.E., Cook, T.E., Kern, M., West, D.M., 1996. *Crosstalk: Citizens, Candidates, and the Media in a Presidential Campaign*. University of Chicago Press, Chicago.
- Lewis-Beck, M.S., 2005. Election forecasting: principles and practice. *British Journal of Politics and International Relations* 7, 145–164.
- Lewis-Beck, M.S., Tien, C., 2002. Presidential election forecasting: the Bush–Gore draw. *Research in Political Sociology* 10, 173–187.
- Lewis-Beck, M.S., Tien, C., 2008. Forecasting presidential elections: when to change the model. *International Journal of Forecasting* 24, 227–236.
- Lewis-Beck, M.S., Tien, C., 2011. Election forecasting. In: Clements, M.P., Hendry, D.F. (Eds.), *The Oxford Handbook of Economic Forecasting*. Oxford University Press, Oxford, UK, pp. 655–671.
- Mendelsohn, M., 1996. The media and interpersonal communications: the priming of issues, leaders and party identification. *Journal of Politics* 58, 112–125.
- Mendelsohn, M., Nadeau, R., 1999. The rise and fall of candidates in Canadian election campaigns. *International Journal of Press/Politics* 4, 63–76.
- Nadeau, R., Lewis-Beck, M.S., 2012. Does a presidential candidate's campaign affect the election outcome? *Foresight* 24, 15–18.
- Nadeau, R., Neville, N., Gidengil, E., Blais, A., 2008. Election campaigns as information campaigns: who learns what and does it matter? *Political Communication* 25, 229–248.
- Nadeau, R., Pétry, F., Bélanger, É., 2010. Issue-based strategies in election campaigns: the case of health care in the 2000 Canadian federal election. *Political Communication* 27, 367–388.
- Norris, P., Curtice, J., Sanders, D., Scammell, M., Semetko, H.A., 1999. *On Message: Communicating the Campaign*. Sage, London.
- Sheafer, T., Weimann, G., 2005. Agenda building, agenda setting, priming, individual voting intentions, and the aggregate results: an analysis of four Israeli elections. *Journal of Communication* 55, 347–365.
- Silver, N., March 26, 2012. Models Based on 'Fundamentals' Have Failed at Predicting Presidential Elections. *Five Thirty Eight Blog*. The New York Times.
- Soroka, S., Bodet, M.A., Young, L., Andrew, B., 2009. Campaign news and vote intentions. *Journal of Elections, Public Opinion and Parties* 19, 359–376.
- Tetlock, P.E., 2005. *Expert Political Judgment: How Good Is It? How Can We Know It?* Princeton University Press, Princeton, NJ.
- van der Eijk, C., 2005. Election forecasting: a sceptical view. *British Journal of Politics and International Relations* 7, 210–214.
- Vavreck, L., 2009. *The Message Matters: The Economy and Presidential Campaigns*. Princeton University Press, Princeton, NJ.
- Walgrave, S., de Swert, K., 2004. The making of the (issues of the) Vlaams Blok. *Political Communication* 21, 479–500.
- Weaver, D.H., 1996. What voters learn from media. *Annals of the American Academy of Political and Social Science* 546, 34–47.
- Wlezien, C., 2010. Election campaigns. In: LeDuc, L., Niemi, R.G., Norris, P. (Eds.), *Comparing Democracies 3: Elections and Voting in the 21st Century*. Sage, Thousand Oaks, CA, pp. 98–117.
- Wlezien, C., Erikson, R.S., 2002. The timeline of presidential election campaigns. *Journal of Politics* 64, 969–993.
- Wlezien, C., Erikson, R.S., 2005. Post-election reflections and our pre-election predictions. *PS: Political Science and Politics* 38, 25–26.
- Young, L., Soroka, S., 2012. Affective news: the automated coding of sentiment in political texts. *Political Communication* 29, 205–231.