THE QUANTITATIVE ANALYSIS OF SOCIAL PROBLEMS

Edited by
EDWARD R. TUFTE
Department of Politics
Princeton University

ATTITUDES AND NON-ATTITUDES: CONTINUATION OF A DIALOGUE

PHILIP E. CONVERSE

Some years ago Carl Hovland (1959) undertook a systematic comparison of conclusions which had been reached concerning the modification of attitudes through communication within two broad research traditions: the *experiment* and the *sample survey*. Despite common substantive interests, these two traditions had remained rather insulated from one another, and as Hovland noted, their separate efforts had produced results easily taken as contradictory. The purpose of Hovland's essay was twofold. First, he wished to point out that these apparently divergent views about the ease of persuading to attitude change could be readily reconciled by proper understanding of the differences in the two types of research designs and the backgrounds of the investigators involved in each. To the best of our knowledge, members of both traditions felt that the reconciliation was well handled. Secondly, however, Hovland seemed concerned in a more general sense with a need to bridge the gap between experiment and survey, and in effect called for a more vigorous dialogue between the two traditions.

This paper is written in an effort to continue the dialogue. It is not our intention here to retrace Hovland's steps in any detail, although a number of things to be said have relevance for his argument. Nor, for that matter, do we wish to rehearse the differential powers and shortcomings of the two methods, for we would subscribe with little amendment to most of the conclusions long since reached by parties interested in the subject. It seems that divergences in outlook are likely to arise between the two traditions if for no other reason than the fact that the experimentalist is able to study what *can* happen in situational configurations which he creates and controls, while survey analysts pay a more passive attention to what *does* happen in an actuarial sense, as a matter of relative empirical frequency. Thus, in the case Hovland discussed, it appears true that when audiences are exposed to certain persuasive communications on certain kinds of issues under certain further experimental conditions, considerable attitude change can be demonstrated. From an actuarial point of view, however, these conditions occur in nature infrequently: few people expose themselves to potentially contrary messages even though a torrent of such messages may be sent, and various other aspects of the experimental condition go unfulfilled as well.

Revised version of a paper read at the Seventeenth International Congress of Psychology in Washington, D.C., August, 1963. Reproduced by permission of the author.

Perhaps the clearest function which sample survey results may fulfill in any dialogue, then, is to remind the experimentalist of the actuarial mainstream, pointing out sources of critical variation in "natural" attitude-change processes to which he may have become insensitive. It is in this spirit that our remarks are made. We will begin by a presentation of results from one analysis of sample-survey data on attitude change. We choose this analysis because it seems to have implications not only for some of the methodological practices common in experimental work on attitude change, but also for the way in which the attitude continuum is to be conceptualized. Carried a step or two further, the results suggest hypotheses which might be worthy of experimental test within the areas described by balance, congruity or dissonance theories of attitude change.

Statistical Properties of Certain "Naturalistic" Attitude Changes

The data to be reported are drawn from a sequence of panel studies conducted on a national cross-section of the adult population of the United States over a four-year period from 1956 to 1960. The immediate occasions for the surveys were the three national elections of 1956, 1958 and 1960. In principle, respondents were interviewed at five different points in time over this period: before and after the election in 1956; after the election in 1958; and once again, before and after the 1960 election. In practice, slightly less than 70 per cent of the original 1956 pool of respondents who were still alive and in possession of their faculties were successfully reinterviewed in the 1960 waves of the panel. However, analyses across a wide variety of social and attitudinal characteristics suggested that the 1960 survivors were a remarkably unbiased subset of the original pool of respondents.

The panel aspects of the study sequence were fully utilized, in the sense that the interview schedules applied at the various points in time contained direct repetitions of a wide array of items bearing on attitudes and social situations. Here we shall focus on some of the time-change aspects of the attitudinal data.

The simplest property to assess from these attitude measurements has to do with the correlation of measurements of the same items over time. Empirically, the range in variation of such test-retest correlations on attitude measurements computed after a two-year interval was very wide, covering a space in the correlation continuum from roughly .10 to .80.[1] Furthermore, the data made clear that these turnover correlations were quite stable for specific items over comparable periods of time: if the correlation of item A between t_1 and t_2 turned out to be .4, for example, it was demonstrable that the same correlation computed for the same item between the interval t_2 and t_3 would rarely be much more than .03 or .04 from .40 as well. In short, then, varying items had very differential but very stable rates of attitude turnover associated with them over the two-year test-retest span.

Cursory inspection of differences in turnover rates attached to varying items was sufficient in many cases to indicate why the turnover rate fell in the levels where it was found. Thus, for example, items which showed test-retest correlations of less than .20 typically had to do with the particular election as an object of estimation. Since three elections were covered, it is quite obvious that while the questionnaire items were identical, the objects of reference—the different elections—were in many respects different objects, so that almost complete turnover of opinion could hardly be taken as surprising. At the very high end of the turnover range, measures of generalized affect toward the major political parties approached a two-year test-retest correlation of .80. Here, of course, the objects are familiar ones evoking strong affect from many people, and are characterized by a reasonable amount of temporal stability in their attributes as perceived by the public.

Included among the attitude items, however, was a battery of eight questions directed at the principal issues of public policy which were being debated during the period of the study. These included attitudinal items of a familiar Likert type on matters of civil rights, social welfare legislation, the relation of government to free enterprise, and problems of foreign policy such as aid to neutral countries and the like.

We shall focus our attention on this set of attitudinal items for two reasons. First, they resemble closely the type of item to which the experimentalist often gravitates in his attitude studies. Secondly, the empirical properties of the items were quite striking. For while the meaning-substance of the items seemed to have changed very little during the period of the study, and while the items were chosen specifically to capture the cleavages in basic psychological commitments in public politics of the period, the test-retest correlations were all within the lower half of the empirical range, running roughly from .23 to .46 within both of the two-year time spans available for inspection.

