International Institute for
Applied Systems Analysis
IIASA www.iiasa.ac.at

Attitude Theory: An Introduction for Physical Scientists and Engineers

Thomas, K.

IIASA Professional Paper
May 1978

Thomas, K. (1978) Attitude Theory: An Introduction for Physical Scientists and Engineers. IIASA Professional Paper. IIASA, Laxenburg, Austria, PP-78-003 Copyright © May 1978 by the author(s).

http://pure.iiasa.ac.at/921/ All rights reserved. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage. All copies must bear this notice and the full citation on the first page. For other purposes, to republish, to post on servers or to redistribute to lists, permission must be sought by contacting repository@iiasa.ac.at

PP-78-3

ATTITUDE THEORY:
AN INTRODUCTION FOR PHYSICAL SCIENTISTS AND ENGINEERS

Kerry Thomas

May 1978

Professional Papers are not official publications of the International Institute for Applied Systems Analysis, but are reproduced and distributed by the Institute as an aid to staff members in furthering their professional activities. Views or opinions expressed herein are those of the author and should not be interpreted as representing the view of either the Institute or the National Member Organizations supporting the Institute.

PREFACE

In our research on the social implications of technical development we have found the conception of 'attitude' a useful tool with which to investigate the nature of public preferences. However, by borrowing this concept from social psychology, we have also had to familiarise ourselves with theoretical and methodological issues surrounding 'attitude theory' in a wider sense. For the social scientists amongst us this is relatively familiar ground but the physical scientists and engineers have had more difficulty in placing the relatively simple model we have adopted within its broader context of attitude theory, and in acquainting themselves with some of the problems inherent in social research. This paper was written as introductory material to help our colleagues assimilate the attitude approach and thus promote useful interdisciplinary interactions.

ABSTRACT

This introduction to attitude theory explores why 'attitude' remains such a central issue in social psychology, and questions the assumption that attitude has a simple causal relation with behaviour. Models which relate attitude to behaviour and the considerable methodological difficulties involved in empirical work on this topic are discussed. The paper also describes the relations between attitudes and underlying belief systems both in terms of deterministic models and general principles of cognitive organisation and attitude change.

TABLE OF CONTENTS

INTRODUCTION

The Basic Questions

The study of attitudes has been of central concern to social psychology since its earliest days. 'Attitude Theory' is reviewed at length in every textbook on social psychology and investigations are carried out under the umbrella of 'attitude research' throughout the social sciences. The treatment of attitude theory and research in the narrower area of what we might call academic social psychology can claim to be a bit more rigorous, in the sense of defining its terms and setting up controlled experiments, but the assumptions that lie underneath the study of attitudes and the questions that we are asking and hoping to solve when we examine people's attitudes are the same across the board. In this paper I believe that it would be inappropriate to try and go into the details of specific attitude theories and the mass of experiments which support or contradict these different schools of thought. Instead I want to take as my general theme three basic questions: (1) what have we been trying to find out by studying 'attitudes'? (2) why have we spent so much time and energy on this concept 'attitude' somewhat to the exclusion of related ideas? and (3) can we perhaps remedy this by shifting attention a little more toward what underlies 'attitude', that is, toward what people believe (their cognitions)? This is an unorthodox way of approaching an introduction to attitude theory but I hope that by asking such fundamental questions attitude theory can be described in a simple, narrative way; and, further, that attitudes can be related to other psychological variables without losing the 'wood', which for me is the study of the social behaviour of socialised humans, amongst the 'trees' of terminologies and experimental data.

Let us begin with the following question: What is it that social psychologists were trying to explain or predict or even describe, when they took the term 'attitude' (stance toward) out of general usage and gave it a special role as a central concept in the study of social behaviour? The term attitude was not applied to some observable event; an attitude cannot be observed. It began as a hypothetical construct, which following the development of satisfactory measurement became an intervening, latent variable, that is, part of an imagined model which serves an explanatory purpose in a sequence between a cause and an effect. This particular 'model' was invoked to help understand and explain observed patterns or regularities in social behaviour, in an analogous way to the use of 'habit' by the learning theorist to explain observed patterning of behaviour in a more general sense. It is important to realise that in a great many respects the 'habits' of the stimulus-response school (the behaviourists) and the 'attitudes' of the cognitive-perceptual psychologist and social

psychologist are very similar ideas. Both are intended to represent 'residues of experience' which in turn act as dispositions to respond consistently to some object or class of objects. Combining these two terms we have the idea of attitude-as-an-acquired-behavioural-disposition. In this sense attitude was thought to be a precursor of observable behaviour and be capable of explaining and predicting that behaviour.

The Ecology of Attitude Research

Behaviourists, the majority of whom work with animals and experimental techniques, can control and design 'experience' for their animals and observe regularities in behaviour, but they cannot question them about perceptual experience or intended action. The behaviourist therefore tends to concentrate on the 'behaviour' end of the sequence and attribute regularities of behaviour to experiences, which in turn the experimenters can define through learning processes and the formation of habits. Cognitive and social psychologists, on the other hand, the majority (but not all) of whom work with articulate humans, try to explain or predict the behaviour of socialised individuals. Such behaviour is often social (that is interactions between individuals) and is usually behaviour with respect to a socially relevant object or group of people. Clearly the 'ecology' of this task is quite different. Usually the social psychologists have relatively little information about the history of the individuals' experience whether in general or specifically with respect to the object (the attitude object) in question; and frequently observation of the behaviours is difficult (i.e., very complex, expensive to study, hypothetical or 'yet to happen'). But the human subject can be questioned and this is the crux of the difference in emphasis in the two approaches. Social psychologists can interview or give questionnaires to their subjects, can find out what they think or believe about the topic in question, ask them factual questions about their experience, their past behaviour and intended behaviour in the future. Using this 'cognitive' material the psychologist tries to build up the individual's own picture or point of view of a topic in terms of the individual's own beliefs and feelings, and the relationships he perceives between the attitude object and other significant aspects of his world. In this way the psychologist tries to make some assessment of attitude toward the object; but even when he has some measure which he calls attitude, can behaviour be predicted, and can it be even partially explained? In the vast majority of cases the answer to this is an unequivocal 'no'; in other cases it is 'irrelevant', because the research was not using 'attitude' to predict or explain behaviour but studying beliefs and opinions for their own sake and with no necessary implication for overt action.

Much of the attitude research that appears in journals outside academic social psychology is of this latter kind. It may be concerned with areas of social science as far removed from each other as market research, politics, clinical psychology or anthropology and while, if carefully carried out, can supply a great deal of factual material, the interpretation of much of this 'opinion' data has little to do with 'attitude' or the prediction of behaviour.

