Noisy and Bursty Opinion Streams: Methods for Analyzing Dynamics of Aggregate Opinion Change

Abe Gong
Ford School of Public Policy,
University of Michigan,
Ann Arbor, Michigan
agong@umich.edu

December 18, 2007

Abstract

This paper introduces statistical methods for analyzing the dynamics of aggregate opinion change. Under the reasonable assumption that information flow is driving these changes, these statistical results can lead to inferences about the information processes behind public opinion. Using data from the 2004 National Annenberg Election Study, I construct a set of twenty opinion streams and compare them with respect to estimated real opinion change and burstiness.

1 Introduction

Virtually every theory of public opinion formation is built on the premise that an opinion is the product of an individual's internal "predispositions" and external "information."

Much has been written about the values, cues, models, heuristics, group memberships, emotions, incentives, and genetics that determine an individual's response to information. However, the nature of the "information flows" surrounding and connecting individuals has been treated much less thoroughly.

This imbalance has been driven largely by methodological constraints: individuals can be caught and interrogated; "flows" are abstract and almost invisible. Nevertheless, the information context that makes the public the public is an important part of any theory of public opinion. This paper demonstrates how empirical techniques from time series analysis and natural language processing can be used to approach the information half of the opinion equation.

2 Literature Review

In this section I outline theories of information flow from the public opinion literature. Although I address a variety of authors, I structure my argument around Zaller's "Receive, Accept, Sample" (RAS) model as outlined in the opening chapters of *The Nature and Origin of Mass Opinion*[11]. I choose Zaller's model because it is representative of many of the dominant traditions in public opinion and has gained broad currency in the last decade. Much of this paper can be seen as a response to and criticism of Zaller's claims about "elite discourse." Also, Zaller's model is unusually clear, comprehensive, and broadly applicable, and therefore makes for a good starting point for discussion.

According to the RAS model, all politically relevant information originates with elites and percolates down through a semi-aware mass public. Although members of the public may react to political information in different ways, they are essentially passive recipients of elite discourse. If elites disagree among themselves, the mass may polarize, but they will never do so on their own. Zaller further postulates that individuals in the mass all drift in the same stream of information; the only difference is that some individuals tend to be more aware of it than others.

Considering the simplicity of the model and the number of very strong assumptions it invokes, Zaller's model gives surprisingly good empirical results. In cases as diverse as medical consensus on the lack of racial correlation to intelligence, various phases of the Vietnam war, and reelection campaigns of Congressional incumbent, Zaller shows that the RAS model nicely predicts empirical results. However, the model fails to anticipate several interesting empirical regularities. Here I highlight four important holes in Zaller's theory which I will address throughout this paper.

First, elites are treated as the sole source of information. This is in contrast to Stimson's "thermostatic democracy" hypothesis, which argues that individuals observe and react directly to the outcomes of policy.[10] Stimson's evidence that approval of government is largely driven by economic performance and his somewhat weaker evidence of a liberal-conservative cycle in national opinion both suggest that elites do not control all relevant political information. To be fair to Zaller, the RAS model doesn't exclude the possibility of such bottom-up effects: in theory elites might gather information from the public and then "discourse" about it. But nothing in the model explicitly predicts this behavior, and including it in this fashion seems much less parsimonious (and normatively desirable) than to simply assume that voters are capable of making basic political inferences directly from personal experience.

Second, Zaller's model ignores the channels by which information is disseminated throughout society. Granted, this simplification works for a first approximation. But it assumes away a host of interesting questions about campaigns, opinion leadership, social connections and cleavages, and "deliberative democracy" in general. Putnam,[8] Gelman[1], Huckfeldt, and Sprague [3] have written at length on this subject, although their work falls short of a comprehensive theory.

Third, Zaller's model makes no prediction as to the frequency of important new information. Years of journalistic lore suggests that the news is made up of scores of little stories ("College Loans by States Face Scrutiny") and the occasional big story ("Nixon Resigns"). In the long term, these stories might average out to a more or less constant stream of information.[1] In the short term, a handful of prominent stories might represent a disproportionate part of the information reaching the public[2]. At the short end of the time scale, Stimson's analysis of post-convention bumps in Presidential campaigns provides one case of a discrete information event that can have a substantial impact on public opinion.

