

The Origin of Mood:
An Autobiographical Account

James A. Stimson
The University of North Carolina at Chapel Hill

Contents

1	Preparing a Seminar	4
1.1	The Seminar Discussion	6
2	The Futile Search	8
3	Black Lava, Blue Ocean, and the Recursive Solution	12
4	Digression on the Technology of 1988	14
5	Writing Code in Oregon	14
5.1	Mood and Thermostats	18
6	Going Public in Political Methodology	19
7	Public Opinion in America	21
8	After Mood	24

Abstract

This is a first person autobiographical account of the forces that came together in the creation of public policy mood. Unlike the standard academic fare which focuses on the product and leaves out the process of creation, this account is frankly personal. I write about what was going on in my life that interacted with scholarship. My goal is to document the facts for my own satisfaction.

This is some thoughts and memories of how mood came to be. It is not intended for publication or even distribution, I am writing it just for the sake of having a record. I intend to post it on a web site where it may or may not be read by stray visitors. I will read it myself which is enough, in retirement, to justify the effort.

1 Preparing a Seminar

I remember precisely where it all started. In the middle of Fall semester of 1987, my first semester at the University of Iowa, I was living in a house on Dearborn Street in Iowa City. Always in my mind it is the “bat house,” because I shared the house with those strange brownish-black winged critters. Owned by a professor of Sanskrit, the house had a strong odor of Indian spices that colors every memory to this day.

In my second floor study on an evening before the meeting of my graduate core seminar in American politics I was preparing for my next day’s class. The assigned reading for that class was John Kingdon’s book, *Agendas, Alternatives, and Public Policies*. I had assigned the book in part because I had not read it, but knew that it was important. It was for a week focused on public policy studies. That was a topic I felt obliged to cover but not an area of personal familiarity. I had never thought of myself as a public policy person.

My teaching style in seminars was to search for provocative questions that I could pose to the seminar that would produce a lively discussion. Those tended to be queries lacking a simple answer but ones that required a comprehensive understanding of the reading to answer. Too simple and the discussion would be extinguished with the correct answer. Too open-ended and it risked being boring.

Kingdon argues that one of the factors that moves policy agenda setting in Washington is what he called “the national mood.” It was, he wrote, a sense that the public underwent periods in which their summary attitudes toward policy issues drifted to the left or right. It was a mood because he believed that it was general, crossing many issues, not just one. Washington players, he wrote, had some level of

consensus on what the mood was at any given moment. Not a central theme of the book, mood was just one of several factors in agenda setting. And Kingdon's only evidence for it was that policy-makers mentioned it in interviews.

And then he wrote something that provoked my interest. Mood, he said, was not the kind of thing you could just measure as in an opinion survey. And there I had my question. How could it be true that national mood could have causal force in decision-making if it was so intangible that it could not be measured? Didn't that mean that a Washington consensus on the current mood was nothing but an exercise in mutual delusion?

The Memo Preparing for class did not exhaust my interest in mood. I went on to start thinking about the role such a concept played in national politics. Perhaps derived from Kingdon's idea that mood influenced policy considerations in Washington, I was taken with the idea that mood influenced policy-making over time, that for example, movements in mood produced change in public policy. I called it dynamic representation.

Representation, in the American context, has been almost wholly captured by the idea of geographical representation, of districts and their representatives in Congress. That is enshrined in the Constitution and surely it is sensible that elected politicians care most about voters who are legally entitled to end their careers. But it is also sensible that representatives see the issue longitudinally, that they first sense a change in national mood and then modify their behavior to be consistent with it. Dynamic representation then was a supplement to geographic representation, not a replacement for it. And it is smart to infer that what is happening in the nation quite probably heralds similar movements in the district.

I was much influenced by the ideas of rational expectations economics. One of those ideas is that professionals in any field would not wait to be hit over the head with new information, but would instead anticipate where that information was going before it got there. There is a premium in electoral success not only for being right, but for being right early. Elected politicians are in a very competitive game where

they are betting their careers on their perceptions of where the public is going. They cannot afford to wait until later to confirm their perceptions because later may be too late. So savvy professionals, whose careers are on the line would know public opinion and know where it is going and respond to it in real time. That means that the process of representation could occur without the necessity of elections as an information mechanism. That, the theory of dynamic representation, was my ultimate goal.

So here I was with a bivariate conception, changing mood producing changing policy outcomes. But I had no measure of either mood or policy outcomes, an $x \rightarrow y$ system in which I didn't know either x or y . That's why I wrote a memo. Maybe someday?

1.1 The Seminar Discussion

Sometimes provocative questions work. They start the discussion that was the goal of the seminar. Sometimes they flop. This was the former; it worked. I became an intellectual traffic cop, trying to keep the traffic in ideas in an orderly lane. Sometime during that hour or so of discussion a thought came to me. Maybe Kingdon was right in that you couldn't just pose a question in a survey because what mood was was many questions about preferred policy choices all running in the same direction.

