

THE SIMULMATICS PROJECT

BY ITHIEL DE SOLA POOL
AND ROBERT ABELSON

The development of electronic computers and of mathematical game theory has greatly stimulated the simulation of human behavior as a means of exploring the implications of theory and extending the analysis of great masses of data provided by surveys and similar sources. Here is the narrative of a fascinating simulation conducted during the presidential election campaign of 1960. It is not only a further step in the automating of opinion research but also a massive test of some theories of opinion formation.

Ithiel de Sola Pool is Professor of Political Science and Director of the International Communication Program of the Center for International Studies at Massachusetts Institute of Technology. Robert Abelson is Professor of Psychology at Yale University.

THIS IS the first report on a program of research conducted for the Democratic Party during the 1960 campaign. The research used a new technique for processing poll data and included computer simulation of likely voter behavior. The immediate goal of the project was to estimate rapidly, during the campaign, the probable impact upon the public, and upon small strategically important groups within the public, of different issues which might arise or which might be used by the candidates.

THE DATA

This study is a "secondary analysis" of old poll results. Students of public opinion are becoming aware that the growing backlog of earlier polls provides a powerful tool to aid in the interpretation of new poll results. Polling has now been routine for three decades, but poll archives are just beginning to be assembled. The main one is the Roper Public Opinion Research Center in Williamstown, the existence of which made feasible the project here described.¹

The first step in the project was to identify in that archive all polls anticipating the elections of 1952, 1954, 1956, and 1958. (Pre-election polls on the 1960 contest were added later when they became available.) We selected those polls which contained standard identification data on region, city size, sex, race, socio-economic status, party, and religion, the last being the item most often missing. Further, we restricted our attention to those polls which asked about vote intention and also about a substantial number of pre-selected issues such as civil rights,

¹ We wish to express our gratitude to that Center, as well as to the MIT Computation Center, and to the men who originally assembled the data, especially George Gallup and Elmo Roper.

foreign affairs, and social legislation. From 1952 to 1958 we found fifty usable surveys covering 85,000 respondents. Sixteen polls anticipating the 1960 elections were added to this number. The sixty-six surveys represented a total of well over 100,000 interviews.

PROCESSING THE DATA

To handle such massive data required substantial innovations in analytic procedures. In essence, the data were reduced to a 480-by-52 matrix. The number 480 represented voter types, each voter type being defined by *socio-economic* characteristics. A single voter type might be "Eastern, metropolitan, lower-income, white, Catholic, female Democrats." Another might be, "Border state, rural, upper-income, white, Protestant, male Independents." Certain types with small numbers of respondents were reconsolidated, yielding the total of 480 types actually used.

The number 52 represented what we called in our private jargon "issue clusters." Most of these were political issues, such as foreign aid, attitudes toward the United Nations, and McCarthyism. Other so-called "issue clusters" included such familiar indicators of public opinion as "Which party is better for people like you?" vote intention, and nonvoting. In sum, the issue clusters were political characteristics on which the voter type would have a distribution.

One can picture the 480-by-52 matrix as containing four numbers in each cell. The first number stated the total number of persons of that voter type asked about that particular item of information. The other three numbers trichotomized those respondents into the percentages pro, anti, and undecided or confused on the issue.

We assembled such a matrix for each biennial election separately and also a consolidated matrix for all elections together. Thus, it was possible by comparison of the separate matrices to examine trends.

The reduction of the raw data to this matrix form was an arduous task. The first step was to identify in each survey those questions which seemed to bear on any of the fifty-two issue clusters we had listed as relevant to the campaign. One such cluster was attitude toward domestic communism or, as we called it for shorthand, McCarthyism. Over the past decade many questions have been asked on this and related matters in many different polls. One survey might ask, "Are you in favor of permitting a Communist to teach in the school system?" Another would ask, "What do you think of Senator McCarthy?" Another would ask, "Do you think McCarthy has done more good or harm?" The problem was to determine which questions tapped essentially the same attitude, domestic anticommunism. The decision was made by a two-step process. First, questions were grouped together a

priori on the basis of intuitive judgment, and then this grouping was empirically tested.

The empirical test was conducted as follows: Replies to each questions were separately trichotomized. Typically, the replies had previously been coded in up to thirteen categories. Where more than three replies had been coded, the codes had to be regrouped. On the McCarthyism issue, replies were classified as McCarthyite, anti-McCarthyite, and indeterminate. A reply opposing retention of a Communist in the school system would be classified as McCarthyite. In the case of such a question as "How well do you like McCarthy?" for which a scale had originally been used, cutting points had to be set depending on the distribution.