Fourth, Zaller's model makes no predictions as to the speed at which information is disseminated through society. This simplification ignores an extensive literature on opinion leadership. Furthermore, it probably limits the validity of Zaller's conclusions to societies where political communication is dominated by mass media.

In conclusion, Zaller uses "elite discourse" as a caricature of information flow. Although this concept is useful as a first approximation, it oversimplifies the sources, channels, frequency, and propagation speed of information through society. My goal in this paper is to introduce analytical methods for thinking about these properties of information flow.

[1] "Might" is the operative word here. Stories might be normally distributed according to size, or they might fall under a scale-free distribution. Sears[9] bases much of his theory of symbolic politics on the idea that rare, massive, shared events shape the consciousness of whole generations.

3 Research Question

In this paper I introduce methods for investigating information flow indirectly by making inferences from streams of aggregate opinion change. As examples, I analyze several different time series of public opinion data. Each series is a daily stream of polling data on a dichotomous variable. Unlike previous studies, which focus almost exclusively on the direction of opinion change, this paper focuses on the timing and magnitude of opinion change. Careful statistical work enables me to answer two questions about each stream:

- How much real opinion change occurs over the span of the stream?
- Does change occur in bursts?

For each of these questions, I introduce metrics and visualizations for thinking about the dynamics of aggregate opinion change. Under the reasonable assumption that information flow is driving these changes, these statistical results can lead to inferences about the information processes behind public opinion. Comparing results from different streams enables me to draw some early conclusions about the nature of societal information flows.

4 Methodology

In this section I describe statistical methods for looking at change in opinion streams. This work can be seen as fitting into a large emerging literature on aggregate public opinion.[1][2] The section is organized around a series of analytical concepts that I have adapted for investigating opinion change.

4.1 Opinion Streams

All of this analysis begins with times series of aggregated responses to a dichotomous question, plus a corresponding series of poll sample sizes. This type of data structure is a common product of rolling cross-sectional surveys. For purposes of this paper, I take several variables from the National Annenberg Election Study (NAES). These data streams span the year leading up to the 2004 general elections. Surveys of national opinion were conducted daily throughout this time, with sample sizes in the dozens early in the survey and growing into the triple digits as election day approached.

When working with data, I will refer to the first series as *proportion vector*, $\vec{P}$, the second series as the *sample size vector*, $\vec{N}$, and both series together as an *opinion stream*. Both $\vec{P}$ and $\vec{N}$ are of length m, where m is the number of time steps in the stream. Elements of $\vec{P}$ and $\vec{N}$ will be referenced as p_i and n_i.

Figure 1: A line graph of $\vec{P}$ is the familiar plot used to show trends in election horse races and approval ratings. Here a smoothed tracking line is shown over raw $\vec{P}$.