Could you find evidence that preferences across different areas of policy were parallel in time? And then was it possible to observe correlated preferences across time and solve for the single common dimension that explained the parallelism? "Aha!" I thought. That is a factor analytic sort of question.

A design emerged from the seminar discussion. The things to be observed over time were preferences about policy choices. That is, they were the national percents who gave liberal or conservative responses to specific questions, should the federal government do more for healthcare, education, housing, racial equality, welfare, and on and on. In each case the liberal response was "yes," or "do more," or "spend more," or whatever. The conservative response was the opposite. Such data existed in the tradition of survey research. They

just needed to be collected and managed.

Mood, in the beginning, was play. The positive aspect, at least for me, is clear. Play is fun. Work is not. The downside of play in this case was my research budget. It was exactly \$0.00.¹ That meant that the project necessarily was based on free, publicly available, data. And that meant data from past and present. If one had the luxury of a budget, one could design a data collection with desirable properties. That would mean a large set of variables (survey questions) each of which was posed every year to a random national cross-section for a large number of years. The cost of such a study would have been astronomical and it was clear that no funding organization would step up to pay it. I never considered proposing such a study.

The publicly available data—thanks to The Roper Center—were both immense and immensely irregular. If one selects all the domestic public policy preference questions (and responses) that have ever been posed in more than one year, it is the case that most possible data is missing. That is, for a typical case about 80% of the preference questions that might have been asked were not, in fact, asked. A first thought was that maybe there exists a year in which all possible questions were posed so that one could score all other times relative to that common year. But no such luck. There isn't even a year in which half of all questions are posed.

And then we must confront the fact that we don't know what a particular survey response means. Because there is no science of question wording, we don't know what the response should be. If say 65% of respondents say "do more" in response to some program, that 65% score is meaningful only with reference to the exact words of the survey question. Since we can't know what response those words should have produced, we can't know whether that particular response is high, medium, or low.

Surely a literature existed on how to analyze such data.

¹In later years when I occupied an endowed chair I had a research fund. But it was only of a size to fund startup sorts of studies. Later the expanded study was funded by the National Science Foundation under the title, "Political Eras and Representation."

2 The Futile Search

I thought that this couldn't be a new problem. There must be a literature out there somewhere which adapted the mathematics of principal components analysis to the case where data were individual time points (rather than respondents) in a series that varied over time, not across people. Each policy preference question would form a time series where individual data were time points, initially years. My task was to find it and then just follow the instructions to a solution.

Well, I did a computer search. And a library search. And I found nothing. I talked to an Iowa colleague, Jae-On Kim, who was an expert in the family of techniques called factor analysis (of which principal components analysis is a specific case). No luck. I had studied factor analysis in graduate school. But all I had learned was how to compute a then obsolete method, centroid, by hand. In the pre-computer era a large part of the teaching of statistical methods centered on ease of hand computation as a desirable feature of alternative estimators. I was a product of that pre-computer era, receiving my Ph.D. in the same year that SPSS was released, which opened the door to computer solutions for social scientists.

I was a computer geek in graduate school. My main focus, beginning in my first semester at UNC, was writing a computer simulation of how members of Congress made voting decisions. See (Matthews & Stimson 1975). For statistical analysis we had a locally written package called TSAR (Tele-Storage and Retrieval) which ran on the University's big mainframe computer, an IBM model 360. Although too many years have passed since I touched it, I don't believe that it had a factor analysis capability. Or at least if it did, I never used it. (And I wouldn't have known how.)

Launched into an enthusiastic search from the seminar discussion. I collided with the reality that nobody seemed to share my interest in dimensional analysis over time. I thought that I would find the literature and then just follow the instructions. But I found nothing and had no instructions to follow. That was almost the end of mood. I conceived the problem as an exercise in dimensional analysis and I didn't know how to do dimensional analysis of time series. And then if anybody was going to tell me how to do the analysis, there was a

second problem for which there was no apparent solution.

I started collecting little scraps of policy preference data. Not a serious data collection, for which I had no reason to believe that any analysis was going to work, I just collected enough to play with. The second problem that emerged is that most of the potential cases were missing. That is, for particular survey questions, including the best ones, in most years they were not posed by any survey organization.

One can think of missing data issues as of two degrees of seriousness. The first, and most common in other applications, is that some cases are missing, but all variables still have ample overlap with other variables. The second degree problem is that not only are some cases missing, but the missingness also includes variables which are never posed in the same time period. For example domestic Communism, which was hot in the early 1950s, never overlaps with climate change issues, which occur in the survey record only in the last three decades. The stark implication is that the correlation between Communism attitudes and climate change does not exist. That means that a correlation matrix between all issues and all other issues is not possible. Numerous cells of that matrix do not exist because the two variables to be correlated never occur at the the same times.

