TRACKING OPINION OVER TIME
A METHOD FOR REDUCING SAMPLING ERROR

DONALD P. GREEN
ALAN S. GERBER
SUZANNA L. DE BOEF

Abstract Across a wide range of applications, the Kalman filtering and smoothing algorithm provides survey researchers with a single, systematic technique by which to generate four kinds of useful information. First, it enables survey analysts to differentiate between random sampling error and true opinion change. Second, Kalman smoothing provides a means by which to accumulate information across surveys, greatly increasing the precision with which public opinion is gauged at any given point in time. Third, this technique provides a rigorous means by which to interpolate missing observations and calculate the uncertainty associated with these interpolations. Finally, the Kalman algorithm improves the accuracy with which public opinion may be forecasted. Our empirical examples, which focus on party identification, show that the Kalman algorithm can dramatically reduce sampling error in survey data. Since software implementing this technique is readily available, survey analysts are encouraged to use it to make more efficient use of the data at their disposal.

Public opinion analysts frequently encounter problems when they attempt to track the opinions of subpopulations over time. A random national sample of 1,000 adults will contain approximately 121 African-Americans, 126 Californians, and 127 respondents over 65 years of age (U.S. Census Bureau 1995). Charting trends in opinion among such small subgroups immediately raises the question of whether movements from poll to poll are due to real opinion change or merely random sampling variability.

DONALD P. GREEN is professor of political science and director of the Institution for Social and Policy Studies at Yale University ALAN S. GERBER is assistant professor of political science at Yale University. SUZANNA L. DE BOEF is assistant professor of political science at Pennsylvania State University. The authors wish to thank Jay Emerson for his programming assistance and the Institution for Social and Policy Studies for its financial support. The data and programs used in this article may be obtained at the Web address http://pantheon.yale.edu/~gogreen.

Public Opinion Quarterly Volume 63:178–192 © 1999 by the American Association for Public Opinion Research
All rights reserved 0033-362X/99/6302-0005$02 50

This article discusses a statistical method for distinguishing genuine movements in public opinion from random movements produced by sampling error. The technique, known as Kalman smoothing, is borrowed from engineering, where is it used to extract signal from noise in a wide array of applications ranging from radio receivers to airplane landing gear to digital imaging (Strang 1986). Although expositions of the Kalman filter abound (Hamilton 1994; Harvey 1986; Meinhold and Singpurwalla 1983), seldom has this methodology been used to analyze public opinion (cf. Beck 1990; Gerber and Green 1998). A series of empirical examples focusing on political partisanship demonstrates how Kalman smoothing may enable researchers to draw meaningful inferences about public opinion trends using small samples and Web-based statistical software. We conclude by discussing the potential benefits, as well as limitations, of this approach.

Preliminary Example

Suppose that polls are conducted in July and August, both independent random samples of the adult population with 1,000 respondents apiece. In July, the poll results indicate that 52 percent of the sample approves of the president's performance in office. In August, this rating falls to 48 percent. Has approval actually declined, or is this change a fluke brought about by sampling variability? If approval has declined, what is our best guess of the state of opinion in the second month?

This scenario repeatedly presents itself to journalists, campaign managers, and poll watchers, yet those asked to formulate an expert judgement based on survey results seldom make use of rigorous analytic methods. There is no single answer to hypothetical questions just posed, but rather a range of answers depending on how one envisions the process by which public opinion evolves over time. If, on the one hand, one supposes that opinion is perfectly stable over time, then the best strategy for assessing the true state of public opinion is to take a simple average of the two polls. On the other hand, if one assumes that presidential popularity may shift up or down in any given month, the calculation required to formulate an optimal statistical estimate of approval in August becomes more complex. While the details of the calculation are deferred to the next section, the intuition behind the formulas is straightforward. The greater the month-to-month variability in the true state of public opinion (''signal'') in relation to variability due to sampling error (''noise''), the more likely that shifting poll results reflect genuine opinion change. In this example, where the samples are moderate in size, our best guess is that true opinion has shifted from 51.4 percent to 48.6 percent, implying that approximately one-third of the decline in the polls is due to random sampling fluctuation.[1]

Where do these numbers come from? They were generated by the Kalman smoothing procedure, which offers a method for constructing a weighted average of the available polling information based on the ratio of signal to noise in each poll. Using this optimal weighting scheme—a series of weights that give the smallest average squared errors when predicting the true state of public opinion—the best estimate of opinion in July (51.4 percent) is obtained by weighting the first and second polls by weights of .85 and .15, respectively. The August estimate (48.6 percent) applies weights of .15 and .85 to the two polls. July's poll results help gauge opinion in August, and vice versa.