In general the concept of 'attitude' has been overworked and misused and has really explained very little in terms of general laws of social behaviour. Even within more rigorous experimental work 'attitude' has become reified as an end in itself rather than a means to understanding the relation between the social experience of the individual, the intentions he forms and the social actions he carries out. Where 'processes' such as attitude change are concerned this concentration on attitude per se can be justified, and I briefly discuss attitude change in the third section below, but first and foremost the question of attitude as a means to understanding consistencies in behaviour must be explored. There is a natural progression from this 'behavioural' end of the process to the second topic of the paper which is an examination of the relation between attitude and inputs from the environment, that is, the means by which the history of experiences of the individual, his beliefs about the world, are organised and translated into attitudes.

ATTITUDES AND BEHAVIOUR

Attitude Measurement

I have managed so far to avoid any serious attempt to define attitude beyond 'residue of experience' and 'acquired behavioural disposition'. Neither of these is really a definition of attitude as it is used now. The most common definition is the 'three component' version according to which attitude is composed of affective (feeling), cognitive (belief) and behavioural components. This tripartite definition has justified measurement of attitude using questionnaires based simply on a collection of statements which appear to tap one or more of these three components. For example, questions directly about feelings, or items which are so emotionally loaded as to enable the researcher to infer 'feelings' from respondents' agreement or disagreement. The so-called cognitive component has usually been tapped by questionnaire items (belief statements) that lack any obvious emotional content, and the behavioural disposition component by questions about previous and intended behaviour or agreement/disagreement with more general statements about behaviour with regard to the attitude object.

Most 'attitude' questionnaires are little more than a cluster of questions which don't hang together, which is another way of saying that when the scores are added up or averaged they do not really mean a great deal since different questions are tapping different things. There are, however, several well established 'paper and pencil' techniques for measuring attitudes, and a few physiological ones which have not been well documented. Three of the most commonly quoted attitude measurement instruments are worth examining here for the light they shed on the concept of 'attitude'. First, there is a scaling technique based on the early work of Thurstone (Thurstone, 1931), who was one of the original attitude theorists and very much involved in questions of measurement. The construction of a Thurstone scale involves collecting or inventing a large pool of statements (or belief items) about the attitude object and asking a panel of judges to sort the statements into eleven categories. The basis of this sorting is what the judges feel to be the degree of favourableness toward the attitude object implied by each item. By calculating the 'average' placement for each item on the 11-point scale (represented by the 11 categories) its average, implied favourability for the attitude object (average for that particular sample of judges) can be represented by a scale value. When the questionnaire is administered to subjects they are asked to agree (score 1) or disagree (score 0) with each statement. By multiplying the scale value of each item by either one or by zero (in which case of course it drops out) an average score can be calculated for each subject which is, in effect, a measure of his favourability toward the attitude object. Thurstone in fact defined attitude in just these terms: "the degree of favourableness expressed toward the attitude 'object'". Thurstone also stated that the overall favourableness which someone might feel toward an attitude object would give an indication of the 'affective tone' of a range of behaviours performed with respect to that object, but would not necessarily predict specific actions.

The second attitude measurement technique, Likert scaling, (Likert, 1932) is based on a different procedure but also measures affect, i.e., the favourableness/unfavourableness of the object in question. Likert, however, puts more weight on the 'agreeing/disagreeing' response of the subject, that is, the subject's strength of belief. Here again the starting point is a pool of items, but in this case the psychologist himself (by some process of common sense perhaps, or from his own experience) assigns to each item a value +1 or -1 to represent implied favourableness or unfavourableness toward the attitude object. The subjects whose attitude is being measured are then asked to indicate the extent to which they agree or disagree with the statement on a 5-point scale. These scores are multiplied by either +1 or -1 (as previously assigned to each item) and the total favourability score is

simply the sum across all items. Once again the figure at the bottom of the page is an indication of affect or like/dislike felt toward the object.

The semantic differential technique (Osgood, Suci and Tannenbaum, 1957) is rather different from the two scales described above but again measures the affective component of attitude. The method by which the attitude is measured is fairly easy to describe. The attitude object is printed above a series of seven-point scales and the ends of each scale are labelled with adjective-pairs of opposite meaning such as good/bad or pleasant/unpleasant. The seven points on the scale represent gradations from one adjective extreme (e. g., extremely good) to the other (extremely bad). The subject is asked to 'rate' the attitude object by placing a mark on one of the seven points of each scale. What is more complex is the theory underlying this technique. This is the work of Osgood, an influential social psychologist, who has worked on attitudes, meaning and attitude change. Osgood was concerned, initially, with how people understand the meaning of concepts. He used the method just described to obtain ratings of many concepts in terms of a large number of adjective pairs, using a great many individuals from different cultures. He then used a mathematical procedure based on correlations between scores obtained for items (factor analysis) to identify the underlying dimensions which could account for observed patterns of intercorrelations. He repeatedly found that the same set of dimensions emerged and called them the 'dimensions of meaning'. They could be identified by examining the 'type' of adjective-pair associated with the dimension. Osgood also found that roughly the same pattern of relative importance amongst these dimensions was obtained for a given concept even when the respondents came from different cultures. The three most pervasive factors are 'evaluation', 'potency' and 'activity'; and of those the evaluative dimension regularly emerges as the most important. In other words, usually a large part of the meaning that any idea or concept has for us is a feeling reaction of goodness/badness - an evaluation. Osgood equates this with attitude. This evaluative dimension of meaning (i. e., attitude) is based on adjective pairs such as like/dislike; good/bad; wise/foolish; ugly/beautiful; pleasant/unpleasant. Clearly, there are close links between this measure of attitude and the overall feelings of goodness/badness tapped by Thurstone and Likert scales.

All of these scaling methods depend on some assumptions about what lies underneath attitude. For Osgood (as for just about everyone else) attitudes are learned, but he specifies in some detail how attitude or the evaluative aspect of meaning is attached to concepts as they are learned. The point I want to make is that the association of 'feeling' or attitude is construed as an automatic part of concept learning. As we build up beliefs and higher-order conceptions about the world

the 'feeling tone' or attitude is developed simultaneously. It derives from the feelings associated with the beliefs we are fitting together, and these in turn are based on our actual encounters with good or bad outcomes (reinforcements in behaviourist language) in our experience of the world. These experiences can be direct or learned through the mediation of language and other people's experience. When measuring attitude most of the beliefs that are used (the questionnaire items) are of the kind which link the attitude object, for example 'bus travel', with some attribute such as 'convenient' (which implies favourability) or 'slow' (which implies unfavourability). But some beliefs are statements of intention, for example 'I am going to make use of the bus service every day next week'. Clearly this statement can also be interpreted as indicating favourability toward bus travel, and a subject can be asked to indicate the truth, for him, of the intention statement. Here the intentional aspect of behavioural disposition is being used to measure attitude, but it is important to note that it is still the 'favourableness' which is being tapped. Although there was a movement away from Thurstone's simple definition of attitude as overall favourability (affect) toward the complex combination of affect, belief, and behaviour which I have already mentioned, which implies measurement of all of these aspects in order to characterise attitude, measurement has always in fact centred on the affective component. At present there is a tendency to return to a definition of attitude simply as overall affect or feeling toward some object; and this affect will, in turn, depend on the beliefs held about the object.

