

ASSESSING THE CAPACITY OF MASS ELECTORATES

Philip E. Converse

*Department of Political Science, University of Michigan, Ann Arbor,
Michigan 48105; e-mail: pconvers@umich.edu*

Key Words elections, issue voting, political information, democratic theory, ideology

■ **Abstract** This is a highly selective review of the huge literature bearing on the capacity of mass electorates for issue voting, in view of the great (mal)distribution of political information across the public, with special attention to the implications of information heterogeneity for alternative methods of research. I trace the twists and turns in understanding the meaning of high levels of response instability on survey policy items from their discovery in the first panel studies of 1940 to the current day. I consider the recent great elaboration of diverse heuristics that voters use to reason with limited information, as well as evidence that the aggregation of preferences so central to democratic process serves to improve the apparent quality of the electoral response. A few recent innovations in design and analysis hold promise of illuminating this topic from helpful new angles.

Never overestimate the information of the American electorate, but never underestimate its intelligence.

(Mark Shields, syndicated political columnist, citing an old aphorism)

INTRODUCTION

In 1997, I was asked to write on the topic “How Dumb Are the Voters Really?” Being revolted by the question formulation, I instantly declined to participate. Long ago I had written two essays (Converse 1964, 1970) to convey limitations on political information in the electorate. Consequently, I found myself typecast, in some quarters at least, as an apostle of voter ignorance. Hence my aversion. Shortly, however, I decided that with a change of title I could take the assignment.

The pithiest truth I have achieved about electorates is that where political information is concerned, the mean level is very low but the variance is very high (Converse 1990). We hardly need argue low information levels any more (e.g. Kinder & Sears 1985, Neuman 1986). Indeed, Delli Carpini & Keeter (1996) have recently provided the most sustained examination of information levels in the

electorate in the literature, in an excellent and thoughtful treatment. They (and I) concur with Luskin (1990) that contrasts in political information have at least three broad sources: ability, motivation, and opportunity. “Dumbness” as commonly conceived is but a part of one of these sources. As this essay proceeds, the impoverishment of the question “how dumb are the voters really?” will become still more apparent.

The essay focuses instead on the second half of my characterization: the extreme variance in political information from the top to the bottom of the public. This is not controversial either. But the degree of this heterogeneity is widely underestimated, and the implications of that dramatic heterogeneity for research seem even less well understood. Hence, I discuss along the way some impacts of this heterogeneity on alternative methods of assessing voter capabilities.

This review emphasizes the relatively recent literature. It also clarifies, where relevant, what some authors still treat as residual mysteries in my two early pieces in the area (Converse 1964, 1970). Moreover, I import several findings from our large mass elite study in France, carried out in the late 1960s but not published until much later (Converse & Pierce 1986). This is an important study in the context of this review for two reasons: (a) It was the first study designed specifically to test the theories in those two early essays, since their hypotheses had merely been suggested by data gathered for other purposes; and (b) crucial results from the French project remain unknown to most students of American voting behavior, presumably because they were studied in a foreign electorate, and who knows what that might mean for external validity.

THE ROLE OF INFORMATION

When in the late 1950s I experimented with analyses stratifying the electorate into “levels of conceptualization” (Campbell et al 1960), I was impressed by the sharpness of the differences from “top to bottom” of the potential electorate in other respects as well. I came to feel that many empirical studies of voting behavior that did not routinely stratify the electorate in some such fashion were actually concealing more than they revealed. In recent research, some form of this stratification has become quite commonplace. The variable I originally thought was probably the clearest differentiator—formal education—had the advantage of being present in most political surveys. Although it is still used, and often to good purpose (e.g. Sniderman et al 1991), I later decided that it gave weaker results than multi-item measures of entities such as political involvement, provided these measures captured enduring interest and not merely the excitement of a specific election. The question of what predictor is best has remained alive, however, and authors using some shorthand for the core variation at issue choose among a wealth of terms to follow the adjective “political”: awareness, attentiveness, expertise, informedness, interest, involvement, knowledge, or sophistication, to name a few. There are different nuances here, but a central construct lurks.

Zaller (1992:333ff) reviews a good deal of experimentation that has led him to prefer a broadly based measure of political information for the crucial discriminating dimension. I heartily applaud the choice. At a theoretical level, I had pointed to the “mass of stored political information” (1962) as a crucial variable in voter decision making, but I never had a good measure of political information to work with in the studies I used, all of which—including the French study—were designed before 1967.

The conventional complaint about measures of political information is that knowledge of minor facts, such as the length of terms of US senators, cannot address what voters actually need to vote properly. This is a tiresome canard. Information measures must be carefully constructed and multi-item, but it does not take much imagination to realize that differences in knowledge of several such “minor” facts are diagnostic of more profound differences in the amount and accuracy of contextual information voters bring to their judgments (Neuman 1986). Absent such imagination, scholars should review Chapter 4 of Delli Carpini & Keeter (1996) for extended proofs. In any event, measurements gauging what these authors denote as “political knowledge,” i.e. “the range of factual information about politics that is stored in long-term memory,” may be the most efficient operationalization of the latent dimension sought.

Evidence of Maldistribution

In my view, the maldistribution of information in the electorate is extreme. Yet Delli Carpini & Keeter, assessing the “Actual Distribution of Political Knowledge” (1996:153), find rather modest differences empirically. A Gini coefficient that I calculate from the distribution of respondents on their main measure (Delli Carpini & Keeter 1996:Table 4.6) shows a weak value of only 0.20. (The Gini coefficient norms between 0.00—when a resource such as information is equally distributed across a population—and 1.00, when one citizen possesses all of it.) This cannot reflect the actual maldistribution in the electorate, which would surely register a Gini coefficient over 0.60 and probably much higher.

At issue here is the easiness of the items making up these authors’ test. It would be possible to devise a test on which everybody sampled could get a perfect score. This would produce a Gini coefficient of 0.00. It would also be possible to use items so arcane that only one person in the sample could get any correct answers, producing a coefficient of 1.00. (This would not mean that subject-matter experts could not answer those items but only that the sample contained at most one such expert.) Of course, no analyst wants to waste time asking questions that do not discriminate, i.e. that nobody or everybody can answer. Indeed, Delli Carpini & Keeter show that their median for correct responses is 49%, proof that the test is nearly optimal to sort on information levels for this sample. Their median does not imply, however, that this is how political information is naturally distributed.

Here is a thought experiment. The universe of all possible political information is, of course, huge and undefinable. But there are subdomains of this universe

that are fairly concrete and listable. One of these is the set of political personages, defined as broadly as one wishes in time and space, such as US figures of the past 40 years. Let us lock up in a room with food and water a sample of respondents asked to list as many such personages as they can resurrect from long-term memory, provided they associate with each name some validating descriptor, however brief. There is reason to expect that, even in an adult cross section, the median number of names exhumed with proper “stub” identification would not be very large: twenty? Thirty? And a few would not get beyond two or three.