For each pair of questions in the presumed cluster we then correlated the percentage "pro," and separately the percentage "anti," across voter types yielding two correlation matrices. (The voter types for this operation were 15, a reconsolidation of the 480. Since this operation dealt with percentages on questions from single surveys, consolidation was essential to obtain base numbers in each voter type large enough so that the percentages being correlated would be reasonably stable.) Only those questions which showed high correlations with each other were retained in a cluster. Thus, our assumption that a question about Communist teachers in the schools could be treated as equivalent to a question about McCarthy was subject to empirical validation.

In many instances questions which *a priori* seemed alike had to be discarded from the clusters. Some clusters had to be broken up into two or more. Indeed, in the particular example we have been using here, it turned out that replies to the identically worded question "How well do you like McCarthy" ceased tapping the same attitudes the minute the Senate censured him. Clusters thus represented questions which could be regarded as in some sense equivalent, both on the grounds of political common sense and on the grounds of empirical correlation.²

² We should qualify. What has been described is what we started out to do and what we did for most issue clusters. In the end, however, we were forced to compromise on certain foreign-policy clusters. This in itself is an interesting finding. On almost all domestic questions, primarily because they were party-linked or left-right linked, it was possible to validate empirically the equivalence of questions which *a priori* seemed alike. On certain foreign-policy issues this was quite impossible. The political scientist looking at a half-dozen questions about foreign aid or about the UN might conclude that they all should reflect a common underlying attitude toward that matter. However, empirically, in many instances the distribution of replies was highly sensitive to conjunctural influences or shades in wording of the question. Rather than completely abandon the hope of doing any analysis of foreign-policy issues in the campaign, we retained some clusters which failed to meet the correlational test, labeling them *a priori* clusters, not sure of what we would do with them (in fact we did very little), but feeling it better to retain them on the computer tape than to discard the data from the start.

It should be emphasized that empirical correlation was not enough. Such a question as "Which party is better for people like you?" and a question about the image of Adlai Stevenson would correlate strongly because they were both party-linked. However, they were not included in a single issue cluster unless they also seemed politically equivalent.

The final step in the preliminary data processing—the step which gave us our matrices—was to take all cards in any one of the 480 voter types for a particular biennial period and tabulate for each issue cluster the number of replies pro, con, and indeterminate, and the number of cards on which such replies appeared. That last number varied for each cluster since some questions (e.g. turnout) were asked on virtually every survey we used, while other questions were asked only occasionally.

PURPOSES OF THE METHOD

The reader may wonder what purposes were served by reorganizing the data into the standard format just described. That handling of the data lent itself to three main uses: (1) A "data bank" was available from which one might draw the answer to any one of a vast number of questions at a moment's notice. (2) The consolidation of separate surveys made available adequate data on small, yet politically significant, subsegments in the population. For example, we wrote a report on Northern Negro voters based upon 4,050 interviews, including 418 with middle-class Negroes. The typical national sample survey contains no more than 100 interviews with Northern Negroes, a number clearly inadequate for refined analysis. (3) The data format and its transfer to high-speed tape facilitated its use in computer simulation of the effects of hypothetical campaign strategies. This aspect of the project is the most novel and is the one to which we shall return later in this article.

THE HISTORY OF THE PROJECT

Before we illustrate those uses of the data, let us detour to examine the history of the project: the fact that it was sponsored and actually used by a partisan group makes the story of its management of some interest to students of public opinion research.

The project was initiated in the early months of 1959 by William McPhee and the authors. Our plan for computer simulation (on a different version of which McPhee had already been working)^{*} was presented to Mr. Edward Greenfield, a New York businessman actively engaged in Democratic politics. Through his intervention, a group of New York reform Democrats who had taken major responsibility for

^{*} William McPhee, *A Model for Analyzing Macro-dynamics in Voting Systems*, Columbia University, Bureau of Applied Social Research, undated.

raising money for the Democratic Advisory Council became interested.⁴ Before this group of private individuals was willing to secure funds, however, they wanted to be sure that the results were likely to be valid and useful. In May of 1959 the project was discussed in Washington at a meeting attended by Mr. Charles Tyroler, Executive Secretary of the Democratic Advisory Council; the members of the Council executive committee; Paul Butler, Chairman of the Democratic National Committee; several other officials of that Committee; Mr. Neil Staebler, Michigan State Chairman; and a number of social science consultants, including Samuel Eldersveld, Morris Janowitz, and Robert Lane. This group was interested but reserved. It was suggested that the project should be supported for four months initially and at the end of this period a further review should be made.

