Reprinted from THE JOURNAL OF SOCIAL ISSUES, Vol. XXI, No. 2, pp. 31-46, 1965.

Manipulation of Human Behavior: An Ethical Dilemma for the Social Scientist*

By Herbert C. Kelman

I

The social scientist today—and particularly the practitioner and investigator of behavior change—finds himself in a situation that has many parallels to that of the nuclear physicist. The knowledge about the control and manipulation of human behavior that he is producing or applying is beset with enormous ethical ambiguities, and he must accept responsibility for its social consequences. Even the pure researcher cannot withdraw into the comforting assurance that knowledge is ethically neutral. While this is true as far as it goes, he must concern himself with the question of how this knowledge is *likely* to be used, given the particular historical context of the society in which it is produced. Nor can the practitioner find ultimate comfort in the assurance that he is helping others and doing good. For, not only is the goodness of doing good in itself a matter of ethical ambiguity—a point to which I shall return shortly—but he also confronts the question of the wider social context in which a given action is taken. The production of change may meet the momentary needs of the client—whether it be an individual, an organization, or a community—yet its long-range consequences and its effects on other units of the system of which this client is a part may be less clearly constructive.

There are several reasons why the ethical problems surrounding the study of behavior change are of increasing concern. First, our knowledge about the control of human behavior is increasing steadily and systematically. Relevant information is being developed in various areas within psychology—clinical, social, and experimental—as

* Paper read at the symposium on "Social responsibilities of the psychologist," held at the meetings of the American Psychological Association in Philadelphia, August 30, 1963. This paper is a product of a research program on social influence and behavior change supported by Public Health Service Research Grant MH-07280 from the National Institute of Mental Health.

well as in sociology and anthropology. Personally, I do not think that the dangers from that direction are imminent. I have the feeling that the power and sensitivity of scientifically based techniques for controlling and shaping complex human behaviors are often exaggerated. Nevertheless, we are constantly working toward a systematization of this knowledge and we must at least anticipate the day when it will have developed to a point where the conditions necessary for producing a particular change in behavior can be specified with relative precision. Second, there is an increasing readiness and eagerness within our society to use whatever systematic information (or misinformation) about the control of human behavior can be made available. This readiness can be found in different quarters and in response to different motivations. It can be found among therapists and pedagogues, among idealists and agitators, among hucksters and image-makers. Third, social scientists are becoming increasingly respectable and many agencies within government, industry, the military, and the fields of public health and social welfare are becoming interested in our potential contributions. Here too there is no imminent danger. We still have a long way to go before becoming truly influential and we may find the road rather bumpy. Nevertheless, we must anticipate the possibility that social scientists will meet with a serious interest in their ideas about behavior control and have an opportunity to put them to the test on a large scale.

For all of these reasons, concern about the implications of our knowledge of behavior control is less and less a matter of hypothetical philosophical speculation. The possibilities are quite real that this knowledge will be used to control human behavior—with varying degrees of legitimacy, effectiveness, and scope. Moreover, this knowledge is being produced in a socio-historical context in which its use on a large scale, for the control of vast populations, is particularly likely. Ours is an age of mass societies in which the requirements of urbanization and industrialization, together with the availability of powerful mass media of communication, provide all the necessary conditions for extensive manipulation and control of the behavior of masses. An interest in controlling the behavior of its population is, of course, a characteristic of every society and by no means unique to our age. What *is* unique is that this is done on a mass scale, in a systematic way, and under the aegis of specialized institutions deliberately assigned to this task. Like the nuclear physicist, then, the social scientist is responsible for knowledge that—in the light of the world situation in which it is being produced—has decided explosive possibilities. It behooves us, therefore, to be concerned with the nature of the product that we are creating and the social process to which we are contributing.