Let us now add a well-informed member of the same population—Nelson Polsby, for example. Even within the 40-year time window, the number of relevant subcategories is huge: the federal establishment, with administrations, agencies, and cabinets; both branches of Congress; the high judiciary; national party leaderships; state gubernatorial administrations and legislatures; city mayors and administrations; other influential party leaders; and so on. Nelson might well achieve a list of several thousand, or three orders of magnitude greater than the least informed citizen. If we relaxed the time limit, so that the whole history of the republic was eligible, Nelson’s edge would simply grow larger still. Critics might say that no current citizen need know anything about earlier personages at all, but surely familiarity with the roles of John Jay, Boss Tweed or Joe Cannon enriches the context Nelson brings to current politics. But this is just the “who.” Another simple category is the “where,” since unit political interactions are enormously affected by geographic relationships. The National Geographic Society has found some respondents cannot even find their home state on a unlabeled US map, and the modal respondent cannot label very many states. Nelson could do the whole map in no time, and most of the world map, too, adding rich geopolitical lore on many subregions.

Yet “who” and “where” are the easy parts. We can move on to much more nebulous but profound contextual matters. “Rules of the game” for various significant agencies of the United States would be one of the simpler categories here. Tougher yet would be descriptions of actual process in Congress, for example, with the large range of parliamentary maneuvers that various situations can activate in both houses. Nelson could presumably write on this topic, from memory and without repetition, as long as the food held out, in contrast to most of the electorate, who would produce nothing at all. Of course Nelson Polsby is unique in many ways, but there are hundreds of thousands of citizens in the most informed percentile of the electorate whose “dumps” from stored political knowledge would have awesome dimensions, although it might take 10 or 20 national samples before as many as one such person fell in a sample. A small fraction of the electorate claims a large fraction of the total political information accessible in memory to anyone, hence the predictably high Gini coefficient.

Why such maldistribution? Downs (1957) pitted information costs against nearly invisible control over outcomes in order to explain low information levels. But to explain maldistribution we must add the aphorism “it takes information to get information.” Consider Paul Revere watching the North Church steeple for

the signal to begin his ride. This signal, in modern terms, transmitted only one bit of information, “one if by land, two if by sea.” To digest that message, however, Revere had to know the code, as well as what it meant in context. So it took much more information to receive this one bit.

The same argument applies easily to much more complex transmissions. Stored information provides pegs, cubbyholes, and other place markers in the mind to locate and attribute meaning to new information coming in. The more pegs and cubbyholes one controls in storage, the lower the cost of ingesting any relevant piece of new information. This is a positive feedback system—“them what has, gets”—and it explodes exponentially, thus explaining extreme maldistribution quite simply. Perhaps people without much stored information on a subject are “dumb,” but that is a rather primitive form of judgment.

Implications of Maldistribution for Research

The extravagant heterogeneity of information levels from top to bottom in the electorate should remind us to interpret research findings in terms of the layers of the electorate generating any particular body of data. There is some recognition of this fact; data on political information from college sophomores are acknowledged to lack external validity. But we too easily assume that “cross-section samples” are totally comparable. They are not.

For one thing, response rates have declined steadily over the past 40 years, and the cumulative difference between samples of the 1950s and those of today is sobering (Brehm 1993). The election studies of the 1950s routinely had response rate percentages in the mid- to upper 80s, and studies financed for exceptional follow-up efforts got as high as 92%. Well-financed efforts nowadays have trouble reaching 75%, and ordinary telephone interviewing response rates lie closer to 60%. I have seen hurry-up phone samples of which college graduates form the majority. This is the secular trend: For comparable levels of effort, nonresponse has approximately tripled since 1955. Within each period, of course, the effort to pursue respondents varies widely from one national study to the next.

When nonresponse grows, the least informed do not drop out en bloc. There is merely a correlation in this direction. Moreover, some surveys try to reweight their samples to restore proportions of the less informed, although it is not always easy to learn whether or not such adjustments have been made. And a recent view challenges whether nonresponse typically affects results much at all. (Evidence that it does not, however, comes mainly from checked-box “opinions,” in which underlying quality of response is always well concealed, rather than from open-ended materials in which quality differences leap out.) In any event, the decline of response rates gives commendable pause to careful scholars such as Delli Carpini & Keeter (1996:66) in comparing measures of information in the electorate from five decades of national samples. All should be wary of the problem.

Outside the regimen of full cross-section sampling, it is even more important to “keep score” on the fraction of the public providing the bases for inference.

Memorable vignettes lead easily to working hypotheses, which in turn harden into convictions as to what goes on in the mind of “persons in the street.” Popkin (1994) provides a charming discussion of some modes of “gut reasoning” about politics (which he sees in clinical settings, e.g. discussions with focus groups during political campaigns) that he calls low-information rationality. It is fun to read between the lines what types of citizens these insights come from. Most obviously, Popkin refers consistently to what “voters” do. Moreover, in context it is clear that these are serial voters, which dismisses roughly the bottom half of the electorate. Further, it turns out that these informants are disproportionately serial voters in presidential primaries, a much more exclusive set. Although voting in primaries is notoriously situational (Norrander 1992), it seems likely that most of the sources for Popkin’s insights are part of, in information terms, the top quartile of the public.

This is no cavil at the Popkin (1994) description but rather a neutral effort to locate the discussion in a broader scheme. I heartily endorse the message that voters reason, and do so all the way down the information ordering of the electorate, in the simple sense of putting two and two together using the information accessible to them. Nor do I object to the label of low-information rationality just because of the high-information stratum of the electorate that Popkin seems to have considered. It takes information to generate new combinations of information, and this is true to the very top of the information hierarchy.

The moral of this section is humble. In the endless argumentation over the capacities of the electorate, the steepness of the information slope in the public means that data provenance must be kept in mind. Rancor sanctified by “data” is mindless when, as is not uncommon, contrasting results actually stem from differences in method.

THE RIDDLE OF RESPONSE INSTABILITY

Undoubtedly the greatest surprise in early survey research came when Lazarsfeld et al (1948) reinterviewed the same people at various points during the 1940 presidential campaign, using a design he christened a panel study. Although the study hinged on measuring preference shifts, it turned out that overall the preference distributions, or “marginals,” rarely showed significant change; but there was a remarkably high rate of change caused by individuals shuffling back and forth across the available response space from one interview to the next, even though the intervals between measurements were very short, such as a few weeks. In sum, a great deal of gross change usually added up to little or no net change. This mystery was one of the factors that led Lazarsfeld to develop “latent structure analysis,” grounded in the view that individual preferences were only probabilistic: Given a set of alternatives, the respondent could not be expected to choose any single one with certainty but had a range of probabilities of acceptance across those alternatives.

The Notorious “Black-and-White” Model

When our Michigan group did a panel study over the 1956, 1958, and 1960 national elections, the same mystery fell in our laps. There was variation in temporal behavior across political attitude responses. Party identification, for example, showed substantial stability as measured by correlations from one wave to the next. Moreover, the correlation of reported partisanship between 1956 and 1960 was not much greater than the product of the two two-year correlations, implying limited amounts of steadily progressive “true” change. On the other hand, eight items measuring the most gripping policy issues of the period showed little net change from election to election, and the temporal intercorrelations were also remarkably low. Furthermore, for these items, the four-year intercorrelations were barely higher than the two-year ones, a configuration thought to signal “no true change, all change is measurement error,” or item unreliability. This high gross change without net change was the essence of response instability.