The conclusion is unavoidable. Since such a correlation matrix is the starting point of principal components analysis, principal components analysis is not possible for such data. And there the search for mood ended. It was just play anyhow, never a serious research proposal. But it was a true cul de sac, fun while it lasted, but a project that had no future. It was set aside while I worked on my real research program of the time, *Issue Evolution* (Carmines & Stimson 1989). But the futility of the principal components solution turned out to be more a beginning than an end. For that “failure” freed me to think about other possible solutions to the estimation problem.

Abandonment of principal components also meant abandoning the thought of solution by off-the-shelf commercial software. That added a spark of excitement because I like to solve problems by writing computer code. Think about puzzles, crossword, sudoku, or jigsaw, for example. What comes after the solution? What comes after is the trash can. When those puzzles are solved the value of the solution is?

Nothing. In contrast, what comes after a piece of code is written is a tool, the solution to an important problem, one nobody else possesses. It is too much fun to call work. But it is problem solving and problem solving is work.

Scoring the Survey Response I needed each survey response to have a single summary score. Instead of the usual two or five or seven possible responses my first decision was to code everything in three categories, liberal responses, conservative responses, and neither. The last was exactly neutral responses, where they existed, plus the various missing data categories always present in survey responses.

I first decided to emulate economic indicators and score each national percentage response as 100 plus liberal percents minus conservative percents. Mike MacKuen, a future coauthor looking over my shoulder so to speak suggested instead $100 * (\text{liberal responses}) / (\text{liberal responses} + \text{conservative responses})$ because it had a more natural interpretation, percent liberal responses. 50 was the neutral point where the numbers of liberal and conservative responses were equal. Greater than 50 indicated relative liberalism while less than 50 indicated conservatism. He was right. It was more natural. And it remains the metric of mood forty some years later.

Solutions Entertained, Solutions Rejected The period that came after dismissal of principal components was a time when many different solutions to the problem of estimating mood were considered. Most were never implemented in code because their limitations and biases were known from the outset. A first idea was just taking the average of available indicators in a particular year. That would have been viable if the missing data problem had been random. But it was not. Polling firms employed systematic criteria, such as newsworthiness, when deciding what questions to field. Since the average value for question series covers almost the whole possible range from zero to 100, the computed value for each year would vary inversely to what items were missing.

Another thought was to center each available series on zero by subtracting out its mean. This gives us an expected value for each

case of zero, which seems to remedy the missing values problem. The immediate problem with this approach was that we don't really know the mean. We have the illusion that we know the mean. After all we can just add up the values and divide by N , the solution we all learned in the early grades of school. The problem is that this computed mean, conventionally called \bar{x} , is only an estimator for μ , the true mean if we can assume that we have a random sample of cases. Alas we cannot make that assumption, because the available cases, just like the missing cases, are not random. Systematic decision-making where time on a survey questionnaire is both expensive and scarce defeats that.

Dummy Variable Regression A proposed use of regression analysis almost solved the estimation problem.² The proposal was to regress the liberalism index on a massive set of dummy variables, one for each of about 250 distinct survey questions to control question effects and a large set (about 35 in 1988 when proposed, 70 now) of year dummy variables.³ The question effect variables were just a control for extraneous variance. They were to be estimated, but not used for anything. The time dummy variables became estimates of mood.

The estimation, which required inverting a massive 285 by 285 matrix, the $[x'x]$ in $\beta = x'y[x'x]^{-1}$ would have been impossible with available desktop computers through the middle 1980s because the matrix to be inverted was over 80K*4 bytes and the PC equivalent computers of the time were limited by a 64K data segment. When implemented in 1988 reasonably priced clones no longer had that limitation and the estimation became possible. But would it work? I didn't know for sure until I had written a linear regression routine in code (compiled BASIC). It seemed possible that it might refuse to invert that giant matrix or that it might undertake the task but take multiple days to do so. It worked. And it did so in minutes, not hours or days or worse.

But it was still imperfect. The reason is that regression is an ana-

²The proposed specification came from Chris Achen.

³My conception of mood at the time was as an annual time series. Subsequently it has been generalized across a range of aggregation intervals, ranging from daily to multi-year intervals.

lytic technique, not a measurement method. The important difference is that in the measurement realm there is a convention that the appropriate way to form a linear composite of multiple variables is to weight each variable by its estimated validity as an indicator of the underlying concept. This regression specification instead implicitly assumes that all variables are perfectly valid indicators of the concept in question. And I knew that the survey indicators did not just vary as indicators of the left-right continuum, some did not tap that domain at all, i.e., they had zero validity. Questions such as national park spending or space travel just did not have any association with liberal vs. conservative at all. So the regression setup would produce a measure of mood—and a pretty decent one at that. But I knew that the perfect validity assumption was less than ideal. Had it not come at about the same time as a solution that worked, I probably would have adopted it.