In their attempts to come to grips with this problem, it seems to me, the practitioner and investigator of behavior change are confronted with a basic dilemma. On the one hand, for those of us who hold the enhancement of man's freedom of choice as a fundamental value, any manipulation of the behavior of others constitutes a violation of their essential humanity. This would be true regardless of the form that the manipulation takes—whether, for example, it be based on threat of punishment or positive reinforcement. Moreover, it would be true regardless of the "goodness" of the cause that this manipulation is designed to serve. Thus, an ethical problem arises not simply from the ends for which behavior control is being used (although this too is a major problem in its own right), but from the very fact that we are using it. On the other hand, effective behavior change inevitably involves some degree of manipulation and control, and at least an implicit imposition of the change agent's values on the client or the person he is influencing. There are many situations in which all of us—depending on our particular values—would consider behavior change desirable: for example, childhood socialization, education, psychotherapy, racial integration, and so on. The two horns of the dilemma, then, are represented by the view that any manipulation of human behavior inherently violates a fundamental value, but that there exists no formula for so structuring an effective change situation that such manipulation is totally absent.

In calling attention to the inevitability of behavior control whenever influence is being exerted, I am not suggesting that we should avoid influence under all circumstances. This is not only impossible if there is to be any social life, but it is also undesirable from the point of view of many important social values. Nor am I suggesting that we need not worry about the manipulation inherent in all influence attempts, simply because it is inevitable. The view that we can forget about this problem, because there is nothing we can do about it anyway, ignores the fact that there are important differences in degree and kind of manipulation and that there are ways of mitigating the manipulative effect of various influence attempts even if the effect cannot be eliminated entirely. This leads me to another very crucial qualification with respect to the first horn of the dilemma that I have presented. In stating that all manipulation of behavior—regardless of its form or of the purpose it is designed to serve—is a violation of the person's essential humanity, I am not suggesting that differences between different types of manipulation are ethically insignificant. The extent to which the influence attempt—despite its manipulative component—allows for or even enhances the person's freedom of choice, the extent to which the relationship between influencer and influencee is reciprocal, the extent to which

the situation is oriented toward the welfare of the influencee rather than the welfare of the influencing agent—all of these are matters of great moment from an ethical point of view. In fact, these differences are the major concern of the present analysis. But I consider it essential, as a prophylactic measure, to keep in mind that even under the most favorable conditions manipulation of the behavior of others is an ethically ambiguous act.

It is this first horn of the dilemma that Skinner seems to ignore, as can be seen from his debate with Rogers, several years ago, on issues concerning the control of human behavior.[1] Rogers, on the other hand, tends to minimize the second horn of the dilemma.

Skinner is well aware of the inevitability of control in human affairs, and argues for a type of control that is based on intelligent planning and positive reinforcement, and is not "exercised for the selfish purposes of the controller." He makes a number of telling points in responding to his critics. For example, he reminds us that, while we object to external controls, we often ignore psychological constraints that limit freedom of choice to the same or an even greater extent. He asks why a state of affairs that would otherwise seem admirable becomes objectionable simply because someone planned it that way. He points out that control based on the threat and exercise of punishment, which is built into our political and legal institutions, is fully accepted, but that use of positive reinforcement by government is regarded with suspicion. I find these and other points useful because they help us to focus on forms of control that often remain unrecognized and to consider forms of control that may be ethically superior to current ones but that we tend to reject because of their unorthodox nature. But Skinner fails to see the basis of many of the criticisms directed at him, because he is concerned about the control of human behavior only when that control is aversive, and when it is misused—i.e., when it is used for the benefit of the controller and to the detriment of the controllee. He seems unable to see any problem in the mere *use* of control, regardless of technique or purpose. This inability is consistent with his value position, which does not recognize the exercise of choice as a good *per se*.[2]

My own statement of the first horn of the dilemma is predicated on the assumption that the freedom and opportunity to choose is a fundamental value. To be fully human means to choose. Complete freedom of choice is, of course, a meaningless concept. But the purpose of education and of the arrangement of the social order, as I see it, is to enable men to live in society while at the same time enhancing their freedom to choose and widening their areas of choice. I therefore regard as ethically ambiguous any action that limits freedom of choice, whether it be through punishment or reward or even through so perfect an arrangement of society that people do not care to choose. I cannot defend this value because it is not logically derived from anything else. I can, of course, offer supporting arguments for it. First, I can try to show that the desire to choose represents a universal human need, which manifests itself under different historical circumstances (not only under conditions of oppression). Second, I can point out that freedom of choice is an inescapable component of other valued states, such as love, creativity, mastery over the environment, or maximization of one's capacities. Third, I can try to argue that valuing free individual choice is a vital protection against tyranny: quite aside from the notion that power corrupts its user, even the well-motivated, unselfish controlling agent will be tempted to ignore human variability and to do what *he* thinks is good for others rather than what they think is good for themselves—and thus in essence become tyrannical—if he is unhampered by the right to choose as a basic human value. While I can offer these supporting arguments, I recognize that freedom of choice is, in the final analysis, a rock-bottom value for me. Skinner is not concerned with the dilemma presented here because apparently he does not share this fundamental value, even though he is strongly committed to certain other related values, such as the rejection of aversive control and selfish exploitation (albeit without recognizing their status as values).