One of the eight issues was more extreme in all of these diagnostic particulars, and I began to consider it a pearl of great price as a *limiting case* of the response instability syndrome. This was the “power and housing” (P&H) issue, an agree-disagree item that read, “the government should leave things like electric power and housing for private businessmen to handle.” The beauty of a limiting case, or boundary condition, is that in the midst of complex problems in which the unknowns outweigh the knowns, it often permits one to set a previous unknown to zero or a constant, thereby getting new leverage for inference. In this instance, if P&H were extreme enough, it would mean true change could be set to zero. And extreme it was: Four-year correlations were essentially the same as two-year ones, and more remote from the product of the two-year correlations (standard deviation = 3.20 in the mean and variance of the other items). As a bonus, the fraction of respondents who, when invited to participate, said they had no opinion on the issue was extremely high (standard deviation = 3.40).

Although it was troubling to posit no true change in attitudes between interviews on any item, it was clear that P&H was the best candidate of all the issues if such an assumption had to be made. The item aimed at measuring the central ideological divide between socialism and capitalism: nationalization versus privatization of industry. This issue had deeply polarized industrial states in the wake of the Great Depression, and nationalization had been ascendant in the 1930s and 1940s. The pendulum would later swing the opposite way, but in the late 1950s the issue was in a kind of repose. The politically attentive had long since chosen sides, while for the inattentive P&H remained a remote issue. Unlike the other issue domains measured, there were no major relevant events or even debates with any likely impact on P&H positions in the whole 1956–1960 period.

The black-and-white model took its name from the division of respondents into two groups: those with fixed true opinions pro or con, and those whose responses were totally unstable, “as though” random. This did not mean they were uncaused, to my mind. Indeed, in this period I enjoyed shocking students by pointing out

that the results of coin flips were also caused. Given enough information about attendant conditions—exact thrust and spin, for starters—the head-tail outcome could be predicted. But it is the resultant of such a large number of forces that throws of an unbiased coin can be treated as random.

Since the proportion of these two groups could be defined between a first and second wave, it was possible to test the model with an independent prediction involving wave three. Changers between the first two waves were all of the error type, and the correlation of their responses between the second and third waves should be 0.00. The second group is a mix of stayers and changers, but in a known proportion, so we could predict that the correlation of their P&H responses between 1958 and 1960 would be 0.47. The observed correlations were 0.004 and 0.489, respectively, an amazing fit.

In presenting these findings I tried to make clear that the P&H issue was a limiting case, because of its location at the extreme boundary. Absent this extreme location, there was no warrant for assuming away true change, which could run in both directions and undoubtedly affected all the other issue items. My explanations clearly did not register. Some supporters wrote to entreat me to stop being so rigid and simplistic about a black-and-white model; if I would just lighten up and admit a range of grays, I would have a much more useful model of attitude change (but, alas, one which would be quite underdetermined for the information available!). Detractors, on the other hand, applied the simplistic model to the other issues despite my advice and, finding garbage results, used them to “prove” that my P&H inference must also be garbage. What both sides had in common was a basic incomprehension of the role of limiting cases in inquiry.

The success of the black-and-white model in illuminating response instability led me to ask whether anybody could answer these policy issue items without huge amounts of response instability. Happily, my colleagues Warren Miller and Donald Stokes had in fact just questioned a sample of US congressmen on these very issues. Considering item intercorrelations to be an indicator of how tightly structured or “constrained” (and free of casual response instability) the respondent’s system of policy attitudes was on these items, I compared such intercorrelations for congressmen with those of the mass sample. The results (Converse 1964:Table VII) showed much higher intercorrelations for the elite respondents. This seemed to establish an extremely plausible direct relationship between information levels and reliability of response. Active politicians simply brought a lot more to the subject matter than many citizens did, and did not have to “guess” at the answers.

I reported these results to two audiences of different interests. The main essay (1964) was written for political scientists. But I read a shorter paper using these results to an International Congress of Psychology in 1963 (later published, 1970). Psychologists at the time were much more versed than political scientists in issues of measurement error, so there was no need to pull punches with them. The issue I wished to highlight stemmed from my student days in psychometrics. Stamping coefficients of reliability on psychological test batteries was then *de rigeur*, and I had been uncomfortable with the apparent precision of these coefficients. It clearly implied that reliability was a fixed attribute of the printed instrument,

invariant over subjects. Of course, psychology was a different world, since for most items in such batteries, the wording was from daily life and the respondent was sovereign—"Do you like carrots?", "Are you uncomfortable in large crowds?" My data on political matters suggested that reliability could vary markedly with the amount of information brought to the subject matter of the test. My climactic statement was: "While the classical view of these matters took 'reliability' to be a property ... attached to the measuring instrument, we could not have a more dramatic example [than the black-and-white results] that reliability in our field of inquiry is instead a joint property of the instrument and the object being measured" (Converse 1970:177).

The phrase "what people bring to" a given subject matter is vague. But it refers to what has anciently been called the "apperceptive mass" in problems of perceptual recognition. In politics, it refers to the stored mass of knowledge mediating what respondents, answering on the fly as always, bring to their decoding of the question.

Critiques and Response

Having been away from this controversy for some time, I was interested in current views of response instability. Zaller (1992) has made major new contributions to the discussion (see below). But Zaller (1992), Page & Shapiro (1992), and others in the current decade conclude that a diagnosis of "just measurement error" seems to have won out over whatever it was that Converse was talking about. The latter is often left a little vague, but since I too was talking mainly about measurement error, I assume the discipline has decided that contrary to my demonstrations, information differences have no impact on measurement reliability. Supporting this verdict, it is true that Achen (1975) and Erikson (1979), frequently cited authorities, were unable to find such differences. But as Zaller has also pointed out, Achen and Erikson are in a minority in this regard. Others have found such differences easily, and in the French project (Converse & Pierce 1986), designed to study the question, they are large indeed.

I am also surprised by descriptions of the measurement-error interpretations of response instability as being relatively novel (Zaller 1992:31). Perhaps I misunderstand the referent, but the view that item responses are probabilistic over a range or "latitude of acceptance" is 70 or 80 years old in psychology, and "latent structure analysis" in sociology dates from the 1950s. That is in fact the view I teathed on, and my chief amendment has been that latitudes of acceptance are broader or narrower according to variations in what respondents bring to the items. Indeed, when I think of "attitude crystallization," the construct refers to the variable breadth of these latitudes. Nor is it true that correcting away error variance is novel. Joreskog has superbly dissected the many facets of error in LISREL, but the root calculation—then called "correction for attenuation"—originated in the 1920s. In fact, it is because of that correction that users of psychological tests by the 1930s began to require the printing of reliability coefficients on test batteries, so that results could be "corrected" up into more impressive regions. I knew that correction well when writing the "Belief Systems" material, and I momentarily

considered applying it at least illustratively, but decided that it would willfully conceal exactly what was substantively important in the responses to issue items.