The Measurement of Behaviour

If we are to try and relate attitudes to behaviour, primarily with the purpose of explaining and predicting the latter, then there is not much point in refining techniques of attitude measurement without paying equivalent attention to a definition of the 'behaviour' we are interested in and how it too might be measured. This problem can be best appreciated with an example which is typical of the sort of study carried out in this area, namely the relation between attitude toward religion and 'religious behaviour'. Let us assume that the attitude has been adequately measured as overall favourability toward religion, now what constitutes the behaviour that one might expect to explain or predict? First, are we going to ask the respondent a question or a series of questions about his intention in religious matters, or about his actual performance of 'religious' behaviours in the past, or are we going to physically observe him as he goes to church or takes part in religious discussions, etc.? Second, are we going to try and predict some particular behaviour (where the level of specificity is quite different from the generality of the attitude) such as attending a particular church service on a particular day (single act) or the frequency with which he attends a particular service

(repeated observation). Or are we going to try and make some statement about his future religious behaviour in the sense of a 'pattern of different behaviours which somehow hang together', for example, frequency of attending church, and arranging religious instruction for his children, and donations of money to his church, and the number of religious books he has in his house, and so on (multiple acts).

The problem with such a multiple-act criterion is that not only do we have to measure many different aspects of religious behaviour but that we also have to find some logical way of combining the scores for each aspect. We might well find ourselves in effect adding 'oranges and apples', our units would not correspond and the different behaviours may well not hang together in any way. It is possible, however, to subject a number of such behaviours to procedures that are the same as those used to construct an attitude scale from belief statements, and the final result is in effect an attitude scale based on behaviours instead of beliefs. Notice that here we have turned the problem around on itself. We have now made a scale to measure (or predict) a general attitude from a large number of behaviours. Further, the implication of 'having a scale' is that there is a clear relation between a general attitude and behaviour, but behaviour in this one sense - a coherent pattern based on many different aspects of behaviour. But the initial problem was phrased in 'the other direction'. Now when two measures relate in some simple way (that is correlate) one can, in theory, make predictions in either direction; but in practical instances this reversability does not always solve a given problem. If we want to know about the religious behaviour of some group, then clearly the effort we have put into constructing a 'multiple act criterion' or a behavioural scale of attitude for that group will tell us a great deal about that behaviour without any need to measure the attitude as well; if on the other hand we just measure the attitude toward religion for this group of people and try to read off what it means by using a religious behaviour scale that had already been constructed on other individuals we still have (as well as the question of whether the two groups are equivalent) the problem of exactly which behaviours will be likely to occur from among the whole set - since the demonstrated relationship from which we are trying to make our prediction is between a general attitude and a total score across a large number of specific behavioural items.

The Relation Between Attitude and Behaviour

Quite recently there has been a long-neglected reassessment of the nature of the relationship between attitude and behaviour and several reviews of just how successful, or in this case unsuccessful, the research in laboratory experiments or 'real life' has been in demonstrating the prediction of behaviour from attitudes (see for example Wicker, 1969; and several other papers

reprinted in Thomas, 1971). For example, beginning with the study that is always quoted, that of LaPierre as long ago as 1934 (LaPierre, 1934) we find that he failed to predict the acceptance of Chinese patrons by hotels and restaurants from attitudes (or rather, what he called attitudes) to Chinese people. In another example, this time a typical laboratory experiment, De Fleur and Westie (1958) asked subjects who were either high or low in prejudice (negative attitude) toward Negroes to indicate their willingness to pose for a photograph with a Negro person of the opposite sex. They were then shown a graded series of seven 'photographic release statements'. These differed in the amount of publicity that could be given to the photograph. Subjects were allowed to sign as many (or as few) of these as they wished and this was taken as an indication of behaviour. In this case a low, but statistically significant relationship was found. Variations on this experiment have been carried out several times, usually with no significant relationship demonstrated between the attitude and the behaviour. In more applied research there has been only a very small degree of success in predicting behaviours such as absenteeism from work, work performance, or dropping out of training programmes from measures of attitudes to job or training programmes.

The reassessment of this work has suggested several reasons for the poor results, and these explanations fall into two main groups. On the one hand, there are what are essentially problems of measurement; and on the other, problems which relate to the basic theoretical assumption that attitudes and behaviour are indeed related. First, the question of inappropriate measures: The reviews of research in this area show quite clearly that in many cases the so-called attitude measure was in fact not a measure of affect in the sense used here, indeed it has ranged from measures of personality to measures of intentions (for example, LaPierre's study assessed attitude by asking 'would you accept members of the Chinese race in your establishment?'). The behaviour measures also varied considerably. Some were intentions rather than actual behaviour (e.g., the photograph release experiment) and the remainder were usually either single act criteria (e.g., LaPierre) or repeated observations, for example, Newton & Newton (1950) examined the relationship between attitude to breast feeding (actually an intention was measured here too) and observed behaviour in the sense of the amount of breast milk received by the baby over six feeds. As might be expected from the earlier discussion of attitude and behaviour measurement, the more successful studies were those which measured attitude as overall affect to the object and predicted behaviour using a multiple behaviour criterion.

There are, however, other methodological problems which have not been mentioned so far and which shade into theoretical questions. The choice both of the object of the attitude and the object of the behaviour can vary considerably along a dimension of generality/specificity, e.g., attitude expressed to a minority group in general or to a particular sub-group;

and similarly, behaviours can be with respect to the group as a whole or (and this is more likely in practice) to specific members of the group. As one might expect, correspondence between the levels of specificity of the attitude and behaviour tend to improve demonstrations of a relationship.