There are at least two possible reactions to such low coefficients. One is to imagine that public opinion on these items must have been in a high state of flux during this period, with a very considerable evolution of attitudes in the wake of changing national events. There were, however, several empirical problems which such an interpretation would encounter. In the first place, the marginal attitude distributions for the various time points were remarkably similar despite high rates of turnover within the tables.[2] It would seem that if national events were exerting systematic forces on opinion in a manner which would produce meaningful evolution of public attitudes, the distributions of opinion should progress in one direction or another over time, rather than remain relatively stable, with almost all of the individual change in one direction being counterbalanced by an equal amount of individual attitude change in the opposing direction. Secondly, the items had been chosen to avoid superficial attitudes toward short-term events, in an attempt to plumb more basic and stable dispositions toward questions of public policy. Hence no very large measure of meaningful change was to be expected. Indeed, examination of the differential rates of turnover within the battery of issue items suggested that the more basic and ideological the issue dimension was, and the more remote its referents were from day-to-day change in national events, the higher the turnover of opinion. Thus, for example, the most unstable issue of all was one having to do with whether or not the federal government should "leave things like electric power and housing for private businessmen to handle." Somehow it seemed implausible that large proportions of the American population between 1956 and 1958, or between 1958 and 1960, had shifted their beliefs from support of creeping socialism to a defense of free enterprise, and that a correspondingly large proportion had moved in the opposite direction, forsaking free enterprise for advocacy of further federal incursions into the private sector.

This being so, the more reasonable reaction to the low coefficients is simply that we were doing a very poor job of tapping the attitudinal dimensions at which we originally aimed, and that our results, viewed now as test-retest correlations in the strict sense, give witness to an incredible degree of measurement unreliability. Hence we should not talk of results at all until we go back and develop better measuring instruments.

There is somewhat more internal evidence supporting this view than was the case for the evolution-of-opinion interpretation. One of these items of evidence is particularly relevant to our argument. Originally we had assumed that some of these items, vis-à-vis some people in the population, would have very little meaning. Hence we had taken elaborate precautions to remove such people from any sense of obligation to respond to items which generated no affect for them. Thus, for example, the battery of issue items was prefaced by a statement which pointed out, among other things, that "different things are important to different people, so we don't expect everyone to have an opinion about all of these (things)." Furthermore, as each item in the battery was read, the respondent was explicitly asked whether or not he had an opinion on the matter and only if he said that he did was he further asked what his opinion was. These precautions may be fruitfully compared with many common attitude-measurement situations in which experimental instructions entreat the subject to express some kind of attitude other than the indifference point even though he may find it hard to do so. After all, "don't know's," equivocations and other forms of missing data are exasperating in analysis and are to be avoided at all cost. However, we felt that it was more important to deal with missing data than with measures laden with "noise."

Our screening procedures were successful enough that variously across the set of attitude items, anywhere from a handful to 35 per cent of the sample confessed that it had no genuine opinion on the matter under consideration. Ironically, however, the low test-retest correlations which we have cited for these items were computed for the subset of people who did lay claim to some opinion; the many no-opinion people were set aside. Furthermore, it was readily discovered that there was almost a perfect correlation between the issue items ordered according to proportions of people who said they had no opinion, and the ordering produced across the same items in terms of response stability.

These pieces of evidence, of course, lead toward a strong suspicion of item unreliability, although it is rather distressing to learn that such unreliability remains after all of our precautions to avoid the measurement of non-existent attitudes. However, let us not be too hasty in our judgment, for the heart of the analysis which we wish to report turns out to throw an odd and unexpected light on the whole question of reliability.

The most revealing statistical property of these attitude-change data emerges when we consider not simply the correlations between the same attitudes over two-year spans, but also the correlation for each attitude between the initial and terminal interviews, a span of four years. For we discover that these t_1-to-t_3 correlations tend to be just about the same magnitude as the t_1-to-t_2 correlations, or the t_2-to-t_3 correlations. That is, surprising though it may be, one could predict 1960 attitudes on most of these issue items fully as well with a knowledge of individual attitudes in 1956 alone as one could with a knowledge of the more proximal 1958 responses. Furthermore, the tendency toward parity of the three correlations is clearest among the issue items with greatest turnover; among the more nearly stable items, the four-year correlation tends to be slightly lower than the two-year correlations, a pattern which is of course much closer to our intuitive expectations.

At this point we will find it useful to shift some of our weight from the vocabulary of correlation to the vocabulary of Markov chains, a body of mathematical theory useful in treating stochastic change processes. Within this vocabulary, the tables representing attitudes at different points in time can be converted into proportions by rows, and considered as empirical matrices of transition probabilities. The similarity of the pair of two-year tables for most of these issue items provides some presumptive evidence for the assumption that whatever else may be said, these matrices as reckoned in summary fashion over the population can usually be considered constant in their probabilities over time.

Proceeding with this assumption, it can readily be demonstrated that the empirical parity of the three time-correlations could not occur if it were true that the issue responses of the total population were properly describable by a single constant matrix of transition probabilities through the two steps from t_1 to t_3. For if such a condition held, then the four-year time-correlations would necessarily take values substantially below those of the two-year correlations, the differences ranging upward from .10 to .20 or more under common conditions.

Hence it is clear that no single transition matrix constant over the two-year stages can account for the observed response behaviors. In other words, if we maintain the assumption of constant matrices over time, then we are forced to conclude that the time paths of response can only be treated as arising from some mixture of two or more transition matrices, and not one alone.

Let us imagine what these matrices might be. At first glance it is apparent that there is a whole range of matrices which, while mathematically possible, are entirely implausible in the empirical situation at hand, and which we can therefore rule out of serious consideration as describing the behavior of *any* member of the population. These are the subset of possible matrices which, if cast as second-order determinants, have negative values, or which in any form would generate negative time-correlations. In content terms, a genuine negative time-correlation would mean that respondents who had agreed with an item in 1956 must have gone out of their way to disagree with that item in 1958, and then remember to agree with it again in 1960. This would seem to be highly unlikely behavior, even if it were conceivable that respondents could remember specific responses they had made to an interview two years earlier; and since this seems quite inconceivable in itself, it seems fair to rule out such matrices as impossible.

Once this constraint is added to the picture, however, we begin to have some strong leverage on our data. In particular, we can now say that if the total population were to be described in terms of a mixture of *two* transition matrices, then there are only two which in combination would generate our phenomenon of three equal time-correlations. These two matrices are (1) the identity matrix, with probabilities of 1.0 in the major diagonals; and (2) what we might call a "random" matrix, of equiprobable responses. And quite naturally, the magnitude of the three equal time-correlations would be a simple function of the relative prevalence of these two matrices in the population.

This model for response behavior we might call a "black-and-white" model, for it posits a very stringent division of the population into two sharply contrasting subsets. In content terms, one portion of the population would be perfectly stable in its responses over time, while the other portion would be given to response time-paths which in a strict statistical sense were random. On the face of it, this would seem to be an unlikely descriptive model. Furthermore, while most of the assumptions we have made about the nature of the data which have led us to this model seem quite palatable, our final assumption was that only two transition matrices were necessary to describe the behavior of the population. This constraint seems somewhat gratuitous: we would much sooner suppose that a continuous shading of different matrices were represented in the population.