“Just Measurement Error”

The “just measurement error” folks present very appealing but misleading pictures of individual-level ranges of response. These probabilities are usually graphed as normally distributed, which is reasonable, except near endpoints. But the response probability distribution is rarely shown as straddling the midpoint between pro and con, which starts the reader wondering; and they typically take up rather small fractions of the possible range of responses (e.g. Page & Shapiro 1992:20). If latitudes of acceptance are on average as narrow as these authors depict, then the test-retest reliability of the item measured over short periods would have to be up in the 0.80s to 0.90s. But what I was talking about were issue items in which the apparent reliability never attained 0.60, averaged about 0.40, and at the bottom was under 0.30. This state of affairs can be diagrammed also, but the pictures look totally different from those used to advertise the plausibility of “just measurement error.” Such realistic latitudes of acceptance sprawl broadly, with notable probability densities over the whole response continuum. In the P&H pro-con case, they show a rectangular distribution over the pro-con alternatives for much of the sample. What kind of issue “positions” are these? I thought I was providing solace by showing that where respondents were familiar with an issue, reliability of measurement was higher.

Achen (1975) and Erikson (1979), as mentioned above, are unable to find any impact of information differences on related measurement error. Their difficulty is worth reviewing in more detail. Both are in the awkward position of needing to disconfirm the substantive hypothesis of information effects on response error rates. This means failing to reject the null hypothesis of no differences in reliability as a function of information. A glance at the two terms in significance tests shows that it will be easiest to disconfirm if the natural variance of the independent variable can be artificially truncated and if the test *Ns* can be minimized. Neither author deals with the full range of variance; there is no elite stratum in their tests, and our French data (Converse & Pierce 1986) suggest that this by itself truncates the variance by about one third. Other steps are taken in the same direction. Most notably, for test purposes, respondents are required to have expressed substantive opinions on all three waves. As observed on the P&H issue, many respondents in any given wave have no opinion; to demand three consecutive opinions reduces test *Ns* dramatically. It also differentially discards cases at the least informed end of the electorate, so that the test variance is further truncated artificially. Thus, this editing gains two for the price of one in favor of disconfirmation.

Erikson’s Critique Erikson (1979) bears on the black-and-white model more directly than Achen (1975). Erikson’s article is a masterpiece of organization, and the questions it asks of this model and the P&H issue in particular are entirely germane. It also shows that the black-and-white model is indeed a spectacular fit

to the dynamics of the P&H issue. Erikson's main point, however, is that rival models produce equally good, and probably preferable, fits with the data. I beg to differ, mainly because crucial tests with the black-and-white model reduce to cross-time transition probabilities, and neither of his main challenges addresses this at all.

The first challenge is that a more likely model would not split error into a 0–100% contrast (the black-and-white way) but would instead spread it evenly over all respondents. This is the view that would preserve the practice of stamping reliability coefficients on measuring instruments. Erikson's Table VI is the proof of this contention, but it fails to represent transitions from one wave to the next, which is where the crucial test centers. We can try to finish the author's argument, however, to see where it leads. If reliability is equal for all respondents, then it must be 0.37, the overall average. If we then ask what the correlation of responses between the second and third waves looks like for Erikson, the answer is simple. For every possible bipartite division of the second-wave population (an extraordinarily large number), the correlation of responses between waves two and three must be 0.37. This is very different from the theory-driven prediction of the black-and-white model, which was a sharp bifurcation into temporal correlations at 0.00 for one subset and at 0.47 for the other. As mentioned earlier, the actual results were 0.004 and 0.489. It is not clear why the author claims his pair of 0.37s would fit the data equally well!

The author's second challenge has a different base, but it too ignores the fact that the crucial test hinges on intertemporal correlations. Erikson (1979:104) argues that if preferences for changers on the P&H issue are truly random (instead of "as though random"), then the responses cannot correlate with anything else. Since he shows nonzero correlations with other issues, our argument seems compromised. His premises are not correct, however.

First, what is randomly generated is a time path, such as the $r = 0.004$ between times 2 and 3. Second, we are not limited to randomness of the equiprobability kind here; within any given p governing the marginals, such as $p = 0.70$ or even $p = 0.9944$, there is a track that shows independence from one trial to the next and would produce a 0.00 correlation (the calculation of chi square "expecteds" goes exactly to this case of independence, for marginals of any lopsidedness). A metaphor for random time tracks where the options are not equiprobable is a sequence of flips of a biased coin. Other early analyses of these agree-disagree issue items had made clear that they were strongly influenced by response set effects of the "acquiescence" type. This response set effect is, of course, stronger for some respondents than for others. In the intertemporal case, this can be thought of as a set of respondents each flipping coins of different biases to answer the questions. Their responses would show temporal correlations of zero, but if the test were performed on two different issues, the correlations between items could be arbitrarily large. So the author's second challenge to the black-and-white model as a limiting case has no more merit than the first.

Erikson goes on to ask whether error in the issue items varies inversely with a multi-item index of information/involvement he labels political sophistication, as

I would predict. The results do not lead him to reject the null hypothesis. Here Erikson's independent variable is well conceived and apparently robust. However, his dependent variable, taken literally, seems not to be a measure of error variance alone but of total natural variance in the responses. If true, this is rather disconcerting. It would include true change, which is lively on some of these issues, and which other empirical work (most notably, Zaller 1992, also Converse 1962) has shown to be usually associated curvilinearly with the information/sophistication hierarchy. If so, then dull results with the author's linear analysis methods might not be surprising.

Achen's Critique Achen (1975) tests the same hypothesis about information differences and error rates, and also fails to reject the null hypothesis. Any scholar addressing this debate should read three "communications" in the December 1976 *American Political Science Review* raising issues about the soundness of the Achen analyses. On the face of it, I prefer Achen's dependent variable for this test to Erikson's because it is an estimate of individual contributions to error variance. But the estimation process here is very murky. On the independent variable side, Achen tests my information-error hypothesis using a global analysis with 12 variables, most of which are face-sheet categories such as urban-rural residence and gender, which have no obvious connection to my theory and dilute the critical test with overcontrols. Only 4 of 12 predictors are in the highly relevant education/involvement department. With these measures predicting to each of the eight issue error variances, Achen reports that the "multiple R ranges from 0.11 to 0.19." This value is so low, he concludes, that response error can have no source in individual differences, such as political informedness. A communication from Hunter & Coggin (1976) points out that given details of its estimation, Achen's dependent variable—individual error variance—cannot even charitably have a reliability of more than 0.25. Noting that Achen wants to correct the issue item intercorrelations for attenuation, they ask why he does not correct his multiple Rs for attenuation also; they would rise to 0.45–0.75 and would be quite eye-catching, suggesting exactly the opposite of his conclusion.

I add that among Achen's 12 predictors, the political involvement/education nexus stands out in predicting even the dilute error variance, although Achen stresses that it does not. First, although Achen and Erikson had the same panel data from which to construct a robust variable to express information/involvement differences, Achen chose the three most feeble involvement measures available. Moreover, instead of combining them in a single index (standard social science practice to maximize their meaningful joint variance), he maintained them as three separate predictors in his array of 12, with the opposite effect of eviscerating that same joint variance. His Table 3 shows that in 7 instances, 1 of the 12 predictors relates to the error measure at a 0.05 level. Five of these seven instances involve an involvement/education predictor, despite the evisceration of those measures! With a more robust involvement variable, significant relationships would clearly have multiplied far enough to have made the disconfirmation verdict untenable.