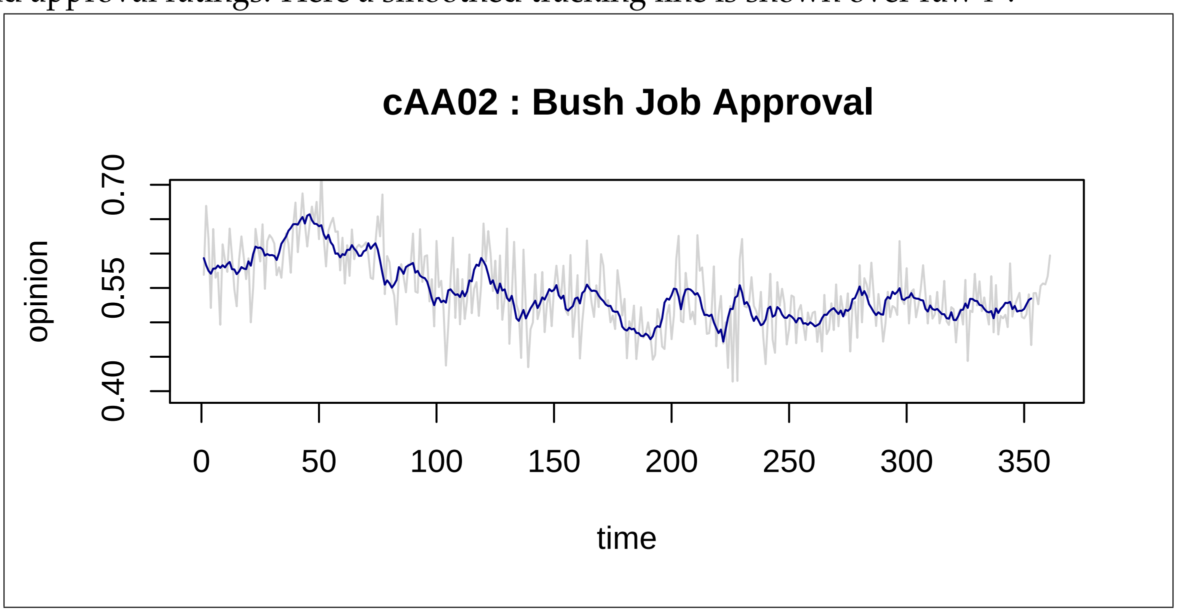

Note that opinion streams are not panels. Respondents in each time step are assumed to be drawn independently from the population. Note also that opinion streams are the product of two levels of aggregation: first, survey responses on a given day, i, are aggregated up to generate p_i and n_i; second, values of p_i and n_i are aggregated across some time span, $i \epsilon [a, b]$, to form opinion streams.

I define *observed opinion change*, q_i in a given time step i as the absolute value of the difference between p_i and p_{i+1}. The vector $\vec{Q}$ of opinion changes has length $m - 1$. The *total opinion change* of a stream is the sum of q_i for all i — opinion change aggregated across the whole stream.

Figure 2: Opinion change over time

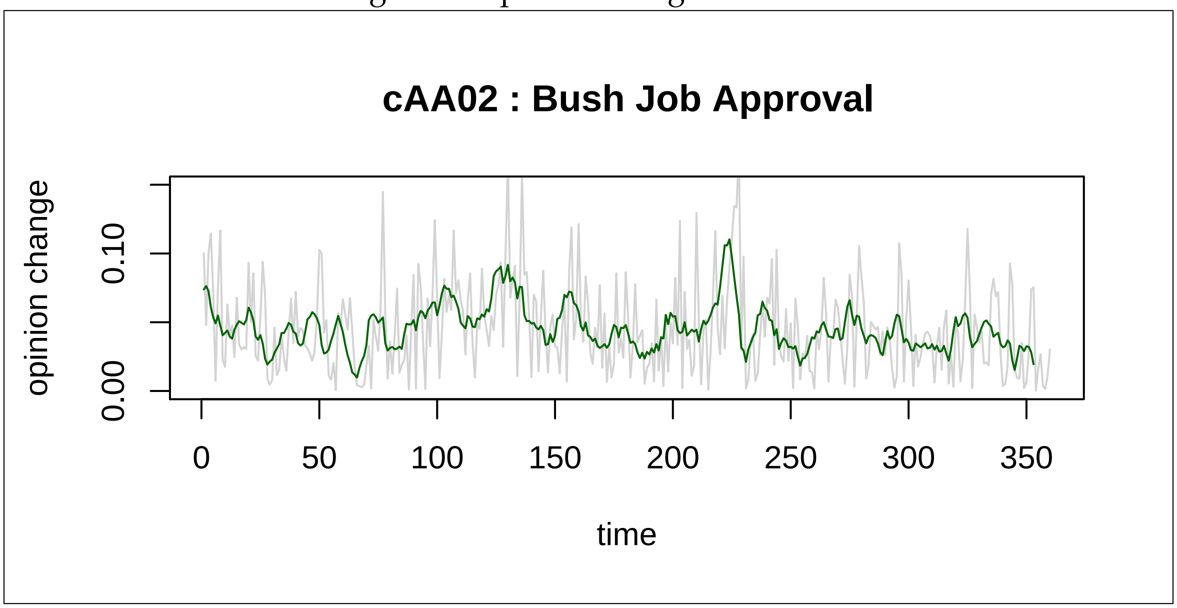

4.2 Noise and Real Opinion Change

Much of the opinion change in a given stream is the result of statistical artifacts, which I lump together as *noise*. Conceptually, noise is the difference between the value of opinion change we would observe in a "rolling census" of public opinion and the value we actually observe. In the NAES data, the major source of noise is sampling error.[2] Because each daily poll is measured with some margin of error, p_i would be expected to differ at least slightly from p_{i+1} for any given i. As a rule, $q_i \neq 0$.