Since less weight is attached to polls with more sampling error, it is informative to consider how the same example would look with different sample sizes. If the two samples had contained 400 respondents apiece (as would commonly be the case if we were to track subgroups in a national sample), the optimal estimates of public opinion in the two periods would have been 50.4 percent and 49.6 percent; almost all of the observed change is attributed to sampling error, and these estimates are close to a simple average in which each poll is weighted equally. Conversely, had the sample sizes been 4,000, thereby driving down the sampling variance, the weighting scheme would have put more emphasis on the contemporaneous poll. In this case, the new estimates would have been 51.8 percent and 48.2 percent, suggesting that almost all of the change recorded in the polls reflects genuine opinion change.

The beauty of the Kalman algorithm is that it provides a systematic method by which to assign weights to a sequence of poll results, taking into account such factors as sample size and the amount of time elapsed between polls. In the next section, the key features of the algorithm are described for interested readers. Those wishing to avoid these technical details may skip to the empirical examples given below, which make use of Web-based software designed to extract signal from noise in survey data.

Overview of the Statistical Model

For purposes of exposition, consider the problem of tracking party identification over time. Call the true, but unknown, percentage of Republicans in the population ξ. We want to estimate ξ, but we only observe X, our

---

[1] The numbers in this example are maximum likelihood estimates based on the assumption that public opinion follows a random walk with an unknown (estimated) disturbance variance and diffuse priors. A more extensive discussion of our methodology is presented below.

survey's estimate of ξ. The problem is that X is off by a certain degree of sampling error (e). In other words,

$$X = \xi + e. \tag{1}$$

This measurement exercise will be repeated at subsequent points in time, so we denote the terms in the equation with a subscript t ($t = 1, 2, \ldots, T$) to indicate different time points. In the first week, for example, the relationship between true opinion and survey results is

$$X_1 = \xi_1 + e_1. \tag{2}$$

We want our model to take into account the possibility that public opinion might change over time. Suppose that current approval ratings are equal to prior approval ratings plus some perturbation factor that we will call u:

$$\xi_t = \xi_{t-1} + u_t. \tag{3}$$

We can write a more general expression of this relationship:

$$\xi_t = \alpha + \gamma\xi_{t-1} + u_t. \tag{4}$$

Equation (4) allows for the possibility that opinion equilibrates to some underlying level, whereas equation (3) assumes that opinion has no equilibrium, following instead a random walk.[2] Just how to characterize opinion dynamics will vary from application to application. We will demonstrate below that with a sufficiently large number of polls, the parameters α and γ can be estimated from the data, so that there will be no need to stipulate them in advance. For purposes of exposition, however, we focus on the simpler (and often quite useful) case in which $\alpha = 0$ and $\gamma = 1$.

In order to differentiate genuine opinion change from apparent changes due to sampling error we need to know (or estimate) the variances of e and u. The variance of e is relatively straightforward when the data are percentages, since the variance of an observed percentage p is simply $\pi(1 - \pi)/N$, where π represents the true underlying percentage and N the sample size.[3] Less apparent is the variance of changes in underlying opinion, a quantity that typically must be estimated using maximum likelihood techniques.

Suppose for the sake of illustration that we know σ_u^2, the variance of u_t. Intuition suggests that the larger this variance in relation to sampling

---

[2] The equilibrium level to which opinion gravitates in equation (4) is given by $\alpha/(1 - \gamma)$. Smaller values of γ mean that opinion returns to this equilibrium at a faster rate following some perturbation.