Happily, in this instance, the data permit a test of the goodness of fit of the simple black-and-white model. There are several statistical manipulations which would accomplish the same general end. We will describe one which is as simple to follow intuitively as any of them.

Let us assume that the black-and-white model comprises an accurate and exhaustive account of the attitude responses generated over time. Then it follows that any individuals who change from one side of an attitude scale (say, the "agree" side) to the other ("disagree") between t_1 and t_2 must of necessity belong to that portion of the population whose responses are random. They are a "pure" random group. However, they do not exhaust the total set of subjects following random response paths. For between t_1 and t_2 certain subjects following random response paths would give two consecutive responses on the same side of the issue by chance alone. Fortunately, the proportion of such subjects in the cells of stable t_1-t_2 attitudes is calculable, given the numbers of subjects who fall in the minor diagonals, or change cells. This does not mean, of course, that we can isolate the "random" subjects from those whose response paths are stable in a meaningful sense. But we can tell in considerable precision how "polluted" with such people the set of apparently stable people is.

Therefore, we can define two sets of people on the basis of the pattern of t_1 and t_2 responses. One is the pure random set. The other is a mixture, in known proportions, of the perfectly stable and the perfectly random respondents.

If the black-and-white model is a proper description of response behavior in the population, then, we can predict with similar precision the nature of the relationship between t_2 and t_3 responses for each subset. The purely random people should show a t_2-t_3 correlation between responses of .00. The polluted subset should show a t_2-t_3 correlation greater than the total-population time-correlations, but falling well short of unity because of the remaining admixture of random respondents. Just how far short of unity this second correlation would fall can of course be readily calculated on the basis of knowledge of relative proportions of random and stable respondents.

As an initial test of the model we chose data arising from the item for which the black-and-white model seemed on a priori grounds to be least inconceivable. That is, we sought the item from our battery for which it seemed likely that *genuine* attitudes would be most deeply ingrained and hence immutable over time, and yet one which would be basically "ideological" enough that the issues posed might be truly beyond the ken of substantial numbers of people in a cross-section population. The item described above concerning the relative roles of government and private business matched these specifications in excellent fashion. It was not entirely coincidental that this was the item on which the largest proportion of respondents had indicated that they had no opinion, and the item which had shown the highest response instability of any in the battery among those who did claim opinions.

Respondents were then sorted into two groups in the manner described above, on the basis of the pattern of t_1 and t_2 responses. The predictions were that the pure random group would show a t_2-t_3 correlation of .00, and that the second group would show a correlation of .47, if the black-and-white model were indeed an appropriate description of the underlying response process. The results of the primary test showed t_2-t_3 correlation values of .004 for the first group, and of .489 for the second group. In other words, the time data generated by this issue item fit the predictions of the black-and-white model with remarkable precision.[3]

We then proceeded to test the same model for a number of the other issue items, although in varying degree they had face content and zero-order time-correlations which suggested in advance that they could be expected to depart from the model somewhat. The test showed that indeed they did, although the bifurcation of t_2-t_3 correlations between the two test groups remains quite extreme. Thus, for example, the key correlation for the putative "random" subset departs from .00, but rarely rises above a figure of about +.09. Hence we might conclude that there remains a "near-fit" for the other items as well.

This fact of "near-fit" may be conceptually more important than meets the eye at first glance. For there are widespread assumptions about processes of attitude measurement, as well as a few formal models, which presuppose some underlying continuum of latent response probabilities vis-à-vis any single attitude item. Thus, for example, in a heterogeneous population, one might expect that ideally one could isolate individuals for whom response probabilities towards taking one of two positions on a given item range continuously between .5 and 1.0. What is intriguing about the black-and-white model, along with the "real" data which fit it, is the demonstration of an absence of such continuity, with two maximally discontinuous classes (or three classes, if one distinguishes between the perfectly stable "pro" class and the equally stable "anti" class).

Once data depart significantly from this simple model, the number and variety of models which could conceivably account for the data become large indeed, and the discriminatory power of our mathematical deductions evaporates accordingly. In these instances it is very easy to fall back on vaguer notions of latent-response probability continua. Perhaps this is appropriate. However, the fact that one set of these data fits the black-and-white model very well, and the other sets of conceptually comparable items only miss a fit with the model in modest degree, suggests that we should not abandon the black-and-white model completely in imagining the processes which underlie the responses to other items in the battery. In other words, it would seem likely that were the truth of the matter isolable, we would discover that a very large proportion of the responses to the other items in the battery could best be understood in terms of two sharply discontinuous classes of respondents, the stable and the random. What is new in these other items, and what leads further data to diverge somewhat from the black-and-white model, is the presence of some few people who are undergoing a meaningful evolution of attitudes on the issue in question. The crucial fact, from the point of view of our argument, is the strong likelihood that even the attitude items straying somewhat from the expectations of the black-and-white model are clouded by large numbers of purely random responses.

What psychological interpretation is to be placed on such random responses? It seems to us most simple to imagine that they came from people with no real attitudes on the matter in question, but who for some reason felt obliged to try a response to the item despite our generous and repeated invitation to disavow any opinion where none was felt. In this vein, it may be useful to analyze exactly what the objects were which we were asking our respondents to evaluate. For example, the key black-and-white item on government and private enterprise posed as an object of potential attitude not just the federal government or private business, but rather a *type of relation* between the two. Furthermore, the manifest content failed to make clear which of the two parties to this relation would feel helped or hurt by it. This means that respondents who may have had some prior feeling of generalized affect toward private business or toward the federal government (or both) could not respond stably, for lack of this further information which the question presupposes. The experimentalist may find this an incredible observation, for the information presupposed—that private business does not generally want further governmental expansions into its economic sector—seems almost a ubiquitous piece of the "common culture." However, the survey analyst rapidly comes to recognize that presupposition of any information about objects which lie beyond the daily ken of the subjects tested will miss the mark for substantial numbers of people in a heterogeneous population. Indeed, in the case of the attitude item fitting the black-and-white model, it can be calculated from the data that something less than 20 per cent of the total sample fell into the category of real and stable attitudes on the item. The remaining 80 per cent represented confessions of "no opinion" or statistically random responses. Unfortunately, it was a *minority* within this 80 per cent which took advantage of our invitation not to bother fabricating an opinion. When attitudes are asked for in such a setting, people are remarkably obliging.