Notice that noise in a stream of opinion change behaves differently than error in a typical time series: because opinion change is the absolute difference in observed opinions, q_i is nonnegative for all i. Consequently, noise due to sampling variation is always positive as well. It does not tend to cancel out over time and only approaches toward zero with astronomical sample sizes. Consistent with typical margins of error, it is impossible to know the amount of noise in a given time step. However, it is possible to estimate the total noise in a stream using Monte Carlo methods.

In order to estimate the total noise in an opinion stream, I sampled 200 *proxy series* as follows: using $\vec{P}$ and $\vec{N}$ from a real opinion stream, I calculated a *standard error series*, $\vec{R}$, with $r_i = \sqrt{p_i(1-p_i)/n_i}$. This is the familiar expression for standard error of a binomial distribution. Next, I created proxy series, $\vec{S}$ with m time steps, each s_i being drawn from a normal distribution centered about zero with a standard deviation of r_i. The changes in this vector approximate the noise in the real opinion stream. By deriving a change vector $\vec{Q^*}$ from $\vec{S}$ and summing total change over $\vec{Q^*}$, I effectively assembled a random draw of possible amounts of total noise in the stream. By repeating this process for 200 proxy series, I created a *simulated noise distribution* for each of the real opinion streams. Estimated noise is the mean of the distribution.

By construction, *real opinion change* is the difference between observed opinion change

[2]In theory, an opinion stream could also be created by "stitching together" a large set of independent polls asking standardized questions and taken across time. This would create at least three more sources of positive and negative noise: apparent change due to poll-specific fixed effects, apparent change due to rounding, apparent change from rolling averages. I leave these methodological challenges for later work.

and noise. Given an estimate of the noise in a stream, we can easily subtract the observed opinion change to arrive at an estimate of *total real opinion change*. Furthermore, the simulated noise distribution can also be used to construct a confidence interval of sorts.

4.3 Burstiness

Burstiness is a concept introduced into the literature of Natural Language Processing by Kleinberg[6]. His original intent was to extract patterns in activity from streams of text, such as email or formal speeches. The key intuition behind burstiness is that some streams of information come in a steady flow, while others are made up of long periods of low information flow punctuated by bursts of high activity.

Here I adopt Kleinberg's intuition, but not his methods.[3] Following the notion that bursty streams are those where change is driven by a few large event, I operationalize burstiness as the Gini coefficient of the vector $\vec{Q}$. This measure is a well-known index of inequality within a population. It is most often used to calculate income inequality. As typically employed, the Gini coefficient takes a value of 0 if all members of the population have the same income. At the other end of the spectrum, the coefficient is 1 in the case of complete inequality: one member of the population is the only one with *any* income. Gini varies smoothly between these two extremes.[4]

The analogue to bursty opinion change is clear: a perfectly bursty opinion stream is one where all of the change happens in a single time step. Applying the Gini coefficient to $\vec{Q}$ in this case would yield a coefficient of 1. A stream where change is spread evenly

[3] Kleinberg bases his measure of burstiness on a discrete Markov process, which requires an arbitrary tuning parameter to identify cutoffs for bursts of different sizes.

[4] A full description of the Gini coefficient is outside the scope of this paper. Believe it or not, wikipedia has an excellent article on the subject.

across time would be completely non-bursty and would have a Gini coefficient of 0.[5]

5 Results

With these measures for analyzing streams in hand, I turn to empirical applications. I draw on 20 opinion streams from the 2004 NAES and focus particularly on the interpretation of two measures: estimated real opinion change and burstiness.

Real opinion change is a measure of inverse stability: the more a stream changes, the less stable it is. Knowing how stable a given stream is can shed light on the aggregate information process driving opinion change. Presumably, an opinion stream that fluctuates widely and often indicates an information flow that delivers new, persuasive[6] messages to a broad audience with some frequency. Also, the stream must include messages that disagree meaningfully with each other, or else opinions would change in one direction and settle there permanently.

A stream that doesn't change might indicate one of two things. First, no change might simply show that no information flow is occurring, or that the information is being ignored by an inattentive or unreceptive public. Alternatively, no aggregate change might be the result of polarization effects of the kind Zaller describes, with aggregate opinion staying constant while comparably sized subgroups sort out to opposite sides. Given sufficient sample sizes, subgroup analysis can sort this out.