With Rogers on the other hand, I feel a complete affinity at the value level. He values "man as a self-actualizing process of becoming" and in general proposes that "we select a set of values that focuses on fluid elements of process rather than static attributes." He favors a society "where individuals carry responsibility for personal decisions." He regards "responsible personal choice" as "the most essential element in being a person." But, as I have pointed out, Rogers tends to minimize the second horn of the dilemma presented here—the inevitability of some degree of manipulation in any influence attempt. He makes what appears to me the unrealistic assumption that by choosing the proper goals and the proper techniques in an influence situation one can completely sidestep the problem of manipulation and control. He seems to argue that when an influencing agent is dedicated to the value of man as a self-actualizing process and selects techniques that are designed to promote this value, he can abrogate his power over the influencee and maintain a relationship untainted by behavior control. This ignores, in my

[1] Carl R. Rogers and B. F. Skinner. Some issues concerning the control of human behavior. *Science*, 1956, *124*, 1057-1066.

[2] This in turn is related to a point stressed by Rogers, namely Skinner's underestimation of the role of value choices in human affairs in general and in the application of science to social problems in particular.

opinion, the dynamics of the influence situation itself. I fully agree that influence attempts designed to enhance the client's freedom of choice and techniques that are consistent with this goal are ethically superior, and that we should continue to push and explore in this direction. But we must remain aware that the nature of the relationship between influencing agent and influencee is such that inevitably, even in these influence situations, a certain degree of control will be exercised. The assumption that we can set up an influence situation in which the problem of manipulation of behavior is removed, because of the stated purpose and formal structure of the situation, is a dangerous one. It makes us blind to the continuities between all types of influence situations and to the subtle ways in which others can be manipulated. It lulls us into the reassuring certainty that what we are doing is, by definition, good. I would regard it as more in keeping with both the realities of behavior change, and the ethical requirements of minimizing manipulation, to accept the inevitability of a certain amount of control as part of our dilemma and to find a *modus vivendi* in the face of the ethical ambiguities thus created.

II

Let me proceed to examine briefly the implications of this general dilemma for each of three roles involving social science knowledge about behavior change: the practitioner, as exemplified by the psychotherapist and the group leader or group process trainer; the applied researcher, such as the social scientist in industry or the public opinion pollster; and the basic researcher, such as the investigator of attitude change. These roles are, of course, highly overlapping, but separating them may help us focus on different nuances of the general dilemma.

The practitioner must remain alert to the possibility that he is imposing his own values on the client; that in the course of helping the client he is actually shaping his behavior in directions that he—the practitioner—has set for him. Thus, psychotherapy, even though it is devoted to what I would consider a highly valuable end—enabling the patient to live more comfortably and achieve his own goals more effectively—is definitely open to the possibility of manipulation. Psychotherapy (at least "good" psychotherapy) is markedly different from brainwashing: the client enters into the relationship voluntarily; the therapist is concerned with helping the patient, rather than with furthering his own ends or the ends of some institution that he represents; influence techniques are designed to free the patient, to enhance his ability to make choices, rather than to narrow his scope. Yet there are some striking similarities between the methods of therapy and those of brainwashing to which the

therapist must always remain alert, lest he overstep what is sometimes a rather thin line. The therapist cannot avoid introducing his own values into the therapeutic process. He cannot be helpful to the patient unless he deliberately tries to influence him in the direction of abandoning some behaviors and trying out others. But in doing so he must beware of two types of dangers. One is the failure to recognize that he is engaged in the control of the client's behavior. The other is intoxication with the goodness of what he is doing for and to the client, which in turn leads to a failure to recognize the ambiguity of the control that he exercises. Only if he recognizes these two conditions is he able to take steps to counteract them.