[3] Note that the difference between true and observed percentages raises a computational issue For ease of exposition, we calculate the sampling variability by substituting p for π. In our more advanced implementations of the Kalman algorithm, we substitute the smoothed estimate of π for π. One statistical issue that we do not address is the fact that stratified probability sampling, nonresponse, and variation in questionnaire format and interview staff affect the sampling errors associated with polls.

variability, the less polls conducted in the past reveal about the current state of opinion. Real change is afoot, and we should update our assessments based on up-to-date information. Conversely, as the ratio of true change to sampling fluctuation diminishes, current polls lose their edge over older information. In the limit, a poll with an infinite sample size will give an exact reading of public opinion; at the opposite extreme, if true opinion never changes ($\sigma_u^2 = 0$), the best guess of where opinion stands at any point in time is a simple average of all the polls conducted before or since. For these extreme cases as well as situations in between, the Kalman filter algorithm provides a systematic way to attach weights to each available poll.

Kalman Filtering and Smoothing

The Kalman smoothing algorithm proceeds in two stages. First, filtered estimates are generated by moving forward in time, beginning with the first poll and iterating until the last poll is reached. Then, working backward in time, smoothed estimates are generated, incorporating the forecast errors made by the filtering procedure. The purpose of each procedure is to construct a system of weights to attach to each poll when formulating an optimal guess about the true state of opinion at each point in time. Filtering constructs weighted averages on polls up to and including a given point in time, which makes it valuable for forecasting. Smoothing, by contrast, makes use of all of the available polls and is therefore suited to retrospective analyses of opinion change. We now briefly describe each procedure.[4]

Filtering. If we have no information about the state of opinion before our first poll, we must take the initial poll results at face value. Our ''filtered'' estimate of the true state of public opinion during period 1, which we will denote F_1, is simply equal to X_1. In effect, we have multiplied X_1 by a weight of one, while other information has been assigned a weight of zero. From this point on, these weights will be denoted W_t. For now, $W_1 = 1$. But even if we put all our faith in this first poll, we still have uncertainty as we make forecasts about the next period. After examining the first poll, this uncertainty stands at σ_u^2 (the uncertainty created by movement in public opinion over time) plus the uncertainty created by sampling error.[5] The notation for this total uncertainty term is P_1.

---

[4] For a complete presentation of the Kalman procedure and derivation of its statistical properties, see Hamilton (1994) and Harvey (1996).

[5] Here we assume that σ_u^2 is known. In practice, σ_u^2 must be estimated from the data, as illustrated below, which in turn introduces uncertainty into the estimated location of public opinion at any given point in time (see Hamilton 1986).

Having set the groundwork for our filtering procedure, we update our weighting factor W_1 to indicate how much emphasis we should place on the new polling information we receive in period 2. This weighting factor is a function of our current uncertainty about the state of opinion (P_1) and random sampling error in our second poll [$Var(e_2)$]:

$$W_2 = \frac{P_1}{P_1 + Var(e_2)}. \tag{5}$$

This formula has a straightforward interpretation. The greater the sampling error in our second poll, the less weight we assign it when judging the current state of opinion.

Using this weighting factor, we adjust the second poll's results (X_2), so that our filtered estimate of the state of opinion becomes

$$F_2 = W_2X_2 + (1 - W_2)F_1. \tag{6}$$

If the second survey is so large that the results have no sampling variability, then $W_2 = 1$, and all of the weight is placed on the new survey results. Otherwise, our best guess of the state of opinion is a weighted average of current and past information. From a computational standpoint, the beauty of this algorithm lies in the fact that once a poll (e.g., X_1) has been incorporated into a filtered estimate (F_1) it no longer enters directly into any subsequent calculations. The filtered estimate has in effect summarized all of the useful information from prior polls.

Finally, we update our uncertainty as we look ahead toward the following week:

$$P_2 = P_1(1 - W_2) + \sigma_u^2. \tag{7}$$

These steps are simply repeated for each subsequent survey, replacing F_2 with F_3, and P_2 with P_3, until one has worked through all T polls. We now have an optimal reading of public opinion at each point in time given information up to and including that time period. For survey analysts who are tracking public opinion as it unfolds during a campaign, filtering gives the best account of where opinion stands today. If, however, the aim is to look back on a series of polls and reconstruct the manner in which public opinion evolved, an additional smoothing step will further improve the accuracy of our estimates.

Smoothing. Once we have a filtered estimate and uncertainty estimate for each observation, we now work backward in time to create the smoothed estimates. The last filtered estimate (for period T) is identical to the last smoothed estimate (since there is no subsequent information).