Table 1 presents twenty opinion streams drawn from Annenberg data. Most of these

[5]For unbiased results, this operationalization assumes that sample sizes are roughly constant across the span of the series. This condition is clearly violated in some of the opinion streams I work with. For instance, questions about John Kerry early in the race had low sample sizes because Kerry's name recognition was still low. It is not clear how much this affects the Gini coefficient.

[6]Keep in mind that "persuasive" is a function of the audience as well as the message.

variables concern the 2004 presidential election; several ask about the economy; a few demographic variables at the end serve as an intuition check. It seems reasonable to expect that race, gender, and household income should exhibit very high stability and therefore very low real change over the course of the campaign.

Table 1: General Results for twenty 2004 NAES opinion streams, sorted by estimated real opinion change.

	Total Change	Est. Noise	Est. Real Change
Bush Job Approval	20.02	17.79	2.23
Approve of Bush Handling Iraq	19.51	17.74	1.77
Kerry Inspiring	26.73	25.29	1.45
Country Going in Right Direction	20.17	18.94	1.23
Voted in 2000 General Election	15.82	14.66	1.16
Bush Cares About People Like Me	20.58	19.46	1.12
Kerry Trustworthy	28.22	27.12	1.09
Registered to Vote	13.94	13.09	0.85
Bush Inspiring	20.34	19.49	0.85
Kerry Favorability	19.00	18.16	0.84
Conservative or Liberal	15.50	14.87	0.63
Bush Favorability	17.85	17.49	0.36
Race	12.91	12.72	0.19
Countrys Economy Today	15.52	15.62	-0.10
Approve of Bush Handling Economy	17.87	18.02	-0.15
Household Income	18.22	18.62	-0.39
Bush Trustworthy	19.23	19.72	-0.49
Party ID	16.42	16.96	-0.53
Sex	16.91	17.44	-0.54
Personal Economic Situation Today	16.86	17.72	-0.85

This hypothesis is validated by the data. All of the demographics are in the bottom third of the variables listed. Most of the economic streams also display little variation over the course of the campaign. Note that several of the variables have negative estimates for real change. This is the result of normal variation in sampling and should be expected.

Unlike many measures of stability, real change has a clear interpretation. For every one unit of real opinion change, every person in the population changes her mind once (on average) at some point over the course of the stream. Viewed in this respect, some of the estimates are remarkably high. For example, the movement in President Bush's job approval is roughly equivalent to the movement that would occur if his job approval fell to zero, soared to 100 percent, and then dropped to zero again. This is a surprising amount of volatility, especially considering that a small subset of "scorekeeper" voters probably accounts for this swing.

It is also worth noting that random changes in individual-level opinion should not affect the stability of the stream as a whole. In this analysis, "opinion" is the percentage of people who agree with a statement on a given day. This is a characteristic of the aggregate, not any one individual. This construction sidesteps the psychological debate about whether opinions are stored as discrete ideas or sampled from an internal attitude space. Since each p_i is an aggregate measurement, it doesn't matter if people change their answers based on priming experiences, as long as those experiences are personal and uncorrelated among individuals. If experiences *are* correlated across individuals, it seems reasonable to say that information is flowing among them and changing aggregate opinion.

Interpreting the measure of burstiness is a little more difficult. Table 2 shows the same twenty opinion streams, this time sorted by burstiness. I abstract this measure loosely as a way to gauge the "mass mediation" of information. That is, certain opinion streams are likely to be influenced by stories that burst onto the national stage and capture the attention of many citizens simultaneously. In our society, such stories are almost certainly conveyed via television. On the other hand, the information flows that move other opinion streams tend to trickle out through other channels: personal experience, local media, or less jarring national media events. As a first approximation, burstiness can be said to be an indicator of information moving via mass media.

Table 2: NAES opinion streams, sorted by burstiness. Estimated real opinion change is included again for convenience.