Similar considerations hold for the group leader. Some of the principles of group leadership developed by social psychologists and variously called applied group dynamics, human relations skills, or group process sensitivity are highly congenial to democratic values. They are designed to involve the group in the decision-making process and to foster self-expression on the part of the individual member. Yet the possibilities for manipulation abound. A skillful group leader may be able not only to manipulate the group into making the decision that he desires, but also to create the feeling that this decision reflects the will of the group discovered through the workings of the democratic process. This need not involve a deliberate Machiavellian deception on the part of the group leader; the leader himself may share the illusion that a group product has emerged over which he has exercised no influence. It is essential, therefore, to be fully aware of the leader's control implicit in these techniques. Some of their proponents argue that, by their very nature, these techniques can be used only for democratic ends. I would question this assumption and, in fact, consider it dangerous because it exempts the group leader from asking those questions that any practitioner of behavior change should keep before his eyes. What am I doing in my relationship to the client? Am I creating a situation in which he can make choices in line with his own values, or am I structuring the situation so that my values dominate?

When the group leader is involved in training others in human relations skills or sensitivity to group process, he is confronted with a further problem. Typically, the trainee is a member of some organization—industrial, governmental, military, educational, religious—in which he will apply the skills he is now learning. The human relations trainer is, thus, in a sense improving the trainee's ability to manipulate others in the service of the organization that he represents. Of course, this is not the goal of the training effort, and trainers always try to communicate the value of the democratic process in group life. But the fact remains that they are training a wide variety of people who will be using these skills for a wide variety of ends.

It can certainly be argued that the widespread introduction of human relations skills is likely to have more positive than negative effects from the point of view of a democratic ideology. Perhaps this is true. But it is dangerous to assume that these skills carry their own built-in protection. There is no substitute for a continued attention, on the trainer's part, to questions such as these: Whom am I training? To what uses will they put the skills that I am placing at their disposal? What are the organizational processes in which I am now becoming a partner?

It is essentially these same questions to which the applied social researcher in the broad field of behavior change must address himself. I am here thinking specifically of applied research in the sense that it is done for a client. While the researcher is merely gathering facts, he is nonetheless participating quite directly in the operations of the organization that employs him. If his work is successful, then his findings will be applied to the formulation and execution of the organization's policies. There is thus the real possibility that the investigator is directly helping the organization in its attempts to manipulate the behavior of others—workers in an industry, consumers, or the voting public.

Let us take, for example, the industrial social scientist who studies factors affecting worker morale. On the basis of his recommendations, and often with his direct assistance, management may become more aware of human relations aspects of industrial work and introduce methods designed to improve morale. Ideally, these methods would consist of increased involvement and participation of workers in decisions relating to their jobs. Critics of this type of approach argue that the social scientist is working for management, providing them with information and introducing procedures that are designed to increase productivity at the worker's expense. The assumption in this criticism, to which I think there is some validity, is that the worker is being manipulated so that he experiences a sense of participation and involvement which is not reflected in the reality of his position within the industrial organization. In response to this criticism it can be argued that—considering the over-all lack of satisfaction in industrial work—it is a net good to give the worker some sense of participation and involvement in the work situation, to give him at least a limited opportunity to make choices and thus find some meaning in the job. To be sure, management is interested in these innovations because they expect the changes to increase productivity, but does that necessarily vitiate the advantages from the worker's point of view? This is a rather convincing defense, but in evaluating the pros and cons we must also take into account the social context in which these changes are introduced. What effect does the human relations approach have on unions, which represent

the only source of independent power of the industrial worker? Does it sidestep them, and will it eventually weaken them? What are the general implications of helping the worker adjust to a situation in which he has no real freedom of choice, in any ultimate sense? These questions are not easy to answer, and every social scientist has to decide for himself whether his work in industry is doing more good than harm. In deciding whether or not—and in what way—to do applied social research in industry or elsewhere, the social scientist must ask himself: Whom am I doing this work for? How are my findings likely to be used? Will they increase or decrease the freedom of choice of the people whose behavior will be influenced? What are the social processes, both short-run and long-run, in which I am participating via my research?