Call S_T the smoothed estimate for the observation at time T. The smoothed estimate for the $T - 1$ observation is

$$S_{T-1} = F_{T-1} + (S_T - F_{T-1})\left(1 - \frac{\sigma_u^2}{P_T}\right). \tag{8}$$

This procedure is repeated until we have generated smoothed estimates back to S_1. Intuitively, this equation suggests that the more F_{T-1} mispredicts F_T, the more one should adjust the smoothed estimate, S_{T-1}. However, the more true opinion changes over time, the less information about period T should alter the assessment of opinion at $T - 1$. Smoothed estimates are more precise than filtered estimates, because the latter are based solely on information up to and including the current time period, whereas smoothed estimates make use of all of the poll readings. If we have modeled movements in public opinion correctly and if the e_t and u_t are distributed normally, the smoothed estimates have the smallest mean squared error of any linear weighting scheme that one could apply to a sequence of poll results (Hamilton 1994).

Of course, performing these calculations for more than a few polls would be unduly time consuming if done by hand. A number of statistical software packages offer Kalman filtering, but none to our knowledge can easily handle polls that are unevenly spaced in time, which is a great drawback for applications involving occasional polls. To fill this gap, we adapted existing Kalman filtering and smoothing algorithms to meet the special needs of survey analysts and created an interactive Web site designed to perform the statistical analyses we have described. Our program also has the advantage of using a grid-search algorithm, so that users are not required to supply starting values.[6]

Empirical Examples

To explicate the nature and advantages of the Kalman procedure, we shall discuss some illustrative examples in which public opinion is charted over time. Our first example concerns Republican partisanship in California from 1981 to 1995. The data are drawn from national CBS/*New York Times* polls (CBS/NYT) conducted during this period. Respondents were asked, ''Generally speaking, do you usually consider yourself a Republican, a Democrat, an Independent, or what?'' Table 1 records the percentage of each sample that self-identifies as Republican.

The CBS/NYT telephone surveys have been aggregated quarterly, yet,

---

[6] This site resides on the server maintained at Yale University's Social Science Statistics Lab and may be accessed at the World Wide Web address http://statlab.stat.yale.edu/~gogreen/samplemiser3.html.

Table 1. Observed, Filtered, and Smoothed Estimates of Republican Partisanship in California: CBS/NYT Polls, 1981–95 (Assumes an Autoregressive Parameter of 1.0)

Quarter	*N*	Observed (%)	SE (Observed %)	Filtered (%)	Smoothed (%)	SE Smoothed Estimate (%)
1981:						
1	90	24	4.5	24.0	33.8	1.7
2	240	37	3.1	32.8	33.9	1.6
3	124	37	4.3	33.9	34.0	1.5
1982:						
6	271	35	2.9	34.4	34.0	1.4
7	153	26	3.6	32.4	34.0	1.4
8	133	33	4.2	32.5	34.1	1.3
1983:						
10	226	36	3.2	33.4	34.5	1.2
12	98	32	4.7	33.2	34.7	1.1
1984:						
13	503	33	2.1	33.1	34.9	1.1
14	230	33	3.1	33.1	35.2	1.0
15	636	36	1.9	34.1	35.5	.8
16	1,700	37	1.2	35.7	35.8	.8
1985:						
19	58	24	5.6	35.1	36.1	1.0
1986:						
22	82	34	5.2	35.0	36.6	1.0
23	310	42	2.8	36.7	36.8	1.0
24	328	33	2.6	35.8	36.8	1.0
1987:						
26	104	30	4.5	35.2	37.2	1.0
28	183	43	3.7	36.5	37.7	1.0
1988:						
29	399	40	2.5	37.5	37.9	.9
30	219	39	3.3	37.7	38.0	.9
31	493	41	2.2	38.7	38.0	.9
32	2,238	37	1.0	37.6	37.9	.8
1989:						
33	426	38	2.4	37.7	38.0	.8
34	120	37	4.4	37.7	38.1	.9
35	256	39	3.1	37.8	38.2	1.0
36	134	42	4.3	38.1	38.3	1.0