The Williamstown Public Opinion Research Center agreed to permit the use of polls in their archives on two conditions: First, all basic data tabulated by Simulmatics from their cards were to be made available to the Center so the Republican Party would have an equal opportunity to use such data if they wanted them. We provided a print-out of the data on the computer tape, but not, of course, the programs for simulation nor supplementary data obtained from other sources (e.g. the census) and used in our system. Second, and demanded by both the Roper Public Opinion Research Center and the social scientists engaged in the study, all results could be published for scientific purposes after the election. This article is part of our program to meet that condition.

Given the green light to carry out the project, the principals organized themselves as The Simulmatics Corporation, for although the objective of the project constituted scientific research, it was clear that universities would not and should not accept financing from politically motivated sources or permit a university project to play an active role in supplying campaign advice to one party.

The summer of 1959 was devoted to the data reduction job described above. In October 1959, when the preliminary data processing had been substantially completed, a review meeting in New York was attended by many of the same persons who had been at the Advisory Council meeting in May plus a number of social science consultants, including Harold Lasswell, Paul Lazarsfeld, Morris Janowitz, and John Tukey. Although the degree of confidence in the basic approach ranged from enthusiasm to doubt, a decision to proceed was quickly reached.

The next step was the development of computer programs, some

⁴ We wish to express our particular thanks to Thomas Finletter, Robert Benjamin, Joseph Baird, and Curtis Roosevelt for encouragement and cooperation.

of which will be discussed below. One objective was to make possible rapid incorporation of new data which might, we hoped, become available during the campaign. Our hope, as we shall see, was only slightly fulfilled.

By June of 1960 we were able to prepare a first report as a sample of the kind of thing which might be done by the Simulmatics process. That was the report on the Negro vote in the North.

Our contractual arrangements with our sponsors ended with the preparation of the process and of this report illustrating it, shortly before the 1960 convention. It was understood that actual use of the service in the form of further reports on specific topics would be purchased by appropriate elements of the party in the pre-campaign and campaign period at their discretion. In the immediate pre-convention period, the National Committee felt that it should not make decisions which would shortly be the business of the nominee. After the convention, the Kennedy organization, contrary to the image created by the press, did not enter the campaign as a well-oiled machine with a well-planned strategy. Except for the registration drive, which had been carefully prepared by Lawrence O'Brien, no strategic or organizational plan existed the day after the nomination. It took until August for the organization to shake down. No campaign research of any significant sort was therefore done in the two months from mid-June to mid-August, either by Simulmatics or by others. In August, a decision was made to ask Louis Harris to make thirty state surveys for the Kennedy campaign. However, because of the late start, data from these surveys would not be available until after Labor Day. On August 11, the National Committee asked The Simulmatics Corporation to prepare three reports: one each on the image of Kennedy, the image of Nixon, and foreign policy as a campaign issue. These three reports were to be delivered in two weeks for use in campaign planning. Along with them we were to conduct a national sample survey which, in the minds of the political decision makers, would serve to bring the Simulmatics data, based as they were on old polls, up to date. (It should be mentioned that one of the most difficult tasks of the Simulmatics project was persuading campaign strategists that data other than current intelligence could be useful to them.) The national survey by telephone was conducted for the project by the Furst Survey Research Center and was indeed extremely useful in guiding the use of the older data. It confirmed the published Gallup finding that Nixon was at that point well in the lead, though we disagreed on the proportion of undecideds (we found 23 per cent). It made us aware that Nixon's lead was due to women. It also persuaded us that voters were largely focusing upon foreign policy at that point in the campaign.

The relationship between the use of such current intelligence and the use of a simulation model developed out of historical data is analogous to the relationship between a climatological model and current weather information. One can predict tomorrow's weather best if one has both historical information about patterns and current information about where one stands in a pattern. While it would be presumptuous to assert that in two weeks of intense activity we approached an effective integration of the two sets of data, that was the ideal we had in mind and which in some limited respects we approximated.

It should be added that the introduction of the national survey data was possible only because of prior preparation for rapid data analysis. The survey was ordered on a Thursday, the field interviewing took place between Saturday and the following Thursday, by Friday morning all cards had been punched, and by Friday night the pre-programmed analysis had been run and preliminary results were given to the National Committee.