If I had adopted the regression approach it would have become necessary to pre-screen the survey variables to make sure that they tapped the left-right dimension of mood. This is the same step that is necessary for recent item response theory estimates of mood See (McGann 2013) and (Caughey & Warshaw 2015).

3 Black Lava, Blue Ocean, and the Recursive Solution

My life at the time was not all work.

Dianne Stimson, my wife, is by profession a hospital executive. In 1987 her career took her to Corvallis, Oregon where she was CEO of Good Samaritan Hospital, a regional medical center. Once a year or so hospital executives hosted a retreat for selected Board members and members of the medical staff. Always in an attractive location, Hawaii was a favorite for West Coast institutions. (The gloomy Pacific Northwest winter climate put a premium on hours of sunlight.) The retreat was an exercise in team building for hospital executives, board members, and medical staff. Spouses traveled with team members. A mix of work and leisure, I enjoyed my role as an escort to the other

spouses during work periods.

On one such occasion I had time to myself and retreated to a “beach” of black lava rocks. I just sat there there looking out to the blue Pacific and thinking, as I had so often over the months of 1987 and 1988, of how to go about extracting the signal of mood from the noise of response to survey questions. My experience as a social scientist was that insights were usually the product of dull grinding hard work, not flashes of inspiration. But on this occasion the setting was so peaceful that when my mind wandered to estimating mood, as it had hundreds of times before, I saw clearly the idea that bedeviled me. I was stuck on the idea that there was no analytic solution, no formula, for estimating mood,

I played in my mind with the arithmetic of estimating the last two points of the time series. The very last point, $mood_T$, I could assign an arbitrary metric, for which I chose 100. What then of $mood_{t-1}$? Well for the set of variables that had values for both T and T-1, one could observe the ratio of liberalism scores between the two in the available indicators. That’s just arithmetic. Then one could assign a value to $mood_{t-1}$ by multiplying that ratio times the value of the series at T (which is 100). So $Mood_{t-1} = 100 * ratio$. So why is that a big deal? It is a big deal because once we have a numeric value for $mood_{t-1}$ we can do the same thing to estimate $mood_{t-2}$. And once we have a value for $mood_{t-2}$, we can project a value for $mood_{t-3}$. And in a nutshell, by this one step at a time, which we will call recursion, we can estimate *all* values of mood.

It was possible to estimate the relationship of a dyad of time points and then connect those to all earlier and later times through recursion. And that meant that mood could be estimated. What was needed then was to turn that thought at the beach into code, and there was a solution to estimating mood.

And that was it, the foundation of what became the dyad ratios algorithm.⁴ The rest of the stay and the travel back to Oregon are a blur in my mind. But I knew that it was possible to estimate mood and my mind was totally focused on how.

⁴For my best explication of the logic of dyad ratios see (Stimson 2018).

4 Digression on the Technology of 1988

At that time I owned a Toshiba laptop running MS-DOS. It lacked a hard disk—way too expensive in that era—but had instead what was called static RAM. That is random access memory which retains its contents until overwritten by something else. In contrast to the better known dynamic RAM, it doesn't lose its contents when the power is off. I no longer remember the memory size, but it would have been expressed in thousands of bytes, not millions. One could store perhaps 20-30 pages of text (plain text with no formatting) at a time. For “permanent” storage there were floppy disks.

Everything electronic from that era is lost. The reason is that the technology of disk storage changed about every five years or so and so replacing a computer with a newer and better model almost always meant that the floppy disks from an earlier machine were unreadable with the current machine.⁵ Hard disks of the era were too small to store anything that was not in current use. The Cloud was several decades in the future. And so most of what was produced in the time of floppy disks was eventually lost forever. The exception was that truly permanent medium, publication on paper. And so my (first) book about mood, *Public Opinion in America: Moods, Cycles, and Swings* survives in paper but all the rest of the work product, data, figures, correspondence, memos, outlines, snippets of computer code, etc. is lost. Printing on paper remains fully intact. Memory remains, but decayed by time. This document is a product of that decayed memory.

5 Writing Code in Oregon

Back in Oregon I sat down to turn my lava beach thoughts into a subroutine that would implement the recursion idea and produce mood as the output. It was not a difficult task. I wrote in the BASIC language using an IBM compiler that I already owned. BASIC was

⁵I still own a dozen or so floppy disks of the most recent technology and can't bear to just discard them. But it has been over a decade since I have owned a computer that could read them.

designed in the first instance as a learning language. It is generally easier to use than alternatives such as `c` or `c++`. The virtue of easy is that you can focus on problem solution and not get bogged down in language details. Because of its origin, it is disdained by professional programmers.⁶

I could write that I was the right person in the right place at the right time to do this task. But the truth is closer to I was the only person positioned to estimate mood at that time and place. Mood was a topic in public opinion and I was a career-long public opinion analyst. It was a time series and for essentially accidental reasons I was a time series analyst, at the time one of a handful in political science. And I wrote computer code—enjoyed writing computer code—again one of a handful. So the project I had on my hands required the very skill set that described me. Add to that the feverish enthusiasm I felt for the challenging puzzle of mood and all the pieces fit perfectly into place.