Similar questions of correspondence bring us to a more fundamental question of whether we should be using general attitude objects while trying to predict behaviours which are usually carried out with respect to some fairly specific object and certainly carried out (inevitably) in particular circumstances which must contribute to the specificity of meaning of the behaviour. Following this line of argument, Fishbein has suggested (1967) that it is more appropriate to predict (and try to understand) particular behaviours by measuring, not the attitude to the object of that behaviour, but instead the attitude to the behaviour itself, that is feelings of favourability toward the actual performance of the behaviour in question, usually carefully defined in terms of the situation in which it will occur. For example, he suggests that it is more appropriate to predict the 'use of public transport for the journey to work in the next month' from the attitude toward 'use of public transport for the journey to work in the next month' than from attitude toward 'public transport'.

This important insight has considerably improved behavioural prediction in many experimental and 'field' situations ranging from 'gaming techniques' in a social psychology laboratory to consumer choices of toothpaste brands, voting behaviour in USA and Great Britain and the off-peak use of suburban bus services.

Further, and this is crucial for the development of my argument in the rest of this paper, the shift to attitude-toward-the-behaviour also means a shift in the sort of cognitions (perceptions and beliefs) which underlie the attitude and therefore which influence the behaviour. Instead of considering attitude as dependent on general attribute beliefs, as in the case of attitude-toward-an-object (for example, 'Candidate X is a Liberal'; 'Candidate X is in favour of spending more money on the National Health Service') we now are concerned essentially with outcome beliefs. These are beliefs about the expected or likely (and note that this use of expectation or likelihood is an indication of the strength of the belief) consequences of performing the particular behaviour in the situation specified. To extend the example above, if voting behaviour was being predicted from attitude to Candidate X (that is, attitude to an object), beliefs like those above might be examined. According to the view taken here, however, it is more appropriate to predict voting from attitude toward the act of 'voting for Candidate X' and here the beliefs might well be quite different and imply a different overall direction of favourability which in turn might imply the opposite voting behaviour. For example, someone with the attribute beliefs described above

and who is a Liberal and thinks that spending money on the National Health Service is 'good' might be expected to have a favourable attitude toward Cadidate X. But, if his outcome beliefs about voting for Candidate X are 'voting for Candidate X will mean throwing away my vote since the Liberal Candidate has no chance in this constituency'; or 'voting for Candidate X will increase the chance of the Labour candidate winning whereas I would rather see the Conservative win (given that the Liberal cannot)', then one might reasonably expect him to have a negative attitude to voting for Candidate X and to vote for someone else.

The topic of underlying beliefs is treated more fully in the third section below; but to complete my discussion of attitude and behaviour I must first return to the other source of difficulty in predicting behaviour from attitude. This is less concerned with how and what is measured and more with the theory underlying the relation between attitudes and behaviour. I want to examine this in the more general context of the precursors of overt behaviour and the formation of intentions.

The Formation of Intentions

Up to this point I have treated intentions simply as either private beliefs about a projected course of action or a public statement of this intent. Assuming for the moment that the intention is formed and/or stated close in time to the behaviour and that no unexpected events intervene to prevent the realisation of the intention, then it follows that, insofar as attitudes are considered as determining behaviour, attitudes can also be treated as determinants of intention. However, the previous section has shown that it has proved very difficult to establish a simple relationship between attitude and behaviour. This is in part due to measurement problems, but several social psychologists are now beginning to take account of 'other variables' which either influence attitudes or act independently of attitudes to bring about the observed inconsistencies between attitude and the behaviour.

The approach described above which has had most success in predicting behaviour from attitude-toward-the-behaviour-specified states that a second factor is also involved. This second factor depends on the individual's beliefs about the expectations of other people who are important to him (significant referents). These often (but not always) influence the formation of an intention over and above the effect of the attitude; and by measuring those so-called 'normative beliefs' and including them alongside attitude-to-the-behaviour a better prediction and a better understanding of intention (and behaviour, providing no unexpected events intervene) can be achieved. This theory also suggests that the relative influence of the attitudinal and normative factors will vary with the sort of behaviour involved. For example, it has been shown that the choice of toothpaste brand is,

to a large extent, under the influence of normative pressure from the expert in question, that is, the dentist, whereas it has recently been shown that voting behaviour in Great Britain is barely influenced at all by the expectations of important social referents. Here, then,we have a clear example of a theoretical position supported by empirical evidence that intentions (and behaviour) do not necessarily have a simple one to one relation with attitude alone.

Other theorists who have continued to use the traditional 'attitude-to-the-object' measure have also considered the effects of additional variables on behaviour. These are often discussed when experimental prediction from attitude alone has failed and so far there has been little systematic research on these 'other variables'. Rokeach and Kliejunas (1972) used, in addition to attitude-to-the-object, attitude-to-the-situation, as a partial cover for several 'other variables'. Here the behaviour studied was students' self-report of missing lectures (other than for reasons of health, bad weather, etc.). Attitude-to-the-object referred to the 'liking for the teacher' and attitude-to-the-situation to 'importance of attending the lecture'. In this case the behaviour was found to be only related to the 'other variable', that is, the 'attitude-to-the-situation', and not to the liking for the teacher. In the light of the earlier discussions of inappropriate measures this finding is not unexpected. More recently a study of donating bone marrow in hypothetical situations has indicated that 'personal norms' or 'moral obligations' may well have some influence on the formation of intention, independent of attitude and more general norms.

The model of intention formation put forward by Fishbein is quite clear that intentions depend only on attitude-to-the-behaviour (based in turn on outcome beliefs) and a general social norm (which can be split up into more specific normative beliefs referring to different referents or sources of pressure) and that all other variables must act through one or other (or both) of these primary determinants. For example, some unusual aspect of the situation might have its effects through beliefs about the outcomes given that circumstance and hence act through the attitudinal variable. Similarly, a sociological variable such as social class or a more personal factor such as age might have its influence through the normative variable via the choice of referent whose pressure was being complied with; or via the exact nature of the normative belief, for example a given behaviour may be perceived, through the expectations of the referent, as appropriate only to a particular age group. There is much empirical evidence to support this model (e.g., Fishbein and Coombs, 1974; Fishbein, Thomas and Jaccard, 1976) but the formation of intentions may well be dependent on other factors that are reflected in these two variables, but which would improve our understanding if we examined them separately. While this is not the place for speculation it seems to me that the role of pressure from others is far more complicated than

would appear from this model; and also that it is largely because the model has been used to examine fairly simple behaviour in well defined contexts where little or no conflict is experienced in the formation of the intention that a high level of success has been achieved with just attitude and simple normative beliefs. Further, the expectations of important others will act as a mirror and for example, will reflect the physical possibilities or impossibilities inherent in a situation. The question of general physical constraints and other external criteria on behaviour and the role of volition (the extent to which the individual can choose which behaviour to perform) may well turn out to contribute to the formation of intention and the performance of purposive behaviour.