Achen ends by famously noting that the reason I had showed much higher issue intercorrelations for congressmen than for constituents was that the questions asked of the elites on each issue were phrased differently than the versions designed for the mass sample. They were more elegant and incisive, whereas the mass items were vague and poorly worded, producing confused and ambiguous answers, full of response instability. It is true that the wordings were different. But Achen's view of the wording difference effects was pure conjecture, with no evidence offered. This conjecture has achieved great currency, but we already knew from our 1967 French data (Converse & Pierce 1986) that it was wrong.

Evidence from France

The conflict over error sources has sometimes been labeled as a problem in differentiating "errors in survey wording" from "errors in respondents." These labels are unfortunate because they imply that errors must arise from one side or the other, whereas I had argued that a joint interaction was involved. But so labeled, it has been suggested that the problem is fundamentally indeterminate. A claim of this kind underlies Sniderman et al's dismissal of the debate as "ontological" (1991:17). However, the problem is technically underidentified only for blind analyses of a grid of numbers. By "blind analyses" I mean number massaging in which side information about the substance of the variables is ruled out of consideration. For example, a parallel indeterminacy has dogged cohort analysis; it is crucial to distinguish conceptually the effects of age, period, and cohort, but with only two variables actually measured, no three-way assignment can technically be determined. But again, this is only true for blind inference from an unlabeled grid of numbers. Side information about the variables involved in such cases often shows that some rival "blind" inferences are, in fact, substantively absurd and can be discarded with a clear conscience (Converse 1976). The issue of error sources is formally equivalent to the cohort analysis problem. And the French study (Converse & Pierce 1986) casts this kind of side illumination on the issue with Achen/Erikson eloquently.

Two improvements on technical shortfalls in our earlier "Belief Systems" data were (a) the addition of a two-wave elite panel (French deputies) to parallel the mass panel, giving for the first time comparative mass-elite stability estimates; and (b) the use of identical wording for deputies and for their mass constituents on some issue questions. The results of the wording changes were directly opposite to the Achen conjecture: Elite responses to mass questions were brighter than elite discursive answers to more sophisticated, open-ended questions on the same policy debates. ("Brighter" here means showing larger correlations with obvious criterion variables.) This is no great mystery; given familiarly simple issue item wordings, our elites assign themselves more incisive and valid positions than remote coders can deduce for them from flowery and "two-handed" ("on the one hand; on the other") mini-speeches fingering the nuances of various component policy options.

For the French study, we routinely subdivided the mass sample into three strata defined on a very robust five-item measure of political involvement, yielding a




Figure 1 Stability of personal locations on the left-right continuum for mass (by political involvement) and elite, France, 1967–1968. (From *Political Representation in France* by Philip Converse and Roy Pierce, copyright © 1986 by the President and Fellows of Harvard College. Reprinted by permission of the Belknap Press of Harvard University Press.)

thin top (15%), a broad middle (57%), and a bottom (28%). When relevant, we superposed the elite stratum above the rest. Variability in both constraint and stability across these strata is typically sharp and, of course, neatly monotonic. Figure 7-3 from the main report (reproduced as Figure 1 in this chapter) the most dramatic of these displays, namely the stability of self-locations on the left-right continuum. (The item is identically worded for mass and elite.) In terms of theory, this should indeed be the sharpest display, because it involves the key ideological measuring stick or “heuristic” device that is so ubiquitous in informed political discourse in France, (but is about as weakly comprehended in the French public as the liberal-conservative dimension is in the United States, as probes of understood meaning have shown).

The differentiation here is indeed exquisite. Survey researchers are often forced to “prove” arguments with 5–8% differences and are thrilled to work with 20% differences, especially when demonstrating important theoretical points. In this display, which is of crucial theoretical consequence, the differentiation is nearly five times larger, spread out over most of the unit space. To be sure, “just measurement error” advocates could artificially remove chemical traces of error from the elite stratum, and 10 or 15 times as much error from the bottom stratum, thus smartly “proving” that even the most uninformed Frenchman has ideological self-locations just as firm and stable as those of his deputies. But such a calculation is ridiculous obfuscation, given a competing theory that predicts this near-total differentiation on grounds of information differences. And here, any alleged indeterminacy in the blind numbers is swept away by copious independent side information charting the steep decline in comprehension of the left-right continuum as political involvement declines. Again, “just measurement error” folks can assert that the question “where do you place yourself on this left-right scale?” is impossibly vague for *citoyens* who do not know what “left” and “right” mean anyway. But how does such an assertion prove that error variances show no interaction with information differences, as Achen and Erikson have convinced so many scholars?

Figure 1’s display from France, 1967–1968, is neatly corroborated by data on “attitude crystallization” (stability) in ideological self-placement by political knowledge in the United States, 1990–1992 (Delli Carpini & Keeter 1996), showing the largest range also, despite lacking an elite stratum for comparison. Other French data relevant to this discussion are displays of factor-analytic structures of issue and ideology items for the elite and the three involvement strata (Converse & Pierce 1986:248). The left-right self-placements are dominant factors for the elite and the most involved 15% of the mass population; this role fades to fourth and fifth place in the broad middle and bottom strata, which display a scattering of much weaker factors [a parallel to Stimson’s (1975) findings for the 1972 National Election Study]. Ironically, in both these lower strata, making up 85% of the electorate, the liveliest component of issue responses is a “methods effect,” an artifact of question wording. Kinder & Sanders (1990) have put more meat on such bones, showing the susceptibility of low-information respondents to “framing” effects. And of course, although the French policy item responses show gradients less steep in both constraint and stability by involvement than gradients where an ideology measure is involved, the slopes are very impressive in their own right.

The Zaller “Receive-Accept-Sample” Model

Zaller’s (1992) “Receive-Accept-Sample” (RAS) model is a pioneering effort to grapple substantively with the long-standing riddle of response instability. Better yet, it is not merely a theoretical argument. It reflects years of empirical probes to test the suspicion that response instability stems from the ill-prepared respondent’s hasty weighting, under the pressure of an interview situation, of diverse top-of-the-head “considerations” that help him arrive at a quick, impromptu answer. Such a

respondent does not, in my lingo, “bring much” to the question as posed; but Zaller has shown forcefully that most respondents at least recognize the topic domain and can intelligently bring to bear some relevant substantive considerations. This model is surely more useful than the “coin-flipping” metaphor. It does not turn response instability into some marvelously new stable base for democratic theory, nor does it claim to. But it gives a more penetrating view of response instability, and it lays out a platform from which a new generation of research can proceed, gaining incisiveness with a more substantive political base.

In one sense the “considerations” view is only a small part of Zaller’s contribution, however. His persistence in stratifying the electorate in terms of very disparate information conditions produces dynamic time traces of opinion formation and change that are simply brilliant and make grosser analyses of the “electorate as a whole” look so clumsy and information-concealing that one wants to demand a recount. At points Zaller skates on very thin ice: *N*s diminish rapidly in tripartite stratifications, and the only solution is more funding and larger samples, which is totally contrary to the tenor of the times. But his work is a centerpiece for the contention that new advances in this field are not cost-free.