Relevance of Findings for Experimental Studies

At this point the experimentalist may well ask, "Of what interest is all this to me? I want to understand the implications of attitudes once people have psychological states worthy of the name. I am a student of attitudes, and not a student of non-attitudes." This is, of course, an impeccable position. It presumes, however, that what he typically studies are indeed attitudes, and not non-attitudes concealed with hastily-fabricated affective judgments, as was the case with a full plurality of our test population. A reading of the experimental literature over the years suggests to me, however, a remarkable insensitivity to this possibility. And when I keep this possibility in mind in reading any given study, I often end up with an interpretation of results that is quite oblique to the interpretation offered by the investigator. In other words, there seems to be food for thought here even for investigators who wish to limit their efforts to the study of genuinely-formed attitudes, but who do little to protect themselves against the measurement of non-attitudes by mistake. Therefore it is appropriate to explore the implications of these survey data several steps farther.

The most obvious implication has to do with instrument reliability. Once we granted that the low-time correlations for these items were not likely to be accounted for by "true" attitude change, but rather should be seen as test-retest coefficients of reliability, we were prepared to send the whole instrument back to the shop for repairs, since a reliability coefficient of .3 is disastrously low. Yet the fit with the black-and-white model suggests that where people actually *had* attitudes, the single item could scarcely be further perfected, for on a trial as stringent as a two-year test-retest, the reliability coefficient was indistinguishable from perfection. From this point of view, what needs repair is not the item but the population. Less facetiously, the moral is clear: where measurement reliability is at issue, the measurement of non-existent states is very unrewarding. And while the classical view of these matters took "reliability" to be a property (or number) attached to the measuring instrument, we could not have a more dramatic example of the fact that reliability in our field of inquiry is instead a joint property of the instrument *and* the object being measured.

Other aspects of psychological measurement may deserve review in this light. Speaking from personal experience, I would hypothesize that such a phenomenon as *test fatigue* is itself a direct consequence of pressures, felt by the subject to search for faint or non-existent bits of affect to fulfill the requirements of the attitude questionnaire. In those rare cases where an attitude item or battery dovetails nicely with thoughts or feelings I have experienced on my own with any strength or clarity before, even such an impersonal process as marking a questionnaire offers the reward of pleasant catharsis. Such pleasure seems somewhat infrequent, however, and the hunt-and-fabricate feeling is fully as familiar.[4] One outcome of such harassment is fatigue; another is a more or less conscious recourse to some response set touched off more by question form than question content.

Underlying the overestimate of who has attitudes about what, perhaps, is the common view held by many social psychologists of the individual as a vibrant bundle of attitudes. Nothing we have said need call this view into question in the least respect: it is certainly an heuristic viewpoint and undoubtedly a faithful one as well. However, it is all too easy to assume from such a view that mere selection of a "familiar" object or controversy as a point of attitude measurement must evoke true attitudes in all or almost all of a test population. There is, of course, a very wide logical leap from the first of these propositions to the second. Possible objects of attitudes are infinite, and a person can be seen as a vibrant bundle of attitudes without any assurance that his attitudes extend to more than a very tiny subset of such objects. Phenomenological differences in information and attention almost ensure the contrary: it may well be difficult to find objects in most domains which will not be matters of non-attitude for many members of the test population.

In sum, then, there is a very real sense in which *attitudes take practice*—practice which is genuine in the sense of having been powered by own psychic energy aside from the kind of transient situation created by the experimentalist or the survey interviewer. Where such practice has not occurred, the state to be measured is non-existent. The measurement of non-existent states gives maximally unreliable results. If the subject himself were helpful to the investigator in refusing to report very *ad hoc* feelings as "attitudes," then the problems would be greatly diminished. This does not occur, however. Hartley (1946) years ago collected a full set of ethnic attitudes toward groups that did not exist. We made a great effort to encourage holders of non-attitudes to bypass such items. Some accepted the invitation, but the majority did not. Whatever our intentions, the attitude questionnaire is approached as though it were an intelligence test, with the "don't know" and "can't decide" confessions of mental incapacity.

It is true that on many grounds the survey analyst is more exposed to the dangers of studying non-attitudes than is the experimentalist. That is, despite great variety in experimental procedures, there are some rather typical aspects of attitude-change studies which intentionally or accidentally provide protection. One example of accidental protection is the use of college students as subjects. Non-attitudes on a wide range of matters which seem "common culture" to the investigator are an inevitable consequence of information impoverishment among the less well-educated strata of heterogeneous populations. While the professor is likely to be impressed that college sophomores are not very well-informed either, they remain, relative to the total population, a fairly alert group. Hence, if a specific attitude item were to show an 80 per cent non-attitude rate in a heterogeneous population, the nature of the population the experimentalist uses might well reduce the rate on the same item to something like 30–50 per cent, among college sophomores.

More intentional steps for protection include such things as the multiple-item battery and the choice of attitude-objects which are "close to home" for all of the subjects, and hence far more likely to have become the object of genuine affect for any population member. Certainly one need not worry greatly about non-attitudes in the sense in which we use the term here when the object of evaluation is "mother" and the dimensions of evaluation so common (as in the more typical uses of the semantic differential) that they are part of anybody's common judgmental vocabulary.

However, in attitude-change experiments *per se* there are pressures away from the havens of protection which either multiple-item batteries or the choice of very homely objects can afford. In the classic format of the "before" measurement, the persuasive message and the "after" measurement, the message itself must be of some limited content scope, and there are in turn only a limited number of attitude items which can be imagined within such a scope. Hence the item base of very many attitude-change measures is extremely limited, if indeed it exceeds one item. Similarly, the necessity of dealing with some "common" object of orientation to which the persuasive message can effectively pertain makes the use of objects which are phenotypically knit into the lives of test subjects ("mother," "my work," "my professor," etc.) rather awkward. Add to these difficulties the need to deal in objects which are controversial, along with an understandable desire to treat "socially significant" attitudes, and the common result is not only the use of a narrow item base, but attitude-objects of political, intellectual or social interest which tend to lie beyond what is very salient for many of the test subjects. It is, of course, in such areas that non-attitudes abound, even for college sophomores.