I had an early collection of policy preference data, some from published sources such as (Niemi, Mueller & Smith 1989), some from a then unpublished draft of what became (Page & Shapiro 1992), some (Trendex data) from Bob Shapiro, and many from the Roper Center, then a clunky mainframe database system (Spires) from the University of Connecticut. My criteria were that they were domestic policy preference questions administered to national samples in the United States.⁷ Questions that made reference to a person or party were excluded as reflecting considerations other than policy.

When I wasn't writing code I was playing around with the preferences raw data. Using a spreadsheet for graphics and computation, I used the spending and priorities questions from the General Social Survey (GSS) with some of the methods I had thought about and discarded. Those data had missing years, when the GSS study was

⁶Because of a confusion with interpreted BASIC, widely used by early hobbyist computer users, it is falsely believed to be slower than other languages. Bench testing the output of optimizing compilers on the same task produces results in elapsed time measures that are indistinguishable from other languages.

⁷In the early years I paid no attention to content except distinguishing foreign from domestic. In later years when culture war issues became numerically dominant, I switched to selection of size and scope of government issues.

not fielded, but no missing data of the typical kind. Every time the survey was fielded every question in the long battery was asked. I tried several manipulations, but one in particular was just to compute the yearly averages across all questions in the battery. Since there are no missing data, that was a valid way to capture the longitudinal pattern. So I had a pretty good idea what to expect when I was able to apply the Dyad Ratios algorithm to a full collection of survey data.

So some time in late 1989 I was able to estimate mood for the first time. It was similar to Figure 1 below.⁸ The strong similarity is validated by Figure 4.1a of (Stimson 1991). See the line labeled “Unsmoothed.” This is public policy mood from 1956 to 1989. The pattern, which I had earlier seen in raw data, was a fairly strong replication of the cultural assumption that the 1960s were a time of liberalism and the 1980s equally a time of conservatism. The intriguing difference is that movements in mood preceded public recognition of shifting preferences. The peak liberalism was in 1961 and not the later colorful decade and peak conservatism was in 1980, the year of Ronald Reagan’s surprisingly strong election. The 1950s, represented by only four years, seemed to anticipate the 1960s liberalism. Later when both data and estimator were improved, the conservatism of the early 1950s emerged.

All in all it was exciting. It seemed clear that this early 34 year series was strongly associated with presidential election outcomes. And that made it plain that all that computation uncovered something somehow connected to the essence of American politics. Though quiet and personal, it was a small Eureka moment.

The backward recursion method that I have described in these pages was an unbiased estimator. But it still had one problem. Inherent in the method is that later cases had more weight in the solution than did earlier ones. Case T affected all earlier cases. Case T-1 affected 1 through T-2, Case T-2 affected cases 1 through T-3 and so forth. Ultimately case 1 affected nothing.

What about reversing the process, starting with an arbitrary value for case 1 and then connecting that to case 2 and connecting 2 and 3 and so forth. So this is forward recursion. It has the opposite

⁸It is only “similar” because the original has been lost for decades.




Figure 1: Public Policy Mood in the United States as first Estimated in 1989: Percent Liberal by Year

weighting issue. Case 1 affects everything that follows. Case 2 affects 3 through T etc.,.Finally case T affects nothing. Now we have two equally good estimates of mood, forward and backward. Although generally highly correlated, the two are different because the case weighting is different. What then to do? You could chose one of the two equally good estimates. But better you can chose the average of both. That has the dual advantages of (1) using everything that you know and (2) of equalizing the weighting of all cases. That became the chosen solution.

Publication? From the first I had no thought of conventional publication. It was a conceptual development and practical measurement exercise, the kind of thing not favored by general political science journals. I never tested that judgement, so I can't know whether or not it was true. But over a long career I was never shy about sending my work to journals for review. The then brand new journal *Political Analysis* was open to publishing pure methodological developments, but its pages were not available to me because I was its editor at the

time. The last thing the fledgling new journal needed was a dubiously peer reviewed article by its editor. The journal was competition for my time—alas its principal effect on mood. It was not an outlet for my work.