Another example of applied social research that raises questions about manipulation of the population is public opinion polling, when used in connection with political campaigns or the political process in general. For instance, in a recent presidential election, computer simulation was used—based on data derived from numerous opinion polls—to predict the responses of various segments of the population to different campaign issues. Information generated by this process was made available to one of the political parties. This type of social research has some troubling implications. It raises the possibility that a candidate might use this information to manipulate the voters by presenting a desirable image—that is, saying what the public presumably wants to hear. In defense against such criticisms, the originators of this technique have pointed out that it represents a systematic way of providing the candidate with relevant information about the interests and concerns of the public, or of particular publics. He can then address himself to those issues with which the public is deeply concerned, thus making his campaign more relevant and meaningful and enhancing the democratic political process. They point out further that this is what candidates try to do anyway—and properly so; all the social scientist does is to help them base their campaigns on more adequate information, rather than on the usually unreliable estimates of politicians. Of course, what assurance do we have that opinion polls and computer simulations based on them will, in fact, be used in this ideal manner to bolster the democratic process, rather than to short-circuit it? The information can be used both to widen and to restrict the citizen's freedom of choice. But, as long as it is information that can help political organizations to manipulate the public more effectively, the researcher must concern himself actively with the question of how it is going to be used and to what kind of process it is going to contribute.

For the man engaged in "basic" research on one or another aspect of behavior change—in contrast to the man who does research

for a specific client—it is much easier to take the position that the knowledge he produces is neutral. Yet, since there is a possibility that his product will be used by others for purposes of manipulation, he cannot be completely absolved from responsibility. He must consider the relative probabilities, given the existing socio-historical context, that this knowledge will be used to enhance or to restrict people's freedom of choice. These considerations must enter into his decision whether or not to carry out a given piece of research, and how to proceed with it.

Take, for example, the area of attitude change, with which I myself am strongly identified. Much of the research in this area is clearly dedicated to the discovery of general principles, which can presumably be applied to many situations with differing goals. Yet, because of the nature of the principles and the experimental settings from which they are derived, they can probably be applied most readily, most directly, and most systematically to mass communications. And, because of the nature of our social order, they are particularly likely to be used for purposes of advertising, public relations, and propaganda—forms of mass communication that are least oriented towards enhancing the listener's freedom of choice. There are, of course, many reasons for continuing this line of research, despite the probability that its findings will be used for manipulative purposes. First, one can argue that extending our general knowledge about processes of attitude change and increasing our understanding of the nature of influence are in themselves liberating forces, whose value outweighs the possibility that this knowledge will be used for undesirable ends. Second, such research may not only increase the knowledge of the potential manipulator, but also help in finding ways to counteract manipulative forces—by providing the information needed for effective resistance to manipulation, or by developing knowledge about forms of influence that enhance freedom of choice. Third, one might argue that information about attitude change, despite its potential for manipulative uses, is important for the achievement of certain socially desirable goals—such as racial integration or international understanding.

I obviously find these arguments convincing enough to continue this line of research. But the nagging thought remains that the knowledge I am producing—if it has any scientific merit—may come to be used for ever more effective manipulation of human behavior. Thus, even the basic researcher in the domain of behavior change must always ask himself: Given the realities of our present society, what are the probable uses to which the products of my research are going to be put? What are the social processes to which I am contributing by the knowledge that I feed into them?

III

The very fact that I have presented my position in the form of a dilemma should make it clear that I do not see an ultimate "solution"—a way of completely avoiding the ethical ambiguity with which practitioners and researchers in the field of behavior change are confronted. I do feel, however, that there are ways of mitigating the dehumanizing effects of new developments in the field of behavior change. I would like to propose three steps that are designed to contribute to this end. Stated in their most general form, they would involve: (1) increasing our own and others' active awareness of the manipulative aspects of our work and the ethical ambiguities inherent therein; (2) deliberately building protection against manipulation or resistance to it into the processes we use or study; and (3) setting the enhancement of freedom of choice as a central positive goal for our practice and research. In order to spell out in somewhat greater detail what these three steps might imply, I would like to examine them from the point of view of each of the three separate (though overlapping) roles that have already been differentiated: the role of the practitioner, of the applied researcher, and of the "basic" researcher in the field of behavior change. The argument that follows is summarized in Table 1.