If it can be granted that many attitude studies have measured an abundant number of non-attitudes on the supposition that they were attitudes, what difference has it really made? Do not the holders of non-attitudes, particularly with multiple-item measures, tend to gravitate toward the zero point of "indifference," where they belong? And even though they may abound, do they do anything more than add "noise" to the results, attenuating them rather than biasing them?

First, let us take the question of the meaning of the attitude continuum, for this is what is at stake when we say that non-attitudes should fall at the "zero-point." It might be mentioned in passing that a frightening number of our random respondents were capable of giving "strong agree" or "strong disagree" responses, probably under pressure to introduce some variety into their strengths as well as to give some attitude. With multiple items, however, it is indeed likely that inconsistencies (which will be frequent in the expression of non-attitudes) will drive the respondent toward the middle of any summary attitude measure. The question becomes, then, whether or not that is really where we want him.

Now there are several ways of imputing meaning to zones on this basic attitude continuum. It is generally assumed that an extreme location has something to do with either an intense or at least a univocal attitude. One way of visualizing the matter is to reduce the molar concept of a generalized attitude toward a complex attitude-object into its molecular parts—the set of affective reactions which the individual holds toward all of the component properties of the object which he perceives. In such a reading, the proper location of the individual on the attitude continuum with respect to the molar object is some algebraic summation of ratio of component valences, weighted in one fashion or another.

While we know remarkably little about what combining rules may pertain, we do know that we can expect some generalization of affect or "strain toward symmetry" (Newcomb, 1953) among these molecular valences, such that the molar attitude toward the object tends to be somewhat more univocal than the same affective components might be if the properties of reference resided in a scatter of dissociated objects. However, we also know that this is no more than a trend, and one which reality very often forestalls. People do maintain mixed attitudes toward very many objects, especially those not lending themselves to any easy dissociation. Indeed, the instance of the perfectly-mixed reaction occurs often enough that it has attracted a term of its own, which is ambivalence, an uncomfortable state but not a non-existent one.

Where should ambivalent people be located on the attitude continuum? Most certainly, it would seem, they deserve a position on the middle, at the zero-point. However, we now see that things are becoming somewhat crowded here, for we have already created something like Sherif's zone of indifference, and located our non-attitudes in the middle of this zone. In some instances, perhaps, the consequences of this overcrowding are few. But there are many instances of experimental treatment of attitude change in which it may matter a great deal.

Take as a concrete example the implication of the Osgood congruity model (Osgood *et al.*, 1955 and 1957) that attitudes near the zero point are moved greater distances by persuasive information than are attitudes initially located towards the extremes of the continuum, a matter which has received experimental confirmation in the typical format of the attitude change study. Perhaps such wider movement under the experimental conditions would be equally true of intense but ambivalent attitudes as well as of non-attitudes. Intuitively, however, we would expect that there would be rather drastic differences among holders of the two types of "zero-attitudes." It is easy to see that people bringing actual "non-attitudes" to the experimental situation (whatever affective answers they may have given initially) would be strongly affected by the information inthe persuasive message, since the experimental situation may be one of the first times they have ever paid attention to any kind of information about the item in controversy. However, if we assume that people can become intensely ambivalent only through an appreciation of the pros and cons of a controversy, it does not seem likely that a brief reiteration of some of the pro arguments will create the same kind of striking attitude change.

If initial non-attitudes are present in such experiments, and if the bulk of the systematic change in measurement is attributable to these people, would we not be more accurate in speaking of "attitude-formation" studies rather than studies of attitude change? Even the term "formation" may be unduly strong, for the non-attitude-holders by definition have little interest in the proposed attitude-object, and save in the case where enforced attention sparks some further self-starting interest in the object, it would seem likely that the nascent attitude expressed in the "after" measurement might decay from memory rather readily.

We would argue that holders of non-attitudes, in the measure that they can be detected at all, deserve no location at the zero-point of the attitude scale, but rather should be located "off the continuum" entirely, in proper recognition of a non-existent state. If they were disposed of in this manner, then the truly ambivalent would be left as a relatively pure group in the center of the continuum. In view of a postulated strain toward symmetry among valences of attitude-object properties, this would mean that more of our attitude distributions would show a U-shaped form when properly measured, and fewer would show the heavy concentrations of responses near the scale midpoint noted by Lorge (1937). In this respect, while the nature of his scales were somewhat different, we heartily applaud the suggestion by Brim (1955) that more intermediate or equivocal responses are likely to be associated with impoverished information about the object in question.

Non-Attitudes and Broader Theories of Attitude Change

Up to this point we have limited ourselves to a consideration of the implications of non-attitudes for certain methodological practices and the interpretation of some experimental findings. We believe that there are a number of implications here which are positive as well, in the sense that they suggest extensions of experimental work within the terms of several current approaches to attitude change.

However, we must first broaden our notion of the significance of the non-attitude, and we shall do so by localizing it as one polar zone on a dimension describing a relationship between the subject and the potential attitude-object. This dimension is intended to be conceptually independent of the primary, positive-negative continuum indicating direction of affect toward the object.

For convenience we shall call this dimension one of centrality of the attitude-object to the subject.[5] Objects of non-attitudes lie at the extreme of low centrality. At the opposite pole, objects which are extremely central to the subject might well be called, in accord with the Sherif-Cantril (1947) usage, "ego-involved."[6]

Some such dimension has received only moderate lip service and remarkably slight experimental attention. Perhaps its most notable experimental use has come in the tradition carried on by Sherif, Hovland, Harvey and others, with particular note due to their efforts to influence attitudes toward alcohol and prohibition held by individuals with deep convictions on the subject in "real life" (Hovland, Harvey, and Sherif, 1957), and the discovery of a marked contrast in susceptibility to influence as a function of the amount of ego-involvement on the issue. Similarly, Rosenberg (1960) has used "issue interest" as an explicit variable in his experimental design.

These are, however, among a small group of exceptions. Hovland (1959), despite his own appreciation of the importance of such a dimension in determining the course of attitude change, remarked with some rue that "the whole concept of ego-involvement is a fuzzy one," and called for further theoretical work on the matter. Or again, Festinger's primary statement of dissonance theory features as a crucial variable the "importance" of relevant cognitive elements to the subject, or what we would call the centrality of their referents for him. Yet as Festinger's discussion moves toward experiment, "importance" becomes not a manipulated property of cognitive elements *per se*, but rather an attribute of a more global decision-event. Even in this second sense relatively little has been done with the variable experimentally, as Brehm and Cohen (1962) have pointed out; and we find the relationship between the second sense and the first quite obscure.[7]

It does seem true as Hovland implied that notions of ego-involvement *per se* do not hold out easy empirical "handles" for the experimentalist, and perhaps this is why there has been no greater spate of experimental activity commissioning such a dimension as an explicit variable. But one of the reasons why we wish to suggest the conceptual dimension of *object-centrality* is that we believe that there are numerous handles for treatment which have been overlooked in the past.