Two final points are relevant here and link this section to the remainder of the paper. First, so far my discussion of attitude theory has been in deterministic language, as though the individual behaves automatically in a way which is 'controlled' by the 'residues of his experience' and perhaps the additional effect of social pressures from important others. But if we move toward a conception of man forming intentions which represent plans of purposive, goal-seeking behaviour, then our approach to the role of attitudes and other factors in the formation of intentions must be somewhat different. We must visualise the individual processing information, defining situations and making choices and decisions in terms of some ultimate goal; and it is this view of intention formation which becomes more salient as the complexity of the beliefs and cognitive structures of individuals are explored. Second, 'attitude', whether toward an object or a behaviour, is essentially the overall feeling associated with some object/behaviour because of the beliefs held about that object/behaviour at that moment. Clearly attitude objects will differ enormously in the complexity and the stability of the beliefs on which they are based, but attitude is still a sort of evaluative or affective summary of the cognitive structure associated with the attitude object. The use of this 'summary', particularly in restricted circumstances such as an attitude toward a very specific behaviour, can be rather uninteresting and contribute little toward understanding social behaviour. By using this summary a great deal of information is frequently lost. In the following section I want to redeem the earlier simplification of 'residues of experience' and the emphasis on the output side of the equation and look instead at the cognitive (input) side of attitude theory.

BELIEFS AND ATTITUDES

The Relationship Between Belief and Attitude

As the individual develops he learns of the relationships between objects in his environment and builds up a version of the 'outside' world which, insofar as it is based on clear perceptual inputs and direct experience, will largely be veridical (true to life) because it is based on rational processing of information. It may, however, be constrained by physical or social factors which limit his learning experience. He will also receive information directly from other people and will further develop his cognitive structure by inferential processes based on the beliefs he has accumulated. He thus builds up belief systems which enable him to 'go beyond the information given' in a particular situation and make predictions about objects, events and relationships which he cannot directly perceive.

The simplest way of treating this topic is to assume that, through learning processes, we build up a vast number of beliefs which relate objects in propositional form (e.g., the cat is white), and that the strength of such a belief is related to the strength of the learning process. We can treat this strength of belief as a probability of the two objects being related in the stated way and assign to it a value between 0 and 1, for example, a probability of .90 that 'the cat is white' might be an accurate measurement of a perception of an extremely dirty white cat . Continuing with this simple example, as more information is received either by observation, or verbally from the owner of the cat, or by inference using our previous experience of the range of colours associated with cats, this belief may well become stronger and reach virtual certainty. In this example, the attribute 'white' does not carry a great deal of affect (feeling) but we can consider further beliefs about the cat amongst which might be 'the cat is angry' and the 'cat is vicious'. These are concepts which we have probably learned in the past in association with unpleasant experiences and thus carry negative affect (a negative attitude). The general model I am describing here suggests that the negative feeling associated with such concepts will become associated with this 'cat' to the extent that we believe that it is indeed angry and vicious. Where the belief is weak, a small amount of negative affect will be associated with the cat, and where the belief is strong then this will 'weight' (a multiplicative relationship) the amount of negative affect that is associated with the cat. It is in this way (although the exact mathematical relations are not clear) overall affect or attitude toward concepts is built up. Attitude toward the cat will depend on the beliefs held about the cat; each belief will associate the cat with an attribute or other concept which implies some degree of evaluation (this can be neutral) and this evaluation will accrue to the attitude object (the cat) in proportion to the strength of the belief.

As learning proceeds a great many beliefs are formed about a great many objects and it seems likely that only some of these (frequently, but not always, those which are most strongly held), will be available at a given moment as descriptions of an event or attributes of an object. In other words, we can only process a limited amount of information at a given time and the beliefs about an object which we are attending to at that time are called the salient beliefs (probably not more than 8 or 9 items and frequently less). In the model described here these are treated as determinants of the attitude. The affect associated with these determinant beliefs is summed over the set of salient beliefs to give an indication of overall attitude.

This type of model, although with different terminology and different mathematical functions, underlies much attitude theory. The version used here is that associated primarily with Fishbein, whose behaviour prediction model was described earlier, but very similar models have been put forward by other social psychologists, perhaps most notably by Rosenberg (1956).

Attitude Change

Since attitudes depend on beliefs and beliefs are learned[1/], we assume that attitudes are learned and are the product of an interaction with our environment (including other people) and inference processes which, to a large degree, are also dependent on the information we have previously obtained. It follows that systematic differences in physical and social environments and experience are likely to lead (via differing views of the world to a greater or lesser extent) to different attitudes. And, insofar as attitudes (and other beliefs, such as those about social pressure) underlie intentions and behaviour, groups of people with different socialisation experiences will tend to have different behaviour patterns. Further, given the general assumption about the relation between attitude and behaviour, new experiences and new information are likely to have some effect on changing these behaviour patterns. A great deal of social psychology has been concerned with this aspect of comparative attitude research and with attitude change. However, all too frequently the content of the different belief systems has taken second place to the study of the attitude (the affective summary) itself, and the extension of such studies into observed behaviour change are rare and usually unsuccessful. Typical experiments in attitude change measure attitude (that is, agreement with a set of belief statements), then expose the subjects to information in the form of a persuasive communication, and then re-measure the attitude. There are virtually no in-

---

[1/] There is little evidence that attitudes can be formed by direct conditioning, most experiments of this kind have involved awareness, that is, the involvement of some beliefs, during the conditioning process.

stances (see Thomas and Tuck, 1975) of the monitoring of specific beliefs during attempts to change attitude. A change in attitude could be due to changes in those beliefs which have been shown to underlie the initial attitude and which may be strengthened or weakened by the communication. It could also be due to new items of information contained in the communication entering a revised salient set of beliefs, perhaps displacing an earlier belief directly and/or changing the affect associated with an earlier belief.

The more traditional area of attitude change research can be roughly separated into two approaches, the very influential Yale studies on persuasion and communication (see for example Hovland, Janis and Kelley, 1953), and more theoretical studies on attitude change which reflect models of belief and attitude organisation. This second approach is discussed later, in the section on 'Principles of Organisation and Change'.

The Yale programme was centred on examination of which situational and personality factors would increase or decrease the effect of communications designed to change attitudes and opinions (stated beliefs). In other words, how to maximise the persuasive process. To do this the group studied the effect on attitudes of different communicators (that is, apparent sources of the message) with different levels of expertise, status, trustworthiness, etc. The messages were also varied, some presented only one side of an argument and others both sides; some set out to evoke emotional responses and some presented rational arguments. The audiences too were systematically studied and changes in attitudes and opinions were related to personality measures, self-esteem, cognitive complexity and many other traits.