HEURISTICS

Much progressive work in this area in the past decade or so, apart from Zaller’s, has been engrossed in the issue of heuristics, the mental shortcuts by which modestly informed voters can bootstrap their contribution to democratic decision making. This use of shortcuts, which Simon terms “satisficing,” is of course ubiquitous, although it does not compete intellectually (Luskin 2000) with the rich contextual information that some sophisticated voters bring to their voting decisions. All of the evidence reviewed above has to do with issue voting; competing candidates offer other attractions that can be assessed with less information, such as smiles and high sincerity.

Nonetheless, much can be said about heuristics that amplify even issue voting. Fair space was given to this subject in the “Belief Systems” essay (Converse 1964). Of first importance were cues about the party background of issue options. The second heuristic emphasized the liberal-conservative dimension. Another was “ideology by proxy,” whereby an ideological vote may be doubled (or *n*-tupled) by personal admirers of a charismatic ideologue, or other “two-step flow” effects from opinion leaders (Katz & Lazarsfeld 1955). A fourth entailed reasoning that pivots on liked or disliked population groups other than parties; this heuristic was highly prevalent in the broad middle among both US and French voters. A fifth heuristic described in the “Belief Systems” essay involved the economies of restricted attention in the issue-public sense.

In the past two decades, some of the above themes have been greatly elaborated, and so many new heuristics devised that the list begins to rival the number of human instincts posited by psychologists early in this century. Fiorina (1981) does a marvelous job documenting one prime addition, labeled “retrospective voting,”

whereby voters simplify decisions by focusing on how well or poorly parties or candidates have performed in the past. Sniderman et al (1991) make other clever additions, such as the “desert heuristic.” Furthermore, in an area where details of reasoning are hard to observe, Sniderman et al have attempted to infer them with intricate analyses. Again, in my opinion, their best insights come from stratifying by education or political sophistication.

In some of these instances, however—retrospective voting is a good example—it is not clear whether a given habit of reasoning has its most natural home among the more or the less sophisticated. As noted above, Popkin (1994) appears to have formed his impressions of heuristics in interactions with more informed voters. Indeed, when he lists what “the voter” will try to learn next (given a few impoverished cues), the very heftiness of the list bears little relationship to the political learning behavior of three quarters of the electorate. On the other hand, short-cut reasoning is not a monopoly of the poorly informed. It is an economy of the species, and it simply takes on different scales at different levels of information. Delli Carpini & Keeter (1996) ask how high elites can store such prodigious amounts of information, and the answer is twofold: nearly constant attention, along with various elegant heuristics for organizing and digesting information, such as an ideological continuum. This answer does not deny that under various circumstances, labels such as liberal or conservative take on huge affective charges with very little bipolar content (Conover & Feldman 1981). There are superb historical examples, including the antagonism of southern whites to “liberalism” in the first decade or two after World War II, when the term had come to mean efforts to protect the rights of blacks.

Sniderman et al (1991:272) conclude that political reasoning is not some generic process; rather, because it is a compensation for lack of information about politics, it depends on the level of information diverse citizens bring to the matter. Obviously I endorse this judgment, which is close to my own argument about response instability. If one reduces the matter to sufficient abstraction, then there are versions of syllogistic reasoning, or “putting two and two together,” in which differences from high elites to citizens totally unaware of politics can be reduced to absurdity. Both do it! On the other hand, if we consider the different raw materials of information brought to the situation, then reasoning will indeed assume hugely different paths across these strata.

As far as I can tell, of the many varieties of heuristics discussed these days, an ideological criterion is the only one whose natural home is not disputed. It is always found among high political elites and remains robust within the most attentive tenth to sixth of the electorate, then weakens rapidly as we look lower in the information hierarchy, despite lingering affective residues (in the Conover & Feldman sense) that have real (but attenuated) effects elsewhere. Sniderman et al (1991) try to build a new synthesis about “ideological reasoning” by noting that both cognition and affect (two antitheses) are important in politics. They decide that the Michigan view of ideology was purely cognitive; and that Conover & Feldman’s view is purely affective, so a synthesis is in order. Conover & Feldman can speak for themselves. As for the Michigan version, the superstructure constituted by the basic ideological dimension is indeed cognitive. But the personally

selected locations on that continuum (which positioning is, after all, the main interest) are saturated in affect. In multi-party systems, the first enemies to be liquidated are the nearest neighbors (10 points away on a 100-point scale), before the party moves on to vanquish still purer evil across the aisle. No lack of affect here, for true believers. The terms of the proposed synthesis are strained from the start.

THE ELECTORATE COLLECTIVELY

Over the years, it has become increasingly clear that electorates grade out better in issue terms when they are viewed collectively, as aggregates, rather than as the sum of individuals revealed in sample surveys (Kramer 1971, Converse 1975, Miller 1986, Converse & Pierce 1986, Wittman 1989). Various revisionist analyses under the macro label, most notably MacKuen et al (1989), profit from the clarity of aggregation as well.

The most extensive recent demonstration of such aggregation effects comes from *The Rational Public* (Page & Shapiro 1992), which mines 50 years of sample surveys for trend data on American policy preferences. Its focus is totally on the electorate taken in the aggregate. Although the original data bases had individual data, the summarized data reported here are marginal distributions of preferences either at the level of the total electorate or, in one chapter, marginal distributions within the larger population groupings defined by face-sheet data. Some featured findings have long been familiar to survey scholars. For instance, aggregate divisions of issue preferences are nearly inert over long periods of time; where short-range shifts do occur, it is usually easy to see the events that putatively touched them off; for longer-range trends, as have occurred for example in race policy, the drift is attributable to turnover of the population more than to changing convictions of the same people, although the latter does occur as well; and when changes occur, most demographic groups respond in parallel to an astounding degree, rather than the more intuitive counterpoint of conflicting interests. It is a great contribution to show how these familiar features hold up in an exhaustive long-term array of survey data on issue items.

Of course, all of these observations have to do with net change; the method conceals gross change, including principally the Brownian motion of response instability. The authors are properly aware of what is hidden and conclude that it is "just measurement error" in the Achen/Erikson sense anyway, so nothing is lost by writing it off. They are also aware that almost all individual position change observed in panel studies, which absolutely dwarfs net change, is this kind of gross change, although they sometimes describe the features of "change" in policy preferences that are conceivably appropriate for net change, but exactly wrong if gross change is taken into account. In any event, since net change is what mainly matters in most political conflict, the findings of Page & Shapiro (1992) are another stunning demonstration of the more reassuring "feel" conveyed by the aggregated electorate. The authors describe this as a transformation "from individual ignorance to collective wisdom" (1992:15).