TABLE 1
STEPS DESIGNED TO MITIGATE THE MANIPULATIVE ASPECTS OF BEHAVIOR CHANGE IN EACH OF THREE SOCIAL SCIENCE ROLES

Desirable Steps	Role of Practitioner	Role of Applied Researcher	Role of Basic Researcher
(1) Increasing awareness of manipulation	Labelling own values to self and clients; allowing client to "talk back"	Evaluating organization that will use findings; considering on whom, how, and in what context they will be used	Predicting probabilities of different uses of research product, given existing sociohistorical context
(2) Building protection against or resistance to manipulation into the process	Minimizing own values and maximizing client's values as dominant criteria for change	Helping target group to protect its interests and resist encroachments on its freedom	Studying processes of resistance to control, and communicating findings to the public
(3) Setting enhancement of freedom of choice as a positive goal	Using professional skills and relationship to increase client's range of choices and ability to choose	Promoting opportunities for increased choice on part of target group as integral features of the planned change	Studying conditions for enhancement of freedom of choice and maximization of individual values

I have already stressed how essential it is for the *practitioner* of behavior change to be aware of the fact that he is controlling the client, that he is introducing his own values both in the definition of the situation and in the setting of standards. Thus, in the therapeutic situation, it is not only inevitable but also useful for the therapist to have certain values about what needs to be done in the situation itself and what are desirable directions in which the patient might move, and to communicate these values to the patient. But he must be clear in his own mind that he is bringing these values into the relationship, and he must label them properly for the patient. By recognizing himself that he is engaged in a certain degree of control—and that this is an ethically ambiguous act, even though his role as therapist requires it—and by making the patient aware of this fact, he provides some safeguards against this control. Among other things, such a recognition would allow the patient, to a limited extent, to "talk back" to the therapist—to argue about the appropriateness of the values that the therapist is introducing. A therapeutic situation is, of course, not a mutual influence situation in the true sense of the word: by definition, it is designed to examine only the patient's values and not those of the therapist. But, from the point of view of reducing the manipulativeness of the situation, it would be important to encourage mutuality at least to the extent of acknowledging that what the therapist introduces into the situation is not entirely based on objective reality, but on an alternative set of values, which are open to question. There may be particular therapeutic relationships in which a therapist finds it difficult to acknowledge the values that he brings to them, because his own motivations have become too deeply involved. There may also be institutional settings in which the therapist is required to present the institutional values as the "right" ones, in contrast to the patient's own "wrong" values. These are danger signals, and the therapist may well consider refraining from entering a therapeutic relationship or working in an institutional setting in which he is not free to acknowledge the contribution of his own values.

Second, in addition to increasing awareness of the manipulative aspects of the situation, it is important to build into the change process itself procedures that will provide protection and resistance against manipulation. For the practitioner of behavior change this means structuring the influence situation in such a way that the client will be encouraged to explore his own values, and to relate new learnings and new behavioral possibilities to his own value system. At the same time, it is important that the practitioner—be he therapist or group leader—keep to a minimum the direct and indirect constraints that he sets on the influencee. Constraints are, of course, necessary to varying degrees, both for the protection of clients

and for keeping the process moving in useful directions. Insofar as possible, however, the situation should be so structured that the influencee determines the direction of the process to a maximal extent. It should be noted that what I am suggesting here is not the same as the use of non-directive techniques. In and of themselves these merely represent a set of formal techniques which may or may not have the desired effect. The crucial point is that the client's own values should be at the center of attention when change is under consideration and should be readily available as criteria against which any induced behavior can be measured. To the extent to which this is true, the patient or the group will be in a better position to resist manipulation in the service of alien values. Often, however, this will require much more than non-interference on the part of the practitioner. It may require active efforts on his part to encourage the client to bring his values to the fore and measure the induced changes against them.