Table 1. *(Continued)*

Quarter	*N*	Observed (%)	SE (Observed %)	Filtered (%)	Smoothed (%)	SE Smoothed Estimate (%)
1990:						
37	273	37	2.9	37.9	38.3	1.0
38	117	43	4.6	38.3	38.4	1.0
39	195	40	3.5	38.5	38.4	1.0
40	540	39	2.1	39.7	38.3	1.0
1991:						
41	835	41	1.7	39.5	38.3	1.0
42	455	38	2.3	39.2	37.9	.9
43	284	36	2.9	38.7	37.6	.8
44	233	37	3.2	38.5	37.3	.8
1992:						
45	836	36	1.7	37.6	37.0	.7
46	531	39	2.1	37.9	36.8	.7
47	1,289	35	1.3	36.7	36.5	.7
48	820	36	1.7	36.5	36.4	.7
1993:						
49	567	36	2.0	36.4	36.4	.7
50	345	36	2.6	36.4	36.4	.8
51	178	39	3.7	36.6	36.4	.8
52	474	37	2.2	36.7	36.4	.8
1994:						
53	991	36	1.5	36.4	36.3	.7
54	326	34	2.6	36.1	36.3	.8
55	700	36	1.8	36.1	36.3	.8
56	395	37	2.4	36.2	36.4	.8
1995:						
57	944	37	1.6	36.5	36.5	.8
58	463	38	2.3	36.8	36.5	.8
59	492	34	2.1	36.2	36.4	.9
60	394	38	2.5	36.5	36.5	1.0

SOURCE.—ICPSR holdings of CBS/*New York Times* polls, aggregated quarterly by authors

Table 2. Maximum Likelihood Estimates of Kalman Smoothing Parameters, Republican Partisanship, Quarterly, 1981–95

	Constrained Model	Unconstrained Model
Autoregressive parameter (γ)	1.0 (fixed)	.880
SE		(.074)
Disturbance variance	.283	.317
SE	(.235)	(.307)
Log-likelihood	−82.915	−80.552

NOTE —Number of surveys = 50. Sample mean = 36.260. Maximum likelihood estimates obtained using interactive software available at the World Wide Web address http://statlab.stat.yale edu/~gogreen/samplemiser3.html. The intercept term in the unconstrained model is determined by the sample mean. So that the first filtered estimate presented in table 1 is identical to the first observed estimate, the constrained estimates are derived based on the assumption that one's expectation before the first poll is equal to the outcome of the first poll, with an (extremely diffuse) variance of 1,000.

as is apparent from table 1, some gaps remain. No polls were conducted, for example, during the fourth quarter of 1981. One of the virtues of Kalman filtering and smoothing is that it can accommodate unevenly spaced observations; indeed, it constitutes a method for gauging the state of opinion during periods when no survey data are available. As discussed below, these missing time points are merely equivalent to polls in which the number of respondents is zero.

With 50 polls spanning 60 quarters, we have sufficient information to estimate both σ_u^2 and γ using maximum likelihood methods described in Beck (1990). In keeping with the previous discussion, however, we begin by assuming γ to be one. In other words, imagine Republican partisanship to be a random walk that is as likely to move up as down at any given time. Table 2 indicates that, although sample readings of Republican partisanship have sometimes fluctuated markedly, party identification has changed very gradually in California. The smoothed estimates indicate that when Ronald Reagan assumed the presidency, 33.8 percent of all California adults were Republicans. Eight years later (quarter 33), this figure had risen to 38.0 percent. This Republican tide subsided after 1991, dropping slowly to 36.5 percent, where it remained more or less from 1993 through 1995. The average change from one poll to the next is 3.4 percentage points. The average change from one smoothed estimate to the next is just .15 percentage points.

This description implies that almost all of the apparent change in partisanship (e.g., from 42 percent to 33 percent between quarters 23 and 24) may be attributed to sampling variability. How can one determine if this characterization is accurate? One method is to compare one-step-ahead forecasts. Recall that our filtered estimates for each time point make no use of subsequently gathered information.[7] A natural comparison therefore exists between filtered estimates and raw poll results: which does a better job of predicting the outcome of the next poll? Our data set of 50 polls permits us to calculate 49 such one-step-ahead forecasts. The mean squared error for forecasts derived from the Kalman filter is 14.8 percentage points. Using the last poll to predict the next poll results in a mean squared error of 23.4 percentage points. It seems clear that to take each poll result at face value leads one to overstate the volatility of public opinion.