I quarrel only with interpretation. The authors see the net change as some proof of the rationality of the public. I prefer to see this type of change more modestly as “coherent,” meaning intelligibly responsive to events. This is in part because of an allergy to the undefined use of the term rational, since most formal definitions involving the maximization of *expected* utility open the door to a tautology whereby any behavior, however self-destructive, is “rational” because the actor must have envisioned the option as personally useful in some sense or he would not have chosen it. But my objection also reflects doubt about all forms of post hoc “explanation.” The epitome for me is in the nightly news’ explanations of why the stock market rose or fell that day. I imagine a homunculus who culls market-relevant news all day, sorting the items into two piles: bullish and bearish. Then, whatever the net market change is at the close, it is clear which pile gives the “reasons.” A real test for the authors (and the homunculus) is whether they could have predicted the amount and direction of change from the events alone, without peeking at the outcome first. Particularly in the more dramatic cases, I agree with Page & Shapiro that the public has shown coherent responsiveness, although I suspect that a real test over all significant net changes would be, as for the homunculus, a pretty mixed bag. It is unlikely that all significant net changes are in some sense inspired or reassuring.

My other concern has to do with the underlying model in our heads when we are pleased by the signs of enhanced competence of electorates in the aggregate. Miller (1986) relates such improvement to Condorcet’s jury theorem; others have followed suit. The Condorcet model may well reflect one force behind gains in apparent competence through aggregation. But it surely is not the most telling model. It assumes, in Bartels’s words (1996), that individuals contributing to the group judgment are “modestly (and equally) well informed.” This does not seem a promising gambit for diagnosing the electorate, given the staggering heterogeneity of informedness across it.

I have thought for years in terms of a rival model, that of signal-to-noise ratios as developed in communications engineering. This is much better suited to electoral heterogeneity. The noise term fits neatly with the huge amount of gross change that has a net worth of zero. Aggregation isolates a signal, large or small, above the noise. The signal, thereby isolated, will necessarily be more intelligible than the total transmission. This smacks of the black-and-white model, although it easily encompasses “true change” as part of the signal and not the noise. The fact that it is still simplistic does not make it useless as a place to start; its complexity certainly advances beyond Condorcet’s “one-probability-fits-all” thought experiment. It also fits the message metaphor: Voting and political polls all have to do with messages from the grass roots. A recent, homologous model of stock market decisions that distinguishes between two classes of participants, the expert traders versus the “noise traders,” has been shown to fit reality better than the assumption of homogeneous information across traders.

Page & Shapiro (1992) imply by descriptions such as “the rational public” that contributions to the actual signal of net change can come equiprobably from any stratum of the electorate. At the same time, net change in policy positions is uncommon and usually limited in magnitude when it does occur. So such change

need only involve a tiny minority of the parent population. At the same time, all data show that a small minority of the population is very well informed and attentive to events. It would be too simplistic to imagine that all net change comes from the most informed, although numerically this would usually be possible. But it would surely not be surprising if a disproportion of these observed net changes did come from those more attentive at least to particular issues, if not always more generally informed. This may be a minor gloss on the Page-Shapiro message. But it does suggest that the “magic” of producing “collective wisdom” from “individual ignorance” may not be the mysterious alchemy it appears, and that there is nothing here to overturn the long-standing picture of great information heterogeneity in the electorate. The fact of collective wisdom remains, however, reassuring for democratic process.

NEW DEPARTURES

A few new research gambits share the goal of improving understanding of the sources and implications of electoral capacity. Some of these may be the research future in this field.

New Issue Measurement

I have registered my doubt that mass publics have trouble giving stable responses to conventional policy issue items simply because questions are objectively vague or poorly worded. I doubt this in part because more informed voters understand the items easily, as do elites. I doubt it also because well-run surveys conduct extensive pretesting to spot confusing terms and to reduce the policy axes to simplest common denominators. None of this rules out, however, the possibility of finding other formats for policy questions that would address issues that matter more in people’s daily lives. It is unfortunate that experimentation in this direction has been limited in the past half century.

I am intrigued by the work of Shanks on a new formatting of issue items, designed to help isolate what policy mandates underpin public voting results (a major focus in Miller & Shanks 1996). Five batteries of structured questions explore governmental objectives in as many policy domains as may seem currently worthwhile. These address (a) perceptions of current conditions, (b) the seriousness of problems by domain, (c) the appropriateness of federal action in the domain, and the relative priority of rival objectives as gauged by government (d) effort and (e) spending. The main drawback of these batteries is the time they take to administer. Nonetheless, a national pilot survey was mounted during the 1996 campaign (Shanks & Strand 1998), and initial results seem to show uncommon liveliness and face validity. Judgment of these innovations would be premature; we await such tests as panel data on short-term stability of responses. I would not expect this form of measurement to erase differences between information strata in matters of issue constraint or stability, but it might shrink the gaps in some reassuring degree. In any event, the effort bears watching.

Simulation of “Higher-Quality” Electorates

For decades, political pollsters have occasionally broken out estimates of “informed opinion” on various issues, simply by discarding substantial fractions of their respondents who can be considered “ill-informed” by some general (i.e. not issue-specific) criterion. In the past few years, much more elegant ways of simulating more informed electorates have begun to appear (Bartels 1996, Althaus 1998). These gain scientific credibility by isolating informed opinion within demographic categories that presumably reflect competing interests in the political arena, thereby preserving the actual numeric weight of those categories in their final solutions rather than distorting proportions by discarding larger fractions of less informed groupings.

There are enough intricacies underlying these simulations (e.g. adjustments linear only, or reflecting nonlinearities?) to encourage a fair amount of controversy. On the other hand, early findings seem to agree that information does matter, although it undoubtedly matters much more in some situations than others, with determinants yet to be investigated. It is not surprising that differences between actual and “informed” electorates are on average more marked with respect to policy issue preferences (Althaus 1998) than when vote decisions, with their grand compounding of often simpler decision criteria, are similarly compared (as in Bartels 1996), although information clearly matters even in the latter case. These are path-breaking efforts that should inspire a wider range of work.

Deliberative Polling

Fishkin (1991) is conducting a growing series of field experiments whereby proper national samples are given standard political interviews and then are brought together at a central location for further deliberation on one or more policy issues. The procedure has varied somewhat, but in the most extensive versions, the convened sample (or as much of it as can attend) receives a range of expert opinion on the topic(s) at hand and/or competing arguments from political figures. In all cases, the plenary sample is divided randomly into much smaller discussion groups that debate the topic(s) at length. In addition to “after” measures to capture attitude change, typically at the end of the deliberations, there are sometimes longer-range follow-up measures to gauge permanence of attitude change.

These field experiments are, not surprisingly, enormously expensive. They have attracted an astonishing barrage of hostility from commercial pollsters, who apparently feel that these funds would be better used to multiply the already overwhelming number of political polls, and who are affronted by the use of the “poll” name, fearing that the public will come to think of polls as a form of manipulation. The experiments also generate staggering amounts of data, not only through their panel waves but also through the material that monitors expert messages and the dynamics of group discussions, which can be seen as a large number of replications of parallel “deliberations.” Material published to date barely scratches the surface of the findings (Fishkin & Luskin 1999).