The three reports that had been ordered on August 11 were delivered on August 25. The speed of the entire operation is, of course, a testimony to the advantages of a high-speed computer system. Nonetheless, such intense pressure is not an optimum condition for research work, even though rapid analysis was one of our objectives from the start. The reader who suspects that under those circumstances clerical errors inevitably occurred is quite right. It was our good fortune that none of those which we have found since in rechecking have turned out to alter any conclusion, but we do not recommend such limited schedules as a normal mode of work. Nevertheless, with well-prepared computerized analysis, it can be done when necessary.

The reader may ask whether the large preparatory investment was justified in terms of the quantitatively limited use of the project. When we planned the project, we—perhaps unrealistically—anticipated active campaign work from the beginning of the summer until about September 15. (Anything done later than that would hardly be useful.) How far the investment was justified by the two weeks of work actually done is a question which we find impossible to answer. An answer depends on an estimate of how much impact the contents of the reports had on the campaign. The reports received an extremely limited elite circulation. They were seen during the campaign by perhaps a dozen to fifteen key decision makers, but they were read intelligently by these talented and literate men.

Despite the contraction of our effort, our own feeling is one of relative satisfaction that the Simulmatics project was able to provide research on demand concerning the key issues at perhaps the critical

moment of the campaign. While campaign strategy, except on a few points, conformed rather closely to the advice in the more than one hundred pages of the three reports, we know full well that this was by no means because of the reports. Others besides ourselves had similar ideas. Yet, if the reports strengthened right decisions on a few critical items, we would consider the investment justified.

EXAMPLES OF USE OF THE SYSTEM

Earlier in this article we listed three uses of the method herein described: providing a "data bank," rapidly available; providing data on small, politically significant groups; permitting computer simulation. The first of these advantages has perhaps already been adequately illustrated. Let us turn to the other two.

Our report on Northern Negro voters did not use a computer simulation but rather illustrated the capability of the process to provide information about small subgroups of the population. Compare here a number of quotations from the report with what we could have said working from a single survey containing responses from perhaps 100 Northern Negroes. The report demonstrated, for example, that between 1954 and 1956

[A] small but significant shift to the Republicans occurred among Northern Negroes, which cost the Democrats about 1 per cent of the total votes in 8 key states [a shift which continued in 1958]. In those years, the Democratic Party loss to the Republican Party was about 7 per cent of the Northern Negro vote—enough to cause a one half per cent loss in the *total* popular vote in the eight key states. In addition, among Northern Negro Independents, only about one quarter actually voted Republican in 1952, but about half voted Republican in 1956, enough of a shift to cause an additional loss of a little less than one half per cent of the total popular vote in the eight key states.

. . .

The shift against the Democrats is more marked among the opinion leading middle class Negroes than among lower-income Negroes.

. . .

Anti catholicism is less prevalent among Negroes than among Northern, urban, Protestant whites.

. . .

The most significant point of all is the fact that the shift is not an Ike-shift: it is a Republican Party shift. It affects Congressional votes as much as Presidential votes.

In addition, the report demonstrated that Northern urban Negroes vote as often as whites of comparable socio-economic status, and that "there is no sharp difference between Negroes and comparable whites in their feelings about Nixon."

This report was made available to all the leading Democratic candidates, to the Democratic National Committee, and to the drafters of the Democratic platform. Probably no one can say what influence, if any, it had upon them. Those men themselves would not know which of the many things they read or heard shaped their decisions. As outside observers, we can assert only that the report was placed in the hands of the platform framers in the ten days preceding the drafting of the problem, and was read.

The most dramatic result, however, was, as indicated above, the finding that Eisenhower had not generated among Negroes the kind of personal following that he had among most white voter types. This suggested that the Negro vote presented far more of a problem to the Democratic campaign than appeared at first glance; it could not be assumed that the losses in recent years would be recovered with Eisenhower out of the picture.

SIMULATIONS

We turn now to what was perhaps the most novel aspect of the study—the use of computer simulations. We describe, first, how we simulated state-by-state results and, second, how we simulated the impact of the religious issue.

One of the benefits gained from the large number of interviews we used was the possibility of approximating state-by-state results. A national sample survey—even a relatively large one—has too few cases from most states to permit any significant analysis of state politics. The same would have been true, however, even for our voluminous data if we had attempted to do a state-by-state analysis in a simple way. We had an average of about 2,000 interviews per state, but that is a misleading figure. In a small state there might have been no more than 300 or 400 interviews, and on a particular issue cluster that had occurred, for example, in only one-tenth of the surveys, there would be too few cases for effective analysis. We therefore developed a system for creating synthetic, or simulated, states.