The model of persuasion underlying this programme is very much based on the attention paid to the message and the learning of its content, followed by a second process of acceptance (or non-acceptance) of what has been learned. Only when acceptance occurs can belief change follow. However, with this two-phase process some of the experimental variables can end up having no appreciable effect overall because they differentially affect each phase. Since the Yale work, McGuire (1969) has pursued this as an explanation for many of the inconsistencies in the results of this programme and later similar work. For example, he has suggested that reception of a message (attention and comprehension) will be greater for high intelligence subjects but that acceptance of (or yielding to) the message will be less. He therefore concludes that, although both reception and yielding have linear relationships with intelligence, a combination of the two produces a non-linear effect such that subjects of both low and high intelligence will show less opinion change than those of intermediate intelligence.

The clearest outcome of the Yale programme and more recent work in the same tradition is that the original questions of who (communicator), what (messages), to whom (audience) have produced complex results which at best can be partially unravelled by looking at interactions between effects rather than simple straightforward effects, and which at worst are just inconsistent. And this muddled picture applies equally to later work in the Yale tradition which has tried to follow the effect of different messages, notably frightening versus non-frightening communications, on changing actual behaviour in areas such as dental care, anti-tetanus protection by innoculation, and giving up smoking. Obviously changing attitudes and behaviour in these sorts of areas is socially important, but the experiment failed to answer questions about how best to effect change. In some cases high-fear messages had least effect and this was attributed to arousing too much anxiety for the message to be remembered; in other cases, especially when a clear-cut course of action to avoid the frightening consequences was also provided (for example, information on exactly how to get anti-tetanus injections and thus avoid the consequences), then the high-fear message had most effect.

This last finding ties in with the approach to attitudes, intentions and behaviour described earlier, that is, when the message is directed at beliefs about the behaviour, as in the anti-tetanus experiment, then more behaviour change should occur. Only if a smoker can be persuaded that cutting down or giving up smoking will actually reduce his chance of lung cancer will he change is behaviour. If he just is told that smoking and lung cancer are related, with no indication that giving up smoking will have any good effect, then he will be much less likely to give up smoking. This has been demonstrated in an experimental study. If we want to change behaviour then the message should be aimed at first changing intention, which in turn requires a change in the beliefs about the outcomes of that behaviour (or perhaps in the normative beliefs if these have already been demonstrated as an important influence on the behaviour). So far very little work has been done using this approach. There is, however, one study which changed the intention and the behaviour of alcoholics with regard to signing-up for a programme of treatment (McArdle, 1972). The persuasive messages were of two kinds, the first aimed at changing beliefs about the outcomes of signing up (good outcomes) or not signing up (bad outcomes); the second kind was in keeping with the earlier experiments on 'fear appeal' emphasising the bad effects of continued drinking. This fear message had least effect on changing both the attitude to signing-up for treatment and actual behaviour. The most effective message for changing attitude and behaviour was the one which linked 'not signing up' with bad consequences. Notice that, in the terms of the Yale tradition, this message also produced most anxiety and least recall of the message, but showed a belief change (compared with control subjects who received no message at all)

in the direction of the message, whereas the fear message did not show a belief change.

Continuing with this theme of attention to the beliefs underlying attitude and their effect on behaviour, a rather different approach to changing a behaviour such as smoking is the technique where the subjects actively participate either by taking part in lectures, by giving speeches supporting 'the other side' (counter-attitudinal) or by role-play, for example taking the part of a heavy smoker who has just been told that he has lung cancer. One explanation of the relative success of participation in changing behaviour is that it forces the subject to re-think the topic and bring into mind beliefs which were previously non-salient; it also forms new beliefs by direct experience of a novel situation. It is likely that the more involved the subject becomes in the role-play (either emotionally or in the sense of actually taking part rather than observing someone else acting), the more belief change and hence behaviour change will occur.

Similar explanations can be applied to 'counter-attitudinal' behaviour, which is debating on the side which you do not support, or consenting to carry out some behaviour that you do not agree with, either for payment, or to please someone, or because you have been 'forced to' (see Festinger and Carlsmith, 1959). But there are other explanations of the changes in attitude (and beliefs and sometimes behaviour) which occur in such situations. One of these is that the subject 'observes' himself doing or saying something that he disagrees with and, in effect, says "If I am doing this, and there is no other good reason for me doing it, then I suppose I must believe in it or like it" and changes his attitude accordingly. But this 'self-attribution' will depend very much on the 'other good reasons'. If he was forced against his will to produce counter-attitudinal behaviour then he has sufficient reason for having done it and need not think it due to some personal disposition. In such instances, in theory, he will not change his attitude. If, on the other hand, he cannot think of a good enough external reason he may conclude that it does in fact reflect his own opinion and thus change his self-attribution and overtly state an attitude or belief which reflects a change.

This 'self-attribution' explanation was put forward by Bem (1967) in opposition to Festinger's influential theory of cognitive dissonance (Festinger, 1957). According to cognitive dissonance theory, changes (if any) in attitudes or opinions which follow 'counter-attitudinal' statements or behaviour are due to an experience of pressure toward consistency. If we believe one thing and do or say something different, according to Festinger, we experience an uncomfortable state of cognitive dissonance and try to remedy this by making a change; and if we have already said or done something overtly that is dissonant with our initial belief or attitude, then the easiest way to

reduce dissonance is to bring this cognitive element into line by changing it. The idea of cognitive dissonance or rather the idea of maintaining cognitive consonance has been considered an important principle of the organisation of cognitions, attitudes and behaviour and has produced a great many experimental studies. Although the details of these and the criticisms they have been subjected to are beyond the scope of this paper, cognitive dissonance theory is briefly discussed below as one of the consistency theories of cognitive organisation.

Belief and Attitude Systems: Principles of Organisation and Change

I have very briefly outlined, in an earlier section, a model according to which information about the world and about transactions with it can be represented in propositional form and subsets of these propositions aggregated into more complex concepts whose substantive content (the actual belief) can be summarised by an overall evaluation (attitude). This sort of model is compatible with both stimulus-response and cognitive views of psychology, but frequently is dismissed as being too simple to account for the complexities of human belief systems and behaviour. I want to make just two points on this subject: First, information processing in propositional form can be extremely complex by virtue of the large numbers of such units which the human brain can deal with, certainly over time. The capacity of the brain is almost always underestimated. Second, the real value of this 'oatmeal' view rests not only on the complexity that can arise from sheer numbers but also on the structures which can thereafter be imposed on that 'oatmeal' and the processes to which these structures can be subjected. This is obviously an extremely difficult subject and one we do not know a great deal about, but certain principles of cognitive organisation have and are being studied and I want to introduce some of these here. I shall first deal with general principles of cognitive organisation, and then the more formalised systems are explored.