However edifying these discussions may be to their immediate participants, they are unlikely to be replicated on any large scale, especially to cover the full range of issues that are likely to be debated in any national election campaign. But their scientific interest is wide-ranging, since they deal very directly with controlled (or at least carefully monitored) manipulation of information conditions affecting issue preferences in a proper microcosm of the natural electorate. Of course this is not the first time the field has profited from smaller-scale experimentation, some of which has been much more tightly controlled and hence incisive in areas such as attitude dynamics (e.g. Lodge & Hamill 1986, Lodge & Steenbergen 1995) or political communications (Iyengar & Kinder 1987). But for those interested in ideals of preference change through increased attention on the one hand, and democratic deliberation on the other, there are intriguing new empirical vistas here to be explored.

Visit the Annual Reviews home page at www.AnnualReviews.org

LITERATURE CITED

- Achen CH. 1975. Mass political attitudes and the survey response. *Am. Polit. Sci. Rev.* 69:1218–31
- Althaus SL. 1998. Information effects in collective preferences. *Am. Pol. Sci. Rev.* 92:545–58
- Bartels L. 1996. Uninformed voters: information effects in presidential elections. *Am. J. Polit. Sci.* 40:194–230
- Brehm J. 1993. *The Phantom Respondents: Opinion Surveys and Political Representation*. Ann Arbor: Univ. Mich. Press. 266 pp.
- Campbell A, Converse PE, Miller WE, Stokes DE. 1960. *The American Voter*. New York: Wiley. 573 pp.
- Conover PJ, Feldman S. 1981. The origins and meaning of liberal/conservative self-identifications. *Am. J. Polit. Sci.* 25:617–45
- Converse PE. 1962. Information flow and the stability of partisan attitudes. *Public Opin. Q.* 26:578–99
- Converse PE. 1964. The nature of belief systems in mass publics. In *Ideology and Discontent*, ed. DE Apter, pp. 206–61. New York: Free
- Converse PE. 1970. Attitudes and non-attitudes: continuation of a dialogue. In *The Quantitative Analysis of Social Problems*, ed. ER Tuftes, pp. 168–89. Reading, MA: Addison-Wesley
- Converse PE. 1975. Public opinion and voting behavior. In *Handbook of Political Science*, ed. FW Greenstein, NW Polsby, 4:75–169. Reading, MA: Addison-Wesley
- Converse PE. 1976. *The Dynamics of Party Support: Cohort Analyzing Party Identification*. Beverly Hills, CA: Sage
- Converse PE. 1990. Popular representation and the distribution of information. In *Information and Democratic Processes*, ed. JA Ferejohn, JH Kuklinski, pp. 369–88. Chicago: Univ. Ill. Press
- Converse PE, Pierce R. 1986. *Political Representation in France*. Cambridge, MA: Belknap Press of Harvard Univ. Press. 996 pp.
- Delli Carpini MX, Keeter S. 1996. *What Americans Know about Politics and Why It Matters*. New Haven, CT: Yale Univ. Press. 397 pp.
- Downs A. 1957. *An Economic Theory of Democracy*. New York: Harper. 310 pp.
- Erikson R. 1979. The SRC panel data and mass political attitudes. *Br. J. Polit. Sci.* 9:89–114
- Fiorina MP. 1981. *Retrospective Voting in American National Elections*. New Haven, CT: Yale Univ. Press

- Fishkin JS. 1991. *Democracy and Deliberation: New Directions for Democratic Reform*. New Haven, CT: Yale Univ. Press
- Fishkin JS, Luskin RC. 1999. Bringing deliberation to the democratic dialogue. In *A Poll with a Human Face: The National Issues Convention Experiment in Political Communication*, ed. M McCombs. Mahwah, NJ: Erlbaum
- Hunter JE, Coggin TD. 1976. Communication. *Am. Polit. Sci. Rev.* 70:1226–29
- Iyengar S, Kinder DR. 1987. *News That Matters: Television and American Opinion*. Chicago: Univ. Chicago Press. 187 pp.
- Katz E, Lazarsfeld PF. 1955. *Personal Influence: The Part Played by People in the Flow of Mass Communication*. New York: Free. 400 pp.
- Kinder DR, Sanders LM. 1990. Mimicking political debate with survey questions: the case of white opinion on affirmative action for blacks. *Soc. Cogn.* 8:73–103
- Kinder DR, Sears DO. 1985. Public opinion and political action. In *Handbook of Social Psychology*, ed. G Lindzey, E Aronson, pp. 659–741. New York: Random House
- Kramer GH. 1971. Short-term fluctuations in U.S. voting behavior, 1896–94. *Am. Polit. Sci. Rev.* 65:131–43
- Lazarsfeld PF, Berelson B, Gaudet S. 1948. *The People's Choice: How the Voter Makes Up His Mind in a Presidential Election Campaign*. New York: Columbia Univ. Press
- Lodge MG, Hamill RC. 1986. A partisan schema for political information-processing. *Am. Polit. Sci. Rev.* 80:505–19
- Lodge M, Steenbergen M. 1995. The responsive voter: campaign information and the dynamics of candidate evaluation. *Am. Polit. Sci. Rev.* 89:309–26.
- Luskin RC. 1990. Explaining political sophistication. *Polit. Behav.* 12:331–61
- Luskin RC. 2000. From denial to extenuation (and finally beyond): political sophistication and citizen performance. In *Thinking About Political Psychology*, ed. JH Kuklinski. New York: Cambridge Univ. Press
- MacKuen M, Erikson RS, Stimson JA. 1989. Macropartisanship. *Am. Polit. Sci. Rev.* 83:1125–42
- Miller NR. 1986. Information, electorates, and democracy: some extensions and interpretations of the Condorcet jury theorem. In *Information Pooling and Group Decision Making*, ed. B Grofman, G Owen, pp. 173–92. Greenwich, CT: JAI
- Miller WE, JM Shanks. 1996. *The New American Voter*. Cambridge, MA: Harvard Univ. Press
- Neuman WR. 1986. *The Paradox of Mass Politics: Knowledge and Opinion in the American Electorate*. Cambridge, MA: Harvard Univ. Press
- Norrander B. 1992. *Super Tuesday: Regional Politics and Presidential Primaries*. Lexington: Univ. Press Kentucky
- Page BI, Shapiro RY. 1992. *The Rational Public*. Chicago: Univ. Chicago Press. 489 pp.
- Popkin S. 1994. *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. Chicago: Univ. Chicago Press. 2nd ed.
- Shanks JM, Strand DA. 1998. *Understanding issue voting in presidential elections*. Presented at Annu. Meet. Am. Assoc. Public Opin. Res., St. Louis, MO
- Sniderman PM, Brody RA, Tetlock PE. 1991. *Reasoning and Choice: Explorations in Political Psychology*. Cambridge, UK: Cambridge Univ. Press
- Stimson JA. 1975. Belief systems: constraint, complexity and the 1972 elections. In *Controversies in American Voting Behavior*, ed. RD Niemi, H Weisberg, pp. 138–59. San Francisco: Freeman
- Wittman DA. 1989. Why democracies produce efficient results. *J. Polit. Econ.* 97:1395–424
- Zaller J. 1992. *The Nature and Origin of Mass Opinion*. Cambridge, UK: Cambridge Univ. Press. 367 pp.