After it was in the public domain papers explaining mood and papers using mood to explain other things became grist for the journal mill both for myself and for others. It was just that initial development that was never submitted for journal publication. It was somewhat unfortunate that scholars were employing a measure that had no public development.⁹

5.1 Mood and Thermostats

Early on I had an audience of one. Chris Wlezien was working on a dissertation (at the University of Iowa) on budgeting games. I was a member of his dissertation committee and, more important, a friend. We went for coffee on at least a weekly basis. In the beginning we argued about the dissertation. But Chris had developed the idea that public opinion was thermostatic. The idea is that public reaction was relative to current policy, not absolute. That makes it possible for a part time electorate to sense “too much” or “too little” without having to know what it wanted. Just as the home thermostat senses too hot and too cold and makes a binary adjustment (turn on or off the heat, turn on or off the AC), all the public needed to know to respond to policy change is whether it was too much or too little. The key insight of the thermostatic model is that it makes democracy possible for an amateur and part time electorate that knows little of what government is doing.

Chris shared my enthusiasm for mood because it was an obvious candidate to test the thermostatic model. Without it, the model was an idea in search of a demonstration. With such a measure in hand the thermostat could be demonstrated.¹⁰ Chris’s enthusiasm for the idea of mood did much to stoke my enthusiasm, particularly at a time when department chair duties pulled me away from my research program.

⁹That has finally been remedied in (Stimson 2018).

¹⁰Chris’s work eventually appeared in (Wlezien 1995).

6 Going Public in Political Methodology

Since the first meeting in Ann Arbor in the summer of 1984 I had always attended the Political Methodology summer meetings.¹¹ Those meetings were the best venue for a critical response to political science papers that existed. So I chose the Political Methodology meetings to make my first public presentation on Mood. I believe it was the summer of 1989, but my memory is cloudy and I might be off by a year.

Left-right ideology is second nature to all political scientists. I had reason to think that an over time measure of where to place the American electorate in ideological space would draw interest from a critical audience. And when I presented mood to the summer gathering of the Political Methodology Society it did draw interest. As one of more than 300 such talks I have given over a fifty year career, I do not have many specific memories. But the methods group was famous for digging into the mathematical underbelly of political research. I had once compared the act of putting an equation in a paper as like feeding bloody red meat to lions. In both cases a violent response is to be expected. The lions roared and ate, as expected. The best one can hope for with such an audience is to survive. I survived. In fact my memories, though scant, are positive.

From the discussion that followed I remember two particular critiques that would leave their mark on the estimation of mood. One was from Gary King. He looked at the graph of mood over time and said that there was too much year to year variation in the early 1950s estimates. I was never quite sure what standard he was using to assess how much variation was to be expected, but he spoke knowingly and I took the comment seriously.

The second comment, apparently unrelated, came from Neal Beck who noted that survey data always had (random) sampling error, the effect of drawing one sample of a particular size instead of another

¹¹I was of an earlier generation than most of the Political Methodology participants and less well trained in hard core methodology than most. The group nonetheless became centrally important to the development of my career.

of the same size. The conventional time series cure for the net effect of such errors is some sort of smoothing. The idea is that nature is smoother than our survey sample measures of it and so whatever values one obtains ought to be subject to smoothing to get closer to what would be expected in a world without sampling error.

Until I got home and had time to reflect I didn't realize that the King and Beck critiques were really the same thing in different language. Excess variability *was* a problem, particularly in a period when data were thin and natural smoothing did not exist. And some sort of smoothing was the solution.

The decade of the 1950s was a time when all, or virtually all, surveys were done in person at the respondent's home. That is vastly more costly than telephone surveys and consequently the Gallup organization and Opinion Research Corporation were the only commercial survey firms doing policy preference studies. That work was supplemented by the Michigan election studies of 1956 and 1958 and by occasional studies from the University of Chicago's National Opinion Research Center. But that is the whole record for a decade. The norm for later decades is at least ten times greater. So the problem of thin data for estimating mood is a 1950s problem only.

One solution to the 1950s issue was to just not try to include that decade. That was unacceptable because we really needed to know what was going on in that quiet decade. The alternative was to implement some kind of smoothing to tame the excess variability. That is what I chose.

When we observe a series of survey measures we are observing nature, the true underlying response, and sampling error intermixed. Since sampling error in a series is statistical "noise," it is a reasonable surmise that unobserved nature is probably smoother than the nature plus noise that we observe. Thus it makes sense to assume that a smoothed version of a survey series is probably closer to unobserved nature than is the actual data series. That is the rationale for introducing smoothing of the forward and backward recursion series (before taking their average).

The exponential smoothing model employed is:

$$y_t = \alpha x_t + (1 - \alpha)x_{t-1} \quad (1)$$

where x is the original (noisy) data, y is the smoothed version and the α parameter is estimated from the x data series.

And that was it. Smoothing was the final piece of the Dyad Ratios algorithm, a subroutine implemented in various languages which survives unchanged from 1989 to the present. There was much code writing to come, but it was all directed toward more user friendly data management issues.

Thank you Gary and Neal. Having critical colleagues is a joy.