By an elaborate analysis of census, poll, and voting data—made more difficult because 1960 census results were not yet available—we developed a set of estimates on the number of persons of each voter type in each state. (Note that since *region* was one of the defining characteristics for the 480 voter types, there were at most only 108 voter types in any given state.) It was assumed that a voter of a given voter type would be identical regardless of the state from which he came. A simulated state therefore consisted of a weighted average of the behaviors of the voter types in that state, the weighting being proportional to the numbers of such persons in that state. For example, we

thus assumed that the difference between Maine and New York is not truly a difference between New Yorkers and inhabitants of Maine as such, but a difference in the proportions of different voter types which make up each state. We assumed that an "upper-income Protestant Republican rural white male" was the same in either state, and that a "small-city Catholic Democratic lower-income female" was also the same in either. This assumption enabled us to use all cases of a voter type from a particular region in arriving at a conclusion for a state.

We do not assert that the assumptions on which this simulation is based are true. On the contrary, we can be sure that they are partly false. The interesting question intellectually is how good were the results obtained with these partially true assumptions. The test is, of course, how far state-by-state predictions made on these assumptions turn out to correspond to reality. To the extent that they do, they suggest that the essential differences between states in a region are in distributions of types rather than in geographic differences, even within a voter type.⁵

Upon this simulation of states was built a second and more interesting simulation, one which attempted to assess the impact of the religious issue. Since the one simulation rests upon the other, the effectiveness of the state simulation is simultaneously tested by examination of the religious simulation. The latter, the main simulation actually carried out during the campaign, represented a hypothetical campaign in which the only issues were party and Catholicism. Our report of this simulation was limited to the North because of the peculiar role of party in the South. The outcome was a ranking of thirty-two states ranging from the one in which we estimated Kennedy would do best to the one in which we estimated he would do worst. The ranking was:

- | | |
|------------------|-------------------|
| 1. Rhode Island | 13. Ohio |
| 2. Massachusetts | 14. Montana |
| 3. New Mexico | 15. Minnesota |
| 4. Connecticut | 16. Missouri |
| 5. New York | 17. Pennsylvania |
| 6. Illinois | 18. Nevada |
| 7. New Jersey | 19. Washington |
| 8. California | 20. New Hampshire |
| 9. Arizona | 21. Wyoming |
| 10. Michigan | 22. Oregon |
| 11. Wisconsin | 23. North Dakota |
| 12. Colorado | 24. Nebraska |

⁵ The states where the simulation was most notably off included Arizona, Nevada, New Mexico, Idaho, and Colorado, states mostly of small population, and states which, in the absence of a "Mountain Region" in our classification, we attempted to treat as Western or Midwestern. Clearly, the assumption of regional uniformity was misleading as applied to them.

- | | |
|------------------|------------|
| 25. Indiana | 29. Kansas |
| 26. South Dakota | 30. Utah |
| 27. Vermont | 31. Idaho |
| 28. Iowa | 32. Maine |

The product-moment correlation over states between the Kennedy index on the simulation (not strictly speaking a per cent) and the actual Kennedy vote in the election was .82. It should be emphasized that this satisfying result was based upon political data not a single item of which was later than October 1958. Surveys on the 1960 election were not available soon enough to be incorporated into this analysis.

The basic method in this simulation was a fairly straightforward application of the cross-pressure findings of earlier election studies.⁶ These findings enabled us to improve our estimate of how a particular voter will behave if we know the cross-pressures he is under. With such knowledge, an analyst should feel more comfortable making guesses about how voters under particular kinds of cross-pressure will shift in an election than he would about making an over-all intuitive guess at the outcome. The method of this simulation was to make a series of such detailed estimates and then let the computer put them together to give an over-all outcome.

To make these detailed estimates we classified our set of 480 voter types into 9 possible cross-pressure subsets arising from a 3-by-3 breakdown on religion and party: Protestants, Catholics, and others; Republicans, Democrats, and Independents. For each of the nine resulting situations we made a prediction. For example, take the Protestant Republicans. They were not under cross-pressure. Since our data had revealed no substantial dislike of Nixon as an individual among such voters, we saw no reason why their vote in 1960 should differ substantially from their vote in 1956, even though Eisenhower was not running. Thus for them we wrote two equations:

$$\begin{aligned} V_k &= P_{56} (1 - P_{55}) \\ V_n &= Q_{56} (1 - P_{55}) \end{aligned}$$

meaning that the predicted Kennedy percentage (V_k) in any voter type of this Protestant-Republican sort would be the percentage of persons in that voter type who had indicated a preference for Stevenson in the 1956 polls (P_{56}), reduced by the nonvoting record of that voter type ($1 - P_{55}$).⁷ The equation for the expected Nixon percentage (V_n)

⁶ Bernard R. Berelson, Paul F. Lazarsfeld, William N. McPhee, *Voting: A Study of Opinion Formation in a Presidential Campaign*, Chicago, University of Chicago Press, 1954.