The dimension of centrality has two faces, as we shall deal with it. The first face is the more familiar one: centrality in a purely motivational sense. This is the flavour carried by the term "ego-involvement," and summarizes also the stabilizing characteristic of attitude structures stressed by the "functional" approach to the dynamics of attitude change by Katz and others (e.g., Sarnoff and Katz, 1954; Katz, 1960). In this latter context, the motivational centrality of a potential attitude-object has to do with the degree to which the object gears into the primary goal or need-structures of the individual. Actually, the functionalists are less likely to speak in terms of "degree" than in terms of the "*way*" in which attitude-objects are linked to diverse underlying needs. While the modes whereby similar attitudes gear into different needs for different people is without question a critical specification in understanding case studies of attitude change, there remains hope that less particularistic statements about attitude change can rest on the more streamlined variable of "degree." However, any good operational work with such a "degree" variable would seem to require a capacity to identify, measure and reliably equate intensities of very disparate needs from person to person, as well as to establish some estimate of the tightness of linkage between the potential attitude-objects and such needs. There is little here that seems within clear operational reach, however accurate we may feel the underlying suppositions about reality to be. Too frequently, thinking in this vein is thrown into a reverse logic: since we cannot move this person's attitude, it must be anchored in some deep underlying need, although we are at a loss to know just what that need may be; or we begin to engage in rather free-floating imagination to "explain" what the need is. Difficulties of this sort are what we have referrred to above as a dearth of empirical "handles" in dealing with the motivational face of the centrality dimension.

The second, or cognitive, face of the centrality dimension has received less attention, although sample-survey data recommend it to our attention. The subject has, however, been adumbrated in a number of comments made by the authors of *Attitude Organization and Change* (Hovland and Rosenberg, 1960), where the assumption is noted on several occasions that sheer "amount of thinking" about particular objects of attitudes has the effect of increasing consistency among attitude components, and some demonstration of the phenomenon seems to occur as a by-product of an experiment by McGuire. Of course, the absence of any significant amount of prior thinking about a potential attitude-object is one of the hallmarks of what we have styled as the "non-attitude," and attempts to measure investigator-shaped components of such psychological states should produce data of maximal inconsistency. In any event, cognitive centrality may ideally be taken to refer to the proportion of "mental time" which is occupied by attention to the attitude-object over substantial periods. In its dependence upon notions of "thinking" or "the forefront of consciousness," it is very close to what is often referred to as the salience of an object, although we avoid this word because of its customary short-term connotations of stimulus-bound and transient arousal. It is only the set of objects which are persistently salient for the actor over substantial periods of time, within or outside of the physical presence of the object itself, which are defined as cognitively central.

Given this "ideal" definition of cognitive centrality, measurement may seem fully as difficult as it was in the case of motivational centrality. Unlike the motivational case, however, some of the reliable concomitants of cognitive centrality (excluding the tautological concomitant of attitude stability) are very easy to measure. That is, we see as reliable concomitants of cognitive centrality at least the following:

1) heightened attention to the object, leading to alertness in singling out information about the object from the total flow of information with which the individual is bombarded in the environment;

2) an increased probability of successful storage or retention of information about the object, leading to increased differentiation of cognitions about the object and the sheer volume of relevant information stored;

3) an increase in the number of "active" associative bonds which tie this object with other cognized objects through various types of "linking" information.

Among these concomitants, perhaps the sheer amount of ancillary information held by the subject with respect to the object (beyond information conveyed by the attitude item itself) lends itself most readily to rapid measurement. Indeed, although there are implicit sampling problems in the choice of items, survey studies suggest that the natural "range of talent" in stored information about the kind of "external" attitude-object under discussion here is likely to be extreme even within populations which are relatively homogeneous on other grounds. This means in turn that measurement of level of information about the object is likely to be far easier and more reliable than the measurement of the attitude itself.

Given this ease of measurement and the seeming importance of centrality in attitude change, it is remarkable indeed how far one can look in the experimental literature for such an explicit variable. It is true that quite recently Rosenberg and Abelson (in Hovland and Rosenberg, 1960), in an excellent discussion of the microprocesses of attitude change, referred to the crucial role played by the "cognitive files" of information relevant to the attitude-object that the subject possesses, and there is of course a perfect identity between the volume of these files and what we are calling the amount of information stored about an object. However, even in this sequence of experiments it does not appear that this variable, or anything much like it, has actually been studied.

It seems a shame to leave the matter here. Somehow when we read an experimental study involving the manipulation of attitudes toward some object of controversy—say, abstract art—we have a great thirst to know which of the experimental subjects were art majors and knew a good deal about the properties of abstract art, as opposed to the experimental subjects who were rather vague as to just what abstract art was, and certainly innocent of the details of the controversy as it may rage in aesthetic circles. Our frustration is compounded by the reflection that an information scale of one, two or three items could make such a discrimination with deadly accuracy.

Ideally, of course, we would want to know not only how *much* information the individual holds, but more precisely *what* that information is in the specific case, for this in a functionalist vein not only gives a better picture of the phenotypic attitude structure, but also would suggest the kinds of information which would be most effective in changing it (see Rosenberg, 1956). However, this approach leads to the same complexities and phenomenological particularities which make notions of motivational centrality difficult to deal with in any clean operational way. Furthermore, it is our belief that attention paid even to such an elemental and simply-measured variable as gross differences in amount of stored information with respect to the object would lead to a large number of important and experimentally verifiable generalizations in the area of attitude change. While all of these generalizations would revolve around the core proposition that attitudes toward objects of high centrality are harder to change, *ceteris paribus*, than are attitudes toward objects of low centrality, many sub-propositions could be tested which would advance our understanding of the mediating processes.