Another important benefit of Kalman filtering and smoothing is a dramatic reduction in uncertainty about the state of opinion at a given point in time. Consider, for example, the last poll result in our series (fourth quarter of 1995), in which 38 percent of the California sample call themselves Republican. Given $N = 394$, a conventional statistical procedure would calculate the standard error of this reading to be 2.4 percentage points. By contrast, the Kalman filtering (or smoothing) suggests an estimate of 36.5 percent with a standard error of just 1.01 percentage points. The dramatic reduction in uncertainty is achieved by making use of the information derived from prior polls, data that are ignored by traditional methods that analyze one opinion poll at a time.

The relationship between the standard errors of each poll, analyzed separately, and the standard errors of the filtered and smoothed estimates is graphed in figure 1. The most striking feature of this graph is the dramatic reduction in sampling uncertainty afforded by the Kalman smoothing procedure. The standard errors of the smoothed estimates are typically less than one-third as large as the one-poll-at-a-time estimates. One way to think of this gain in precision is to imagine the implicit sample sizes associated with the smoothed estimates. In quarter 40, for example, the conventional standard error for the poll showing 39 percent GOP with a sample size of 540 is 2.1. The corresponding smoothed standard error, .98, is what one would obtain with conventional methods assuming a sample size of 2,427.

What about the state of opinion between polls? As noted earlier, we can obtain interpolations and accompanying standard errors simply by inserting poll "results" based on an arbitrarily small sample size. These fictitious polls have no effect on the other smoothed or filtered estimates

---

[7] For purposes of this exercise, we use filtered estimates that assume the sampling variability of each poll to be $p(1 - p)/N$. An alternative approach is to filter and smooth in two stages. In the first stage, one filters and smooths using the algorithms described in the text. In the second stage, one calculates the sampling variability at each time point by substituting the smoothed estimate for p. (Monte Carlo simulations indicate that this two-step procedure provides slightly more accurate parameter estimates)

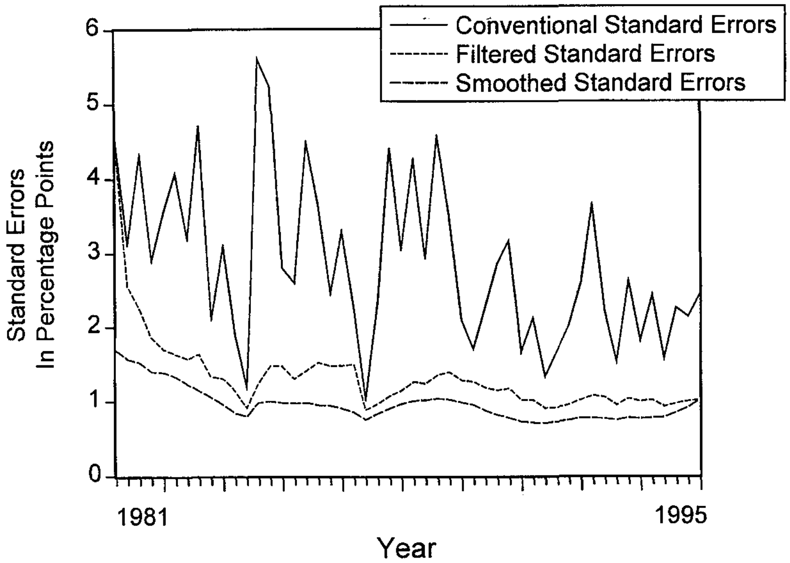

Figure 1. Conventional standard errors and standard errors associated with Kalman filtering and smoothing, for CBS/*New York Times* Poll measures of Democratic partisanship, 1981–85.

because their small sample size will cause them to be assigned a weight of zero in the Kalman filter algorithm. Suppose we want to know the state of opinion in quarter 25, a missing observation in our time series. Using the parameter estimates obtained from our analysis of actual poll results, we include a fictitious poll with a sample size of .0001 (some arbitrarily small number greater than zero). The smoothed estimate for period 25 is 37.0 percent with a standard error of .98. Again, to achieve that degree of accuracy would ordinarily require a sample of over 2,400 respondents.