⁷ Since we trichotomized results, $P_{56} + Q_{56}$ do not add up to 100 per cent. The reader may wonder why a turnout correction is added: are not the residuals the nonvoters? The answer is that a turnout correction is needed because many more persons express a candidate preference on a poll than actually turn out to vote.

was the same except that it used the 1956 Eisenhower supporters (Q_{56}).

The above was the simplest set of equations used. Let us now turn to a more complicated set, that for a group under cross-pressure—Protestant Democrats. First, we decided that, barring the religious issue, 1958 vote intentions would be a better index of the Protestant Democrats' 1960 vote than would their 1956 vote intentions. Too many of them were Eisenhower defectors in 1956 for us to believe that 1956 was a good indicator of normal behavior. On the other hand, 1958 polls would overestimate their Democratic vote, since many of them would defect again against a Catholic. However, it would not suffice merely to subtract the percentage who gave anti-Catholic replies on poll questions, for perhaps those very Democrats who were anti-Catholic were the ones who in practice voted Republican anyway. In short, the question was: Were the bigot defectors right wingers whose vote the Democrats would lose even without a Catholic candidate? Our system could not give us that information for each respondent incorporated into our data. While one respondent in a voter type might have been polled in a survey in 1958 about his vote intentions, another man of the same voter type, on a different survey, might have been polled on whether he would vote for a Catholic for President. To estimate the correlation between these two variables we had to find one or more surveys on which both questions appeared. We then ran anti-Catholicism by 1958 vote for each of the more numerous Protestant Democrat voter types. We found that among them the ratio ad/bc in the following fourfold table averaged about .6. With that information we could estimate how many of the anti-Catholics were hopeless cases anyhow (i.e. had gone Republican even in 1958) and how many would be net losses only in a campaign dominated by the religious issue.

<i>1958 Vote Intentions</i>	<i>Anti- Catholic</i>	<i>Not Anti- Catholic</i>
Democratic	<i>a</i>	<i>b</i>
Republican	<i>c</i>	<i>d</i>

It should be added here that we decided to take poll replies on the religious issue at face value. We were not so naïve as to believe that this was realistic, but since we were not trying to predict absolute percentages, but only relative ones, all that mattered was that the true extent of anti-Catholicism, voter type by voter type, should be linearly related to the percentage overtly expressed. Even this could only be assumed as a promising guess.

Finally, in predicting the vote of the Protestant Democrat voter types, we took account of the established finding that voters under

cross-pressure stay home on election day more often than voters whose pressures are consistent. Therefore, for our 1960 estimate we doubled the historically established nonvoting index for these types.

Thus we arrived at equations applied to each Protestant Democratic voter type:

$$\begin{aligned} V_h &= (P_{58} - a)(1 - 2P_{35}) \\ V_n &= (Q_{58} + a)(1 - 2P_{35}) \end{aligned}$$

The estimate of anti-Catholic 1958 Democratic voters (i.e. persons in cell a in the fourfold table above) was arrived at by the computer, given that

$$\begin{aligned} a + b &= P_{58} & a + c &= P_{14}(P_{58} + Q_{58}) \\ P_{14} &= \text{per cent anti-Catholic} & \text{and } \frac{ad}{bc} &= .6 \end{aligned}$$

Space precludes a similar examination of each of the other of the nine conditions.⁸ Suffice it to say that one other set of serious guesses had to be made, namely what proportion of those Democratic Catholics who had voted Republican in 1958 would switch back to their party

⁸ With the above information, the remaining equations should be decipherable and are reported here for the record:

Protestant Independents, same equations as Protestant Democrats.