Take, for example, the notion of incredulity, with which all students of the persuasion process have been obliged to deal in one way or another, at least as an appended correction function. While it is undoubtedly true, as many have observed, that the location of a point of incredulity in any given case depends in some degree on the distance between the individual's own attitude position and that represented by the message, it seems likely that the location of such a point must depend as heavily upon the amount of information held by the subject with respect to the attitude-object prior to the receipt of the message. Thus the impact of the message should be greater if it contains information new to the subject, rather than old and incorporated (or subjectively discredited) information. And of course the probability that *any* given message will contain arguments new to the subject is logically the perfect inverse of the amount of information he already holds relative to the object. Similarly, the subject's implicit estimate as to the reliability of the sender as a source of information would seem likely to vary widely according to the ratio of own stored information to the apparent amount of information which the sender possesses on the subject. Even an opinionated amateur may feel obliged to credit the contrary information of an expert; an expert may discredit new information transmitted by an obvious amateur even though it implies support for some of his predispositions. In short, then, the amount of information the individual has stored about the object prior to the persuasion attempt should profoundly modify the course of such a process, and the introduction of such a factor in experimental studies would help illuminate the phenomenon of incredulity as well as the more general aspects of susceptibility to attitude change.

The Rosenberg-Abelson discussion (in Hovland and Rosenberg, 1960) of microprocesses underlying the redress of cognitive imbalance, resting as it does upon a procedure of "search" through the "cognitive files," can yield a variety of other propositions worth experimental test with amount of information as a measured variable. And even earlier in the process, it is likely that the individual who is well-informed about an object of controversy can screen out channels of potential incoming information more effectively from "afar" than the more poorly-informed bystander. The well-informed person *knows* which newspaper columnists are so wrong they are not worth reading; the poorly informed person, while less likely to read any columnists, perhaps, will be unable to exercise such knowledgeable selection on the occasions when he does.

Much more generally, the whole notion of centrality, cognitive or motivational, would seem crucial for balance theories in predictions of *which* attitude will change (or will change *more*) when the familiar triangle of valences (such as between self-other-object or self-source-message) becomes imbalanced. When attitudes toward one or the other external object must change to restore balance, the less central of the two should undergo the primary adjustment. To some degree, the precise predictions of the Osgood congruity model can be taken as implying this very thing. However, this implication follows only if it be true that objects which are extremely central for any individual (e.g., "mother," "self") cannot be the object of ambivalent feelings, but are necessarily the objects of extreme, univocal attitudes. As indicated above, however, we are reluctant to accept this assumption, and would enjoy seeing what difference might occur in experimental results if there were greater assurance that the center of the Osgood scales relative to the remote objects used in attitude-change studies were restricted to the ambivalent, with non-attitudes discarded. However, the basic generalization about centrality and balance would seem an important one.

Without too many additional assumptions, it can be deduced from balance theory itself that when the attitude toward one of two objects must change, the object which is the less central (as we have defined the matter) is more likely to undergo change. This can be seen by taking account of the fact that objects involved in the balance triangle are also involved in cognitive associations with other objects as well. Festinger (1957) has suggested that a cognitive element which is dissonant with some other may resist change toward consonance because such change would create other dissonances. Since the more central the object, the more numerous and active are its associative bonds with other cognitive objects, it follows directly that change in attitude toward the less central of two objects will result in fewer new imbalances with further associated cognitive systems. Imbalance being unpleasant, such change will be preferred if any occurs at all.

In addition to the possibility of changing own attitudes toward one or another of two objects, such as a source and a message or another person and an object of mutual attention, a further mode of redressing imbalance involves some dissociation of the external objects or refusal to accept the immediate evidence which associates them in a way creating imbalance. The role which object centrality plays in this broader scheme of possibilities might well be the following:

Relative centrality of attitude objects *A* and *B*	Outcome of imbalance
1) *A* more central than *B*	Attitude toward *B* changes
2) *A* and *B* of equal *low* centrality	Imbalance not noticed or, if noticed, readily tolerated
3) *A* and *B* of equal *high* centrality	Some form of dissociation or denial of external evidence concerning *A–B* bond

The assumptions underlying this scheme are not particularly new, and indeed are at least implicit in most discussions of the subject. Dissonance theory posits quite directly that the subjective importance (or centrality) of the objects involved in the dissonant relationship not only guides the outcome of the attitude change, but also is a key determinant of the intensity of the discomfort which the dissonance occasions. Similarly, discussions of pressures toward the restoration of balance through attitude change often append the observation that these pressures presuppose that the objects are of some concern to the individual (item 2). It might be observed as well that most treatments of dynamic conflict which involve notable reality distortions are concerned principally with attitude-objects which are extremely central for the individual (the self, intimate associates, crucially valued objects), and hence operate almost exclusively in the domain defined by item 3.[8]

What we wish to stress, then, is less the novelty of the scheme than the gross discrepancy between the widespread common-sense assumption in discursive and theoretical work that some such dimension plays a prime modifying role in situations of incongruity or imbalance, and the scanty experimental attention paid to it.[9] We know very little in a systematic way about the truth of this bit of common sense. We know little as well about the sources of variation in centrality of the same objects for different people.

Nor, we may note in closing, do we know how flagrantly we have "cut corners" on reality in the foregoing discussion by dealing with cognitive and motivational centrality as though they represented a unitary dimension. The short-term defense for this tactic is a simple one: across the universe of possible attitude-objects for any given individual, the empirical correlation by object between the two types of centrality is certain to be very high. That is, experimental results, survey results and common sense join in suggesting that objects which are motivationally central to a person are likely to be cognitively central to him as well; and objects that are not central in one sense are not likely to be central in the other. Hence our blurring of the difference between the two types of centrality can suffice as a first approximation. But in the degree of fit between the two lie issues which seem of fundamental theoretical importance, yet which the survey analyst is ill-equipped to illuminate. Perhaps a resumption of our dialogue at a later date can include commentary on this score as well.

NOTES

1. All correlation coefficients cited in this paper are Kendall tau-betas (see Kendall, 1955) computed on square ($k \times k$) tables. For the benefit of the reader more accustomed to the Pearson product-moment coefficients, we might note that with non-pathological bivariate distributions, the Pearson coefficient computed on the same tables with linear scoring tends to give consistently liberal (high) estimates of the degree of association, relative to the tau-beta. The difference between the two estimates is slight (.02) in the low ranges (correlations of .1 or .2) but increases considerably in the higher ranges, so that a tau-beta coefficient of .75 might be interpreted as a Pearson coefficient of .85 or .90.

2. This rather reliable property has deeply impressed other survey analysts dealing with panel data. See, for example, Wiggins (1955).