As beliefs are learned the selection of information from the environment (including the social environment) is influenced by the concepts already formed; this is called 'prior entry' effect. We tend to pay more attention to topics which are of interest to us and about which we already have some knowledge or perhaps some anxiety. We tend, to some degree, to select information which hangs together and we build up interrelated beliefs and value systems into a 'view of the world' which, at the most general level, we call ideologies. These may reflect prior dispositions or even personality traits. This sort of view of cognitive organisation has been called a 'functional approach to attitudes'. It does not exclude more objective and rational processes but allows for considerable individual variation in the extent to which personality and motivation affect our belief systems. But more recently the emphasis has shifted.

There is now considerable evidence accumulating that the integration of information (beliefs or propositions) and the processes by which we use beliefs to infer or predict other beliefs can be fairly well mirrored by mathematical models of information processing, analyses in terms of conditional probabilities and formal logic for which mathematical functions already exist. These models do not provide an exact fit, an obvious example is that logic is apparently influenced by "wishful thinking' to some variable extent, but if we think of beliefs in this way and measure them as probabilities attached to propositions (i.e., the likelihood that A is a B), then we have at our disposal a means for considerably enlarging our understanding of belief systems.

The assumption which underlies the possibility of mathematical analysis of belief systems is that, to the best of our abilities, we process information in rational ways, learning the probabilities of interrelations between concepts and making logical inferences. For example, when faced with a decision we make a probabilistic estimate of the outcomes of each alternative. A further example refers to the attitude scales described earlier. When faced with a questionnaire made up of a list of belief statements (for example about attributes of a political party) some will be beliefs we already hold and can easily be rated for agreement/disagreement on the basis of what we actually know. But when faced with an item that is new, then a decision to express agreement or disagreement will involve a consideration of other beliefs already held about that political party and inference will be based perhaps on an 'implicit theory' about the kind of attributes that probably 'go together' in the political area. This response will reflect 'probabilistic' (rational) consistency. But if faced with an item about which we know absolutely nothing, then we have no alternate (other than a random response or no response) but to make use of a different principle of organisation, namely evaluative consistency. In other words, we will rate the item in terms of (or based on) our existing feelings about the political party. If we like the party (positive attitude) and the item indicates something good then we will tend to express agreement.

The idea of consistency has been extremely important in social psychology and several theories, notably Heider's 'Balance Theory' (Heider, 1958), Abelson and Rosenberg's 'Affective-Cognitive Consistency Theory' (Abelson and Rosenberg, 1958), and Festinger's 'Cognitive Dissonance Theory' are based on some notion of internal consistency amongst beliefs and attitudes and amongst higher order 'subsets' of cognitions. Consistency is a difficult concept to define. Previously I have talked about beliefs 'going together' and in very simple language this does get close to the notion of consistency. Consistency can be thought of as a condition in which one is able to make some prediction or inference (possible involving several steps) from one or more of the elements (beliefs, attitudes, etc.) concerning other elements in the system, or about the structure

of the system itself. Probabilistic consistency allows one to predict relationships on the basis of learned probabilities of association; logical consistency enables one to predict on the basis of a set of formal rules; evaluative consistency suggests inferences based simply on 'good things going together' or being associated with other good things, bad things being somehow associated with other bad things, and good and bad things being disassociated.

I cannot possibly do justice to the complexities of consistency theories, but very briefly, Heider's main thesis can be illustrated by the following example. If I see that Mary (whom I like) likes Jane (whom I don't know) then I will like Jane. This is balanced triad; there is Mary and Jane and myself and three positive links. The unit between Mary and Jane could have been some other form of association, such as 'similar to'. And we need not always be concerned with people only, thus 'I like Mary; Mary made a dress; I like the dress' is also a balanced triad. An unbalanced state would occur if, in this example, I did not like the dress. A balanced state is one where the relations among the entities fit together harmoniously and in this case there is no 'force toward change'; but when relations among entities (which can be translated into 'beliefs and attitudes') are not balanced then there is pressure to change. This is where consistency as a principle of cognitive organisation has implications for attitude change. If some new event (expressed as a belief or as a feeling) does not fit in a balanced way with existing beliefs and attitudes then according to several theorists, belief or attitude change will occur. Osgood has formalised this idea of 'imbalance leads to change' in a semi-quantitative way to predict the extent of attitude change that might be expected in given circumstances. The affective-cognitive consistency model of Rosenberg and Abelson is similar to balance theory but more developed. They have formulated the balance principle in a set of rules; not formal logical but what they call 'psycho-logic', and by applying these to relations between cognitions they claim it is possible to read-off predictions about what sort of cognitive change will occur.

Experiments have been carried out on all these versions of consistency theory and there is plenty of evidence that changes in attitudes and beliefs often do occur when imbalance is created; but it is difficult to say in advance which sort of change will occur and in which cognitive element. Also it is plain that frequently, if not usually, we all exist relatively happily with many imbalanced belief systems. It is difficult to conceive of an universal automatic pressure (homeostastis) toward cognitive balance which can only be avoided by 'not thinking' about the inconsistency.

The final consistency theory I want to mention is the one which has probably stimulated more research than any other single theory in social psychology: Festinger's cognitive dissonance theory states that "two (cognitive) elements are in a dissonant relation if, considering these two alone, the obverse of one element would follow from the other"; for example, if I believe that smoking causes cancer and I know that I smoke, then these two elements, which are dissonant, lead to an unpleasant or tension producing experience of cognitive dissonance. Festinger is more explicit than the other consistency theorists that dissonance will be greater when there are more dissonant elements and that, up to a point, dissonance will be tolerated. Only when dissonance exceeds some threshold will a change occur, such as giving up smoking or changing to a low-tar brand. Dissonance theory is supposed to be relevant to many situations. The 'counter-attitudinal' example has been described already; another illustration is decision-making where a choice has to be made between two desired alternatives. This produces the dissonant cognitions: 'I want X, (but) I have chosen Y'. The theory suggests that this dissonance can be reduced by either decreasing the evaluation of X, the not-chosen alternative, or increasing the evaluation of Y, the chosen one.