Catholic Democrats and Catholic Independents:

$$V_h = (P_{58} + \frac{Q_{58}}{3})(1 - P_m)$$

$$V_n = \frac{2Q_{58}}{3}(1 - P_m)$$

Catholic Republicans:

$$V_h = (P_{58} + \frac{Q_{58}}{3})(1 - 2P_m)$$

$$V_n = \frac{2Q_{58}}{3}(1 - 3P_m)$$

All others:

$$\frac{ad}{bc} = .6$$

$$a + b = \frac{P_{58} + P_{58}}{2}$$

$$c + d = \frac{Q_{58} + Q_{58}}{2}$$

$$a + c = P_{14} \frac{(P_{58} + Q_{58} + P_{58} + Q_{58})}{2}$$

$$b + d = (1 - P_{14}) \frac{(P_{58} + Q_{58} + P_{58} + Q_{58})}{2}$$

$$V_h = \left(\frac{P_{58} + P_{58}}{2} - a \right) \left((1 - 2P_{35}) \right)$$

$$V_n = \left(\frac{Q_{58} + Q_{58}}{2} + a \right) \left((1 - 2P_{35}) \right)$$

to vote for Kennedy and what proportion of Republican Catholics who had voted Republican in 1958 would also switch to Kennedy. After an examination of the trial-heat data from polls which asked about Kennedy vs. Nixon, we decided to use one-third as the proportion in each case, and to use that figure also as an estimate of the proportion of Catholic independents who would be won back by the religious issue.

The simulation required that the computer make 480 separate calculations, each one using the appropriate set of equations from above. During each of the 480 calculations, the computer put into the equations values for turnout record, 1958 vote intention, 1956 vote intention, and anti-Catholicism, derived from the data which had been assembled about that particular voter type. This gave a 1960 vote estimate for each voter type for the particular hypothetical campaign being investigated. Weighted averages of these gave the state-by-state estimates.

These estimates, as we have already noted, turned out to be close to the actual November outcome. They were not intended to be predictions. Or, rather, they were *contingent* predictions only. They were predictions of what would happen *if* the religious issue dominated the campaign. We did not predict that this would happen. We were describing one out of a set of possible types of campaign situation. But by August, when we took our national survey, comparison of our simulation and the survey results showed that this situation was actually beginning to occur. And the closeness of our contingent prediction to the final November result suggests that, indeed, the religious issue was of prime importance.

How close was the religious-issue simulation to the actual outcome compared to alternative bases of prediction? A full exploration of this remains to be made. We must, for example, further vary the parameters used in the simulation to determine which ones affect the results most critically and which values of those give the best prediction. For the present we look only at the one set of values and equations on which we relied during the campaign and which has already been described. (A few variations were tried and dismissed during the campaign, but none that made much difference.) How did this one simulation compare with other predictive data?

An obvious comparison is with the Kennedy-Nixon trial heats on polls taken at the same time as the latest polls used in the simulation. The correlation between the state-by-state result of these polls and the actual outcome is but .53 as compared to .82 for the simulation. The simulation, in short, portrayed trends which actually took place between the time the data were collected and election day. The uncorrected polls two years before the election explained but one-fourth of the variance in the real results, while intelligent use of them taking

into account the cross-pressure theory of voting behavior allowed us to explain nearly two-thirds of the variance.

A more stringent comparison would be with Kennedy-Nixon trial heats run in August 1960, when the simulation was run on the computer. Such a comparison would answer the question of whether the Democratic Party would have gotten as good information at that date by the conventional means of up-to-the-minute field interviewing as it got by reanalysis of old data. Very likely it could have, if it had chosen to invest in a large enough national sample survey to give it state-by-state results, for as far as we can now tell the Catholic issue exerted most of its impact by shortly after the conventions. However, until poll data for that period becomes available we can only speculate. We wish to emphasize, however, that at some point in the history of the campaign, poll data certainly came into close correlation with the November election results and thus with our simulation. The date the raw poll results became as or more predictive than the simulation would be the point in the campaign at which mechanisms of voter behavior anticipated in the simulation became reality.

Besides simulation and polls, what other indices might have forecast long in advance the state-by-state order of voting in 1960? Results of previous elections would be one such index. Perhaps the rank order of the states in a previous election is a good forecast of rank order in future ones, even if the electoral outcome changes. (The whole country could move one way or the other, leaving the order of the states much the same.) But, if one is to use this device, which election should one use? The year 1956 was a presidential election year, as was 1960, but in 1956 the Eisenhower phenomenon was operating. 1958, although more recent and less affected by Eisenhower's idiosyncratic appeal, was a Congressional election year. In our simulation, too, we faced this problem. We resolved it for some voter types one way, for some another. But what happens if one relies on a single simple over-all assumption of continuity between elections? The result is not very good, though slightly better using 1956 than 1958. The product-moment correlation of Northern results between 1956 and 1960 was .39, between 1958 and 1960, .37. The multiple correlation using both earlier years was .44 with the 1960 election. So far our simulation clearly was superior as a forecast.