3. It should be noted that the initial test was carried out in a preconceived format with little attention to the fact that the test might have been carried out in alternative formats. The alternatives depend primarily on the way in which the data may be collapsed to carry out the critical test. That is, the attitude turnover tables for t_1–t_2 must inevitably be reduced from their "raw" format (5×5 matrices) to a 2×2 format, with the tiny handful of "undecideds" dropping out of the middle row and column of the table, if one is to define the two subsets of the population necessary to make the test. The question arises as to what size tables should be employed in computing the t_2–t_3 rank-order correlations. Initially we computed these critical correlations for the two test groups on the basis of 4×4 tables, utilizing the maximal remaining information. The results of this test are the ones cited in the text. Later we ran the computations a second time, evaluating the t_2–t_3 correlations on the basis of collapsed 2×2 tables, rather than the 4×4 form. The results of this test are slightly less pretty, as the t_2–t_3 correlation for the random group slips somewhat farther off the "dead center" of .000 into the negative correlation region, although the departure is not great (−.05). Another version of the test, suggested by Wiggins (1955), also utilizes only the dichotomous responsibility over the three time points. Here one posits a three-class response model:

	Probability of a "plus" response
Class I	$.5 + a$
Class II	$.5$
Class III	$.5 - a$

Then one uses the configuration of three-stage data to solve for a. The issue item under discussion yields an estimate of .52 for a, which once again is essentially indistinguishable from our black-and-white model.

4. These felt difficulties are true not only of items requiring attitude locations, but also of measurements of a perceptual or personality type, which require a location of oneself, others, or other objects on some judgmental dimension. Once again, where the investigator's dimension is some very customary item in my own evaluative toolkit, such locations are easy to make, and I would feel that I could reproduce them at a later date quite accurately. If, however, the dimension involved is not one in terms of which I customarily locate objects, even though I may understand its meaning, the search process is more tedious, and I have less faith that I could reproduce my responses short of rote memory of the response marked before.

5. The usage here departs notably from that developed for the same term by Rokeach (1960).

6. Of course it is scarcely coincidental that the concept of ego-involved attitudes came to be emphasized in an earlier moment of collaboration between the sample-survey and experimental traditions. The fact that we have come upon the matter from the opposing pole—that of non-attitudes—should not conceal the fact that sample surveys were proffering the same indications about attitude change twenty years ago that we wish to develop in more detail here.
7. Brehm and Cohen discuss Zimbardo's (1960) experimental manipulation of "issue importance" as one of the infrequent examples. This manipulation clearly operationalizes decision-importance (second sense) and not element-importance. Indeed, in explicating the experiment, Brehm and Cohen are forced to rely on *unmanipulated* differences in amount of "prior commitment" to various presumed cognitive elements in the decision. And here momentarily their "prior commitment" means precisely what we mean by centrality, and probably Festinger's "importance" (first sense). (See Brehm and Cohen, 1962, especially pp. 58–59.)
8. We exclude from consideration here the apparent reality distortions demonstrable with objects of slight centrality which represent casual guesses in the absence of objective information rather than a more dynamic defense against a clear and present reality.
9. One of the few clear exceptions is Pilisuk (1962). Although the centrality dimension is not used as an experimental variable, the experimental hypothesis has to do with resistance to change of attitudes toward highly central objects after induction of some imbalance. Hence the experimental design lies within the class of situations defined by item (3) in the schema above. The results—that attitudes toward the key objects do not change, but that the subjects doubt the genuineness of the situation or rationalize in a variety of ways to "decouple" the imbalanced terms—fall generally within our expectations.

REFERENCES

Brehm, J., and A. R. Cohen, *Explorations in Cognitive Dissonance*, New York: Wiley, 1962.

Brim, O. G., "Attitude Content-Intensity," *American Sociological Review* **20** (1955), 68–76.

Festinger, L., *A Theory of Cognitive Dissonance*, Stanford: Stanford U. Press, 1957.

Hartley, E. L., *Problems in Prejudice*, New York: King's Crown Press, 1946.

Hovland, C. I., "Reconciling Conflicting Results Derived from Experimental and Survey Studies of Attitude Change," *American Psychologist* **14** (1959), 8–17.

Hovland, C. I., O. J. Harvey, and M. Sherif, "Assimilation and Contrast Effects in Reactions to Communication and Attitude Change," *Journal of Abnormal and Social Psychology* **55** (1957), 244–252.

Hovland, C. I., and M. Rosenberg, Eds., *Attitude Organization and Change*, New Haven, Conn.: Yale U. Press, 1960.

Katz, D., "The Functional Approach to the Study of Attitudes," *Public Opinion Quarterly* **24** (1960), 163–204.

Kendall, M., *Rank Correlation Methods*, New York: Hafner, 1955.

Lorge, I., "Gen-like: Halo or Reality," *Psychological Bulletin* **34** (1937), 545–546.

Newcomb, T. M., "An Approach to the Study of Communicative Acts," *Psychological Review* **60** (1953), 393–404.

Osgood, C. E., and P. H. Tannenbaum, "The Principle of Congruity in the Prediction of Attitude Change," *Psychological Review* **62** (1955), 42–55.

Osgood, C. E., G. J. Suci, and P. H. Tannenbaum, *The Measurement of Meaning*, Urbana: Univ. of Illinois Press, 1957.

Pilisuk, M., "Cognitive Balance and Self-Relevant Attitudes," *Journal of Abnormal and Social Psychology* **65** (1962), 95–103.

Rokeach, M., *The Open and Closed Mind*, New York: Basic Books, 1960.

Rosenberg, M. J., "Cognitive Structure and Attitudinal Affect," *Journal of Abnormal and Social Psychology* **53** (1956) 367–372.

Rosenberg, M. J., "Cognitive Reorganization in Response to the Hypnotic Reversal of Attitudinal Affect," *Journal of Personality* **28** (1960), 39–63.

Sarnoff, I., and D. Katz, "The Motivational Bases of Attitude Change," *Journal of Abnormal and Social Psychology* **49** (1954), 115–124.

Sherif, M., and H. Cantril, *The Psychology of Ego Involvements*, New York: Wiley, 1947.

Wiggins, L. M., "Mathematical Models for the Interpretation of Attitudes and Behavior Change: The Analysis of Multi-Wave Panels," unpublished doctoral dissertation, Columbia University, 1955.

Zimbardo, P. G., "Involvement and Communication Discrepancy as Determinants of Opinion Change," *Journal of Abnormal and Social Psychology* **60** (1960), 86–94.