	Est. Real Change	Burstiness
Kerry Trustworthy	1.09	0.472
Race	0.19	0.455
Kerry Inspiring	1.45	0.443
Sex	-0.54	0.441
Country Going in Right Direction	1.23	0.435
Conservative or Liberal	0.63	0.435
Countrys Economy Today	-0.10	0.435
Household Income	-0.39	0.434
Voted in 2000 General Election	1.16	0.427
Registered to Vote	0.85	0.426
Kerry Favorability	0.84	0.426
Party ID	-0.53	0.425
Bush Inspiring	0.85	0.421
Personal Economic Situation Today	-0.85	0.418
Bush Trustworthy	-0.49	0.418
Approve of Bush Handling Economy	-0.15	0.414
Bush Cares About People Like Me	1.12	0.411
Approve of Bush Handling Iraq	1.77	0.405
Bush Favorability	0.36	0.402
Bush Job Approval	2.23	0.395

This interpretation is further complicated by the presence of sampling variation. Since normal curves decay exponentially, streams in which apparent change is entirely driven by sampling variation exhibit high burstiness. Race and sex, for example, are very high on the list of bursty variables. For streams with low real opinion change, a high Gini coefficient is more likely to be the result of statistical noise than true information flow.

However, other streams do seem to be driven by true bursty information dynamics. Kerry's popularity and the notion that the country is headed in the right direction are both variables where national media coverage seem likely to be particularly important drivers of opinion. Interestingly, the Bush job approval streams that showed the highest real change have the lowest burstiness. It seems that presidential job approval is driven by many small information events rather than a few prominent ones. This result is consistent with Stimson's finding that approval of the President is largely an expression of satisfaction with government in general.

6 Conclusions

In this paper, I have introduced new techniques for analyzing aggregate opinion streams. By looking at the aggregate change in opinion and burstiness of several streams. One finding that emerged from this analysis is that the scale of opinion change can be substantial. In a given year, a substantial number of people will change their political views at least once. Another important finding is that political information shows a range of burstiness, indicating that not all communications channels are mass mediated equally.

More broadly, this paper can be seen as a demonstration of how aggregation gain can enable aggregate analysis of information flows. This work suggests at least three areas for moving forward. First, aggregate analysis needs to be extended to subgroup analysis in order to identify more precisely where information flows are causeing opinion change. Second, most of my analysis rests on the assumption that information flow is the driving factor behind opinon change. This seems reasonable, but needs to be validated. Finally, more precise statistical measures need to be developed. Burstiness, for example, is a useful concept, but needs to be calculated in a way that is independent from sampling variation.

Information flows remain one of the great unexplored frontiers of public opinion research. For some time now, the field has focused on opinion formation, relying on simplifying assumptions about mass media when thinking about sources of political information. As information technologies and modes of communication proliferate, these oversimplified models of information flow are likely to become less and less accurate. Understanding opinion formation in a world of asychronous, microtargeted, and participatory media will require the develoment of new methodological and theoretical tools.

References

[1] Andrew Gelman; Gary King. 1993. "Why Are American Presidential Election Campaign Polls so Variable When Votes Are so Predictable?" British Journal of Political Science, Vol. 23, No. 4. (Oct., 1993), pp. 409-451.

[2] D. Sunshine Hillygus, Simon Jackman. 2003. "Voter Decision Making in Election 2000: Campaign Effects, Partisan Activation, and the Clinton Legacy." American Journal of Political Science 47 (4), 583596.

[3] Huckfeldt, Robert and John Sprague. 1995. "Citizens, Politics, and Social Communication: Information and Influence in an Election Campaign." Cambridge University Press.

[4] Iyengar, Shanto, and Kinder, Donald R. 1987. *News that Matters*. Chicago: University of Chicago Press.

[5] Kinder, Donald R. 2002. "Communication and politics in the information age." In David O Sears, Leonnie Huddy, and Robert Jervis (editors), *Handbook of Political Psychology*, Oxford University Press.

[6] Kleinberg, Jon. 2004. "Bursty and Hierarchical Structure in Streams." Data Mining and Knowledge Discovery, Vol 7:4.

[7] Lodge, Milton, Steenbergen, Marco, and Brau, Shawn. 1995. "The responsive voter: Campaign information and the dynamics of candidate evaluation." American Political Science Review 89:309-326.

[8] Putnam, Robert. 2000. *Bowling Alone: The Collapse and Revival of American Community.* Simon and Schuster.

[9] Sears, David and Nicholas Valentino. 1997. "Politics Matters: Political Events as Catalysts for Preadult Socialization." American Political Science Review 91:45-66.

[10] Stimson, James. 2004. *Tides of Consent*. New York: Cambridge University Press.

[11] Zaller, John. 1992. *The Nature and Origins of Mass Opinion,* Cambridge University Press.