Perhaps one might have made a good prediction of the impact of the religious issue by a simple slide-rule method of calculation instead of by an elaborate computer procedure. One could correct the 1956 or 1958 vote by some crude percentage of the Catholic population of each state. That would have worked and worked well if, by some act of intuitive insight, one could have hit on the right percentage correction. One would have had to decide first of all to use 1956, not 1958, as the

base, for no simple correction of the 1958 results gives a good correlation with the actual outcome. If one had that correct flash of intuition, one could have surpassed our complex simulation with a correction of exactly 34 per cent of the Catholic percentage of the population added to the Democratic vote. The correlation with the actual outcome achieved by this process is .83. The simulation was better, however, than any correction except 34 per cent. It is better than 33 or 35 per cent. At corrections of 32 and 40 per cent the coefficients of correlation for the simple correction procedure drop below .80.

There was, in other words, a "lucky guess" way of estimating the effect of the religious issue in the campaign which would have given an excellent prediction. But even if we had tried to make such an overall estimate and had somehow arrived at the right "lucky guess," we could not have defended it against skeptics. What the simulation did was to allow competent political analysts, operating without inspired guesses, to make sober, scientifically explicable estimates that they were willing to commit to paper before the facts. As the accompanying table shows, the simulation gave results about as good as the very best which *hindsight* now tells us could have been reached by simpler methods if infused by the right lucky guesses.

CORRELATIONS WITH ACTUAL ELECTION RESULTS

Trial heats contemporaneous with simulation data	.53
Continuity with 1956	.39
Continuity with 1958	.37
Continuity with 1956 and 1958	.44
1958 results with optimum, or "lucky guess," correction for Catholic vote	.83
Simulation as done during campaign	.82

The essence of the simulation was to treat each voter type separately. Under what conditions should one expect that procedure to obtain a better result than an optimal across-the-board correction applied to the total? Clearly, if the process at work in each voter type was uniform it would make no difference whether we applied correction factors voter type by voter type or to the total. One could add 34 per cent of each Catholic voter type to the Democratic vote for that type or add 34 per cent of total Catholics to the total Democratic vote and come out with the same result. Where there are complex interactions of several variables on a voter type, however, then calculations done the two ways are no longer equal. If, for example, turnout varies between voter types and party voting also varies, then an equation applied to each voter type could not equally well be applied to total voters.

It is clear that we did not use the most predictive values for all parameters in our simulation. Determining what these were with the aid of hindsight is part of our present research program. But before election day we had no way of knowing what they were. (The one-third of 1958 Catholics casting Republican votes likely to go Democratic in 1960 according to our equations should not be confused with the 34 per cent of all 1956 Catholic voters, which turned out to be a good across-the-board correction.) The fact that our result came out on a par with the optimum simple correction which hindsight has enabled us to make is a crude measure of the gain from working voter type by voter type, with account taken of interactions within each type, that is, the gain from the computer operations.

The test of any new method of research is successful use. The outcome of the present study gives reason to hope that computer simulation may indeed open up the possibility of using survey data in ways far more complex than has been customary in the past. The political "pros" who commissioned this abstruse study were daring men to gamble on the use of a new and untried technique in the heat of a campaign. The researchers who undertook this job faced a rigorous test, for they undertook to do both basic and applied research at once. The study relied upon social science theories and data to represent the complexity of actual human behavior to a degree that would permit the explicit presentation of the consequences of policy alternatives.

This kind of research could not have been conducted ten years ago. Three new elements have entered the picture to make it possible: first, a body of sociological and psychological theories about voting and other decisions; second, a vast mine of empirical survey data now for the first time available in an archive; third, the existence of high-speed computers with large memories. The social science theories allow us to specify with some confidence what processes will come to work in a decision situation. The backlog of survey data permits us to estimate the parameters of these processes with fair precision and great detail for each small element of our national population. The computer makes possible the handling of this mine of data. More important still, it makes possible the precise carrying out of long and complex chains of reasoning about the interactions among the different processes. In summary, we believe that conditions now exist for use of survey data in research far more ambitious than social scientists are used to. If it is possible to reproduce, through computer simulation, much of the complexity of a whole society going through processes of change, and to do so rapidly, then the opportunities to put social science to work are vastly increased. It is our belief that this is now possible which was put to a test by the campaign research reported here.