To this point, our exposition has assumed that $\gamma = 1$, which implies that public opinion follows a simple random walk. An alternative conjecture holds that public opinion equilibrates to some level—a "natural" division of Democrats and Republicans that would occur if opinion were unperturbed by outside forces. Several studies of aggregate partisanship (see Green, Palmquist, and Schickler 1998) have found γ to be in the neighborhood of .95, which suggests that we might improve the fit of our smoothed estimates by treating γ as a free parameter and thereby allowing partisanship to return (slowly) to its long-term equilibrium. The Kalman filter model can be easily generalized to allow the estimation of γ. Consistent with previous results, table 2 shows that the maximum likelihood estimate of γ is .88 with a standard error of .07. Freeing this additional parameter increases the log-likelihood from −82.92 to −80.55, a statistically significant improvement $[-2(L_0 - L_1) = 4.74$, 1 *df*, $p < .05]$. Nevertheless, when we recompute the estimates in table 1 using the revised parameters, we find that allowing γ to depart from 1.0 does little to change the pattern of filtered and smoothed estimates. Apart from the very first observation, the two sets of estimates lie within 1 percentage point. Our experience in working with the Kalman filter has been that when examining the over-time trajectories of slowly evolving series, such as partisanship, the assumption that $\gamma = 1$ typically provides a good approximation. However, researchers working with more volatile series that seem to reequilibrate rapidly in the wake of short-term shocks may find it necessary to estimate γ.

Limitations

The best opportunities to apply Kalman filtering and smoothing arise when we have large numbers of surveys, each with 100 or more respondents. Thus, the analysis of Republican partisanship in California worked well when we attempted to estimate either or both the disturbance variance or the autoregressive parameter. Based on simulation analyses (available from the authors on request), we have found that it is often infeasible to estimate both parameters based on fewer than 25 polls. In such cases, one must either gather more data or stipulate the value of γ or σ_u^2.

Another limitation arises when the input percentages get close to 0 percent or 100 percent. The Kalman model described here assumes a linear model of opinion change. As percentages approach 0 percent or 100 percent, this linearity assumption breaks down. It is in some sense harder to move opinion from 94 percent to 99 percent than it is to move from 94 percent to 89 percent, and more complex nonlinear models would be required to handle the special properties of percentages in such cases. For this reason, users working with multiple response categories should typically use the modal category as input for the Kalman procedure.

But there is a more fundamental limitation with which to contend. The Kalman smoothing algorithm used here presupposes one of two patterns of true opinion change. The first pattern obtains when $\gamma = 1$, in which case opinion wanders without returning to any particular equilibrium. The second occurs when $0 \leq \gamma < 1$ and opinion gravitates toward a single equilibrium. What about cases in which opinion moves from one equilibrium to another? In such cases, the intercept (α) in effect changes over time. For example, when we look at partisanship in the South using National Election Study (NES) data, we find that during the 1950s, approximately 60 percent of all NES respondents identified themselves as Democrats (see fig. 2). In 1994, this figure stood at 36.6 percent. Applying the Kalman smoothing estimator to the time span 1952–94, making no