Third, it is important to go beyond providing protection and resistance against manipulation that would encroach on the client's freedom of choice. The actual *enhancement* of freedom of choice should, ideally, be one of the positive goals of any influence attempt. Thus, the therapist should use his professional skills and his relationship to the patient to provide him with new experiences that enhance his ability to choose (and thus to maximize his own values) and with new information that widens his range of choices. Similarly, the group leader should attempt to bring the group to a point where members can make more effective and creative choices, conducive to the achievement of individual and group goals. The enhancement of freedom and creativity as the positive value toward which behavior change should be directed has been discussed most eloquently by Rogers.[3]

Needless to say, it would be essential to include in the training of practitioners of behavior change and in their professional standards some consideration of these three desiderata for mitigating the manipulative aspects of their activities. If they learn to acknowledge the role of their own values in the situation, to make active efforts at keeping the client's values in the foreground, and to regard increased freedom of choice as a primary goal, they are less likely to make full use—either unwittingly or by design—of the potential for manipulation that they possess.

Now let me turn to the *applied researcher*. In deciding whether to take on a particular piece of research, he must keep in mind that the information he is being asked to supply may be used for the manipulation of others—e.g., workers in an industry for whom he

[3] For example, in his debate with Skinner, *ibid.*

is doing a morale survey, or the voting public if he is working with poll data. The question of *who* is employing him becomes crucial, therefore. He must evaluate the organizations that will be using his findings, and consider how they are likely to use them, whose behavior they will attemp to influence, and in what context this influence will occur. He must consider the probable uses of these findings not only in the short run, but also in the long run. Thus, for example, he cannot simply rely on the fact that his contact man in an organization is someone he trusts. If this man is in a peripheral position within the organization, and if the organization is generally undemocratic and exploitative in its orientation, then the long-run prospects are not too reassuring. There is, of course, the possibility that the research itself will have a liberalizing effect on the organization; the probability that this will, in fact, happen must also be estimated. In the final analysis, there can be no foolproof guarantees, but the investigator must at least feel reasonably certain that the net effect of his research will not be a reduction in the freedom of choice of a segment of the population. Each investigator has to draw his own line, both with respect to the probability and the amount of manipulation that he is willing to tolerate. If they are likely to go beyond this line, then he must consider turning down the assignment. Once a researcher has decided to take on an assignment, he must continue to keep the manipulative potential of his findings in mind, and try to counteract it by the way he communicates his findings and the recommendations he bases on them. If his research is, indeed, to have a liberalizing effect on the organization, then he will have to take active steps in this direction.

In order to build some protection against manipulation into the change procedures based on his findings, the researcher should make it a rule to communicate directly with the target group (i.e., the group that is to be influenced) and to involve it in the research, and in the change process insofar as he has charge of it. Thus, an industrial social scientist employed by management might insist on informing the workers in detail about the purposes and findings of the research and the attempted changes that are likely to result from it. In giving them this information, he would try to help them protect their interests against undue attempts at manipulation and to offer them specific recommendations for resisting encroachments on their freedom of choice. Furthermore—in order to promote freedom of choice as a positive goal—he should make a concerted effort to influence the planned change that will be based on his research so that it will actually leave the target group with greater choice than it had before. In submitting his findings and recommendations to the organization that contracted for the research, he should actively seek and point up opportunities for enhancing freedom of choice on the part of the target group that can be integrated into the planned change.

The two last points both imply a rather active role for the researcher in the planning of change based on his research. I would not want to say that the researcher must always participate directly in the change process itself; there are many times when this would be impossible or inappropriate. But since he is providing information that will (at least in principle) be directly translated into action, it is his responsibility to take some stand with respect to this action. The uses to which the information is put are not only the problem of the contracting organization, but also very much the problem of the man who supplied the information. The researcher should be clear about this, and he should have the support of his profession when he takes an active part in expressing his point of view.

Let me finally, and more briefly, turn to the *basic researcher.* I have already stated my position that, even though the products of pure research are in a sense neutral, the investigator cannot escape responsibility for their probable consequences. The student of attitude change, for example, must keep in