7 Public Opinion in America

I like to write. I especially like to write books. Books give their author a freedom, subject to the constraint that you have to keep the reader interested, to do pretty much whatever you want. And I needed to go public with mood. Conference papers, such as what I had written for Political Methodology, are for a very high quality audience, but a very small one. They are not totally invisible, but pretty close to that. I wanted to introduce mood to a broad wave of of working political scientists, beyond the discipline to fellow travelers in public opinion, and even inside the beltway to the people who practice politics. That last had something of a pie in the sky character.

Around the time of the Political Methodology meeting I was approached by Larry Dodd, series editor for the Westview Press *Transforming American Politics* series to do a book in the series. He introduced me to Jennifer Knerr, then Westview's editor. I had warm working relations with both. I agreed.

“We don't care about the aggregate level of opinion, we care about why one voter has a different opinion than another.” That idea, drawn from what we would now call political psychology, is what I had been taught as an undergraduate. And I found myself, more or less by accident, asserting the opposite. I found myself saying “I don't care

about why one voter (or citizen) differs from another, but for what the public, taken as a whole, does and thinks. I found myself standing on the opposite side of the fence from the scholar I had been trained to be.

I didn't need to rebel from public opinion orthodoxy. I could have maintained the primacy of individual psychology and noted that my use of aggregate opinion was just a measurement technique. But in truth I had never bought into the idea that explaining individual attitudes and behavior was the chief goal of opinion studies. I quote from my introductory prose.

What we found is that citizens, taken one at a time and in the norm did not seem to be competent by the standards of democratic theory. . . . They could not act as individuals in the prescribed manner. And that may well be true. But should we take them *one at a time*? Is that the nature of politics? Or is public life instead the life of the herd? We carry normative blinders about this issue. Terms like *herd* or *groupthink* express them. We like individuals, distrust aggregations. But one of the plainer facts of everyday life is that individuals do not function as individuals. They are enmeshed in a social environment. They interact. They give, receive, borrow, and steal things from one another. Some of those things are ideas. Some ideas are about politics.

We corner individuals in their living rooms or on their telephones, all by themselves, for a survey. And we come to think of them as individuals, "the respondent." And "the respondent" is an abstraction of a real person, normally not by himself or herself, whose ideas come only partly—and maybe in pretty small measure—from self. We forget that we are interviewing not a self-contained individual but a spokesperson for the herd. . . .

If our topic is people, then this abstraction as individuals is relatively harmless. If we want to know how people think, then perhaps an order and harmony not of their own making may be omitted. But if the topic is politics, then it matters a great deal if the aggregate opinions are more orderly and meaningful than individual ones. For it is the

aggregate that matters in politics. To think otherwise is to confuse our normative preference for individualism with the reality of social order.(Stimson, 1991, page 2)

The book was *Public Opinion in America: Moods, Cycles, and Swings*. In its preface I described it as an unfinished essay. “Essay” was the style of presentation. The “unfinished” part alerted the reader to the fact that I had a grand design of telling the story of public opinion moving government policy and I had, as yet, no measure or analysis of policy. Forging the link between the two I saw as a two to three year process and I didn’t want to wait for completion to tell the public opinion story. The finish was to come in an article entitled “Dynamic Representation” (Stimson, MacKuen & Erikson 1995) and a similar analysis in the coauthored book *The Macro Polity* (Erikson, MacKuen & Stimson 2002).

Was there a pattern to be found in mood? For an example of what was available to see in 1989 when I started writing just take the data of Figure 1 and rearrange it a little. Instead of actual year as in the figure, group mood by its relationship to the four or eight year presidential election cycle. Assign the initial election year to 0, the first year in office 1, and so forth. Then subtract out the mean of mood for the period so that positive deviations indicate liberalism and negative deviations indicate conservatism. Then multiply those scores by minus one if the year is a time of a Republican presidential win or Republican control of the White House.

First, if there is no relationship between mood and the electoral cycle we would expect to see small random deviations around zero with roughly equal numbers of positive and negative values. Instead all values are positive, indicating that mood on average tends to support the policy stance of the White House party. Leaving out the somewhat strange four year term of Jimmy Carter the pattern is perfect. All presidential election years deviate in the direction of the winning party. And then the remaining seven years are a perfect pattern of exponential decay, with maximum movement in year one and a steady tail-off every year thereafter.