Dissonance theory has received a lot of experimental support but the complexities of more recent experiments make it difficult to interpret the many inconclusive results. While dissonance theory is similar to the other consistency models it should be noted that the cognitive elements can be confined entirely to beliefs, even though these are frequently beliefs-about-attitudes (I know I like smoking) or beliefs-about- behaving (I know I smoke). This suggests that Festinger's definition which centres on "the obverse of one element would follow from the other" (and no-one has really decided what obverse means), could well be compatible with 'probabilistic' and/or logical consistency. In other words, obverse might refer to "on the basis of any past experience unlikely to follow from ..."; or "in terms of strict logic unlikely to follow from". It seems then that dissonance theory may well turn out to be less a product of evaluative consistency than are other consistency models. This is in keeping with Bem's suggestion that cognitive changes which occur in so-called dissonance experiments can be explained by a self-attribution process based on beliefs and realistic self-observation. This explanation again emphasises the importance of information, inherent in all situations, and which can be rationally processed in propositional (or 'belief') form.

CONCLUSIONS

Any attempt to confine the whole topic of attitude theory, even at the level of an overview, to one relatively short paper means selection and omission and inevitable bias. But I hope

that I have at least made my bias explicit and that the (mild) elements of controversy introduced here will provide 'coat hangers' on which to hang the vast literature on attitude theory. To recapitulate, I began by asking two fundamental questions: What have social psychologists been trying to find out when studying attitudes? and why has so much time and energy been given to the study of attitudes somewhat at the expense of beliefs? I hope that the narrative approach I have adopted has nevertheless allowed the answers to these questions to emerge. Attitude theory was first conceived in the hope of explaining behaviour but met with little or no success. Meanwhile the 'attitude' concept itself became the topic of research producing studies of differential attitudes across groups, and a great deal of work on the process of attitude change. 'Attitude' as the focus of attention is partly a misnomer because frequently it was beliefs that were under investigation but in a non-theoretical 'loose' way, due largely to no real attempt at definition of belief, and poor belief measurement.

From the perspective of the present climate in social psychology, I believe it is evident that social and other applied psychologists have put misplaced emphasis on 'attitude'. 'Attitude' has been used inappropriately and/or measured wrongly and used to attempt to predict behaviours to which it hardly relates. But those psychologists who have tried to use attitude within a strict theoretical framework and attempted to arrive at general statements about behaviour have found that the only way that this can be achieved with any real degree of success is by specifying the conditions for behavioural predictions so tightly that the generality of the law is lost. Attitude itself becomes reified as the initial cause of the behaviour, rather than a mediation of important social inputs, that is, the socially influenced cognitive and perceptual experiences of the individual.

Originally (when the then current movement in psychology was rigorous behaviourism and avoidance of mentalistic concepts such as instinct, intentions, and introspection) social psychologists began the task of describing and explaining social behaviour from general laws in a way directly corresponding to general laws in the natural and physical sciences. It seems to me that it is now time to reassess the role of 'attitude' in the understanding of social behaviour. To do this we need a different approach to explanation, one which considers the human subject not as 'behaving' in a way analogous to a laboratory rat but as capable of forming purposive plans of action (social action in most cases). Although his 'residues of experience' can to some extent be measured as 'attitude', and although his intentions are to some degree dependent on these attitudes, direct and perceptual experiences and the resulting cognitive systems need to be studied in their own right as reflections of the social environment of the individual and the society in which he lives; these are 'fine grain'residues of experience for which attitude can be no more than an affective summary.

It is at this point that I return to beliefs. I have given some indication throughout this paper of the role and importance of beliefs, the sorts of ways in which belief systems can be modelled and these models be examined empirically. More attention can be paid to variations in beliefs, rather than losing this information by summarising it into 'attitude'. Retention and study of such local complexity in its own right will almost certainly require new methodologies. This disaggregate approach suggests a move toward induction from the beliefs and intentions of individuals to those of groups; and the aggregation criteria chosen, and the equivalences imposed on the beliefs will themselves be crucial variables in explanations of social processes and social behaviour.

REFERENCES

Abelson, R.P. and M.J. Rosenberg (1958). Symbolic Psycho-Logic: A Model of Attitudinal Cognition. *Behavioural Science*, 3, 1-13.

Bem, D.J. (1967). Self-Perception: An Alternative Interpretation of Cognitive Dissonance Phenomena. *Psychological Review*, 74, 183-200.

DeFleur, M.L. and F.R. Westie (1958). Verbal Attitudes and Overt Acts: An Experiment on the Salience of Attitudes. *American Sociological Review*, 23, 667-673.

Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Evanston, Ill., Row, Peterson.

Festinger, L. and J.M. Carlsmith (1959). Cognitive Consequences of Forced Compliance. *Journal of Abnormal and Social Psychology*, 58, 203-210.

Fishbein, M. (ed.) (1967). *Readings in Attitude Theory and Measurement*. New York, Wiley.

Fishbein, M. and F.S. Coombs (1974). Basis for Decision: An Attitudinal Analysis of Voting Behavior. *Journal of Applied Social Psychology*, 4, 95-124.

Fishbein, M., K. Thomas and J.J. Jaccard (1976). Voting Behaviour in Britain: An Attitudinal Analysis. *Occasional Papers in Survey Research*, No. 7. Published by SSRC Survey Research Unit, London.

Heider, F. (1958). *The Psychology of Interpersonal Relations*. New York, Wiley.

LaPiere, R.T. (1934). Attitudes vs. Actions. *Social Forces*, 13, 230-237.

Likert, R. (1932). A Technique for the Measurement of Attitudes. *Archives of Psychology*, No. 140.

McArdle, J.B. (1972). *Positive and Negative Communications and Subsequent Attitude and Behavior Change in Alcoholics*. Unpublished doctoral dissertation, University of Illinois.

McGuire, W.J. (1969). The nature of attitudes and attitude change. In G. Lindzey and E. Aronson (eds.), *The Handbook of Social Psychology*, 2nd ed., Vol. 3. Reading, Mass., Addison-Wesley, 136-314.

Rokeach, M. and P. Kliejunas (1972). Behavior as a Function of Attitude-Toward-Object and Attitude-Toward-Situation. *Journal of Personality and Social Psychology*, 22, 194-201.

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

Thomas, K. (ed.) (1971). *Attitudes and Behaviour*. Penguin Books, London.

Thomas K. and M. Tuck (1975). An Exploratory Study of Determinant and Indicant Beliefs in Attitude Measurement. *European Journal of Social Psychology*, 5, 167-187.

Thurstone, L.L. (1931). The Measurement of Attitudes. *Journal of Abnormal and Social Psychology*, 26, 249-269.

Wicker, A.W. (1969). Attitudes vs. Actions: The Relationship of Verbal and Overt Behavioral Responses to Attitude Objects. *Journal of Social Issues*, 25, 41-78.