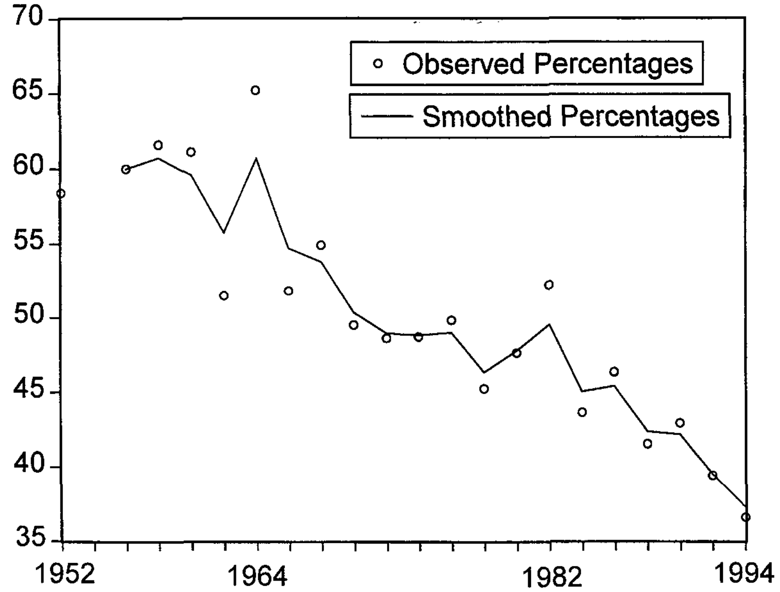

Figure 2. Observed and smoothed estimates of Democratic partisanship among southerners (assumes an autoregressive parameter of 1.0). Log-likelihood = −46.221; estimated disturbance variance = 5.710 (standard error = 3.424); sample mean = 50.305. Source: American National Election Studies Cumulative File, 1952–94, variables v4 (year), v301 (party identification), and v113 (South/non-South).

allowance for this shift in equilibria, produces estimates suggesting that partisanship is generally quite volatile, when in fact it was very stable over the period 1966–82. When we analyze the entire 1952–94 period assuming that $\gamma = 1$, we obtain a disturbance variance estimate of 5.71, implying that one should place a great deal of faith in the year-to-year shifts in partisanship gauged by these NES surveys. However, a comparable analysis from 1966–82 generates an estimate of 1.47. When we analyze the entire period, we are in effect asking the Kalman algorithm to smooth the data based on an oversimplified model of opinion change.

The problem of one-time shifts in equilibrium suggests that caution is warranted when applying the Kalman algorithm. If there is reason to believe that equilibria have changed, the analyst may wish to introduce a covariate into the analysis—perhaps a time trend or substantive predictor, such as evaluations of the economy. Adding regressors to allow for convention bounces or changing administrations involves a straightforward extension of equation (4) (see Hamilton 1994). Indeed, such an analysis can accommodate changes from one measurement regime to another, as when a survey house switches from personal interviewing to phone surveys or when it periodically refreshes its sampling frame. In this respect, analysis using the Kalman filter is entirely compatible with more traditional modes of public opinion analysis, in which the researcher begins by constructing a causal model to account for observed changes in opinion.

In summary, across a wide range of applications, the Kalman algorithm provides survey researchers with a single, systematic technique by which to guide analysis. The benefits of the Kalman procedure are fourfold. First, it enables tracking poll analysts to differentiate between random sampling error and true opinion change. Second, Kalman smoothing provides a means by which to accumulate information across surveys, greatly increasing the precision with which public opinion is gauged at any given point in time. Third, this technique provides a rigorous means by which to interpolate missing observations and calculate the uncertainty associated with these interpolations. Finally, the Kalman algorithm improves the accuracy with which public opinion may be forecasted. Since software implementing this technique is readily available, survey analysts are encouraged to employ it to make more efficient use of the data at their disposal.

References

Beck, Nathaniel 1990 ''Estimating Dynamic Models Using Kalman Filtering.'' In Political Analysis, vol. 1, ed. James A. Stimson, pp 121–56. Ann Arbor: University of Michigan Press.

Gerber, Alan, and Donald P Green. 1998. ''Rational Learning and Partisan Attitudes.'' *American Journal of Political Science* 42:794–818.

Green, Donald P., Bradley Palmquist, and Eric Schickler. 1998. ''Macropartisanship: A Replication and Critique.'' *American Political Science Review* 92:883–99.

Hamilton, James D. 1986. ''A Standard Error for the Estimated State Vector of a State-Space Model.'' *Journal of Econometrics* 33:387–97.

——— 1994. *Time Series Analysis.* Princeton, NJ: Princeton University Press.

Harvey, Andrew C. 1996. *Forecasting, Structural Time Series Models and the Kalman Filter.* New York: Cambridge University Press

Meinhold, Richard J. and Nozer D. Singpurwalla 1983. ''Understanding the Kalman Filter '' *American Statistician* 37:123–27.

Strang, Gilbert. 1986. *Introduction to Applied Mathematics.* Wellesley, MA: Wellesley-Cambridge Press.

U.S Census Bureau 1995. *Statistical Abstract of the United States: 1995.* Washington, DC: Department of Commerce.