And the over time pattern is not a bit random. Instead it tracks a positive impulse at the outset of the term which regularly decays

as the term progresses. That is not quite a test of the Wlezien thermometer idea because that would necessitate an actual measure of policy change. What we have here instead is consistent with an assumption that parties follow their normal tendencies in policy making, Democrats liberal and Republicans conservative. Then just by knowing who is in office the electorate can infer the direction of change and respond thermostatically.¹²

What Moves Mood? Forget about “the bully pulpit.” Forget about alleged “great communicators.” What moves mood is reaction to the party in power. And that reaction is always and wholly negative. And this relentless negativity is true for both parties. When Democrats occupy the White House mood becomes more conservative, year by year. When Republicans rule mood becomes more liberal, year by year. These patterns are very strong and very predictable. And recall that mood is purely a matter of policy preferences. The questions that compose it make no reference to party or partisanship. Equally they do not refer to the approval or disapproval of incumbents. So this observed negative response means exactly that policy preferences, in general, move away from the position of the party in power.¹³

A great number of trees are felled to make all the paper on which pundits assert that true opinion change in America is toward liberalism or, with wholly equal falsity, toward conservatism. The behavior of mood suggests the true answer. Opinion does not trend in either direction; it cycles contrary to party control of government.¹⁴

8 After Mood

Let me be clear about my motives. Mood was my idea and my creation (with due credit to John Kingdon for its origin). I wanted it to be used by others. It is possible to create an idea or measure or estimator that is esteemed by readers and approved by the always critical journal

¹²This response to inference was named the “implied thermostat” in (Atkinson et al. 2021).

¹³See (Durr 1993) for a development of this idea.

¹⁴A small number of exceptions are noted in (Atkinson et al. 2021)

reviewers and yet die an orphan because it is so complex or difficult to replicate that readers are tempted to read and then put it down. If something is inherently difficult, it is possible to die by publication. To truly have an impact you want readers to use, replicate, and even criticize a development. I wanted mood to have an impact.

While one might be tempted like a child to keep a toy for your own exclusive play, a private toy will never have a big impact. So a first decision was to make it public, to put the measure on a web site where anyone could access it. And then a follow-up decision is to keep it current, so it will have greater value to the scholarly community. I did a first update in 1990. And here I am, approaching 2024 and still planning to update mood one more time the better part of a lifetime later.

And I have committed myself all these years to produce software that permits scholars of all kinds to estimate “moods” on topics far beyond anything I ever conceived. Three pieces of software, wcalc.exe for Windows, mcalc for Mac, and extract for the R platform, (all available on the web) bring the ability to estimate mood from their own datasets for users all over the world. I have some knowledge of how wide usage is from the inquiries I get and from citations in published work. The thousands of hours spent of those projects don’t add lines to my c.v., but they do expand the impact of mood. I have participated in bringing mood estimation to Great Britain and France. Others, working beyond my knowledge, have similarly done the same for many other countries. Some appear in languages I cannot read. Those extensions are now so extensive that I cannot monitor them. Are these efforts to increase the impact of mood worthwhile? Yes.

December 18, 2023

jas

References

- Atkinson, Mary Layton, K. Elizabeth Coggins, James A. Stimson & Frank R. Baumgartner. 2021. The Dynamics of Public Opinion. New York: Cambridge University Press: Elements.
- Carmines, Edward G. & James A. Stimson. 1989. Issue Evolution: Race and the Transformation of American Politics. Princeton, NJ: Princeton University Press.
- Caughey, Devin & Christopher Warshaw. 2015. “Dynamic Estimation of Latent Opinion Using a Hierarchical Group-Level IRT Model.” Political Analysis 23(2):197–211.
- Durr, Robert H. 1993. “What Moves Policy Sentiment?” American Political Science Review 87:158–170.
- Erikson, Robert S., Michael B. MacKuen & James A. Stimson. 2002. The Macro Polity. New York: Cambridge University Press.
- Kingdon, John W. 1984. Agendas, Alternatives, and Public Policies. Boston: Little-Brown.
- Matthews, Donald R. & James A. Stimson. 1975. Yeas and Nays: Normal Decision-Making in the U.S. House of Representatives. New York: Wiley InterScience.
- McGann, Anthony J. 2013. “Estimating the Political Center from Aggregate Data: An Item Response Theory Alternative to the Stimson Dyad Ratios Algorithm.” Political Analysis 22(1):115–129.
- Niemi, Richard G., John Mueller & Tom W. Smith. 1989. Trends in Public Opinion: A Compendium of Survey Data. Westport, Connecticut: Greenwood Press.
- Page, Benjamin I. & Robert Y. Shapiro. 1992. The Rational Public: Fifty Years of Trends in Americans’ Policy Preferences. Chicago: University of Chicago Press.
- Stimson, James A. 1991. Public Opinion in America: Moods, Cycles, and Swings. Boulder, CO: Westview Press, 1991.
- Stimson, James A. 2018. “The Dyad Ratios Algorithm for Estimating Latent Public Opinion: Estimation, Testing, and Comparison to Other Approaches.” Bulletin of Sociological Methodology Bulletin de Méthodologie Sociologique 137-138:201–218.

- Stimson, James A., Michael B. MacKuen & Robert S. Erikson. 1995. "Dynamic Representation." American Political Science Review 89:543–565.
- Wlezien, Christopher. 1995. "The Public as Thermostat: Dynamics of Preferences for Spending." American Journal of Political Science 39(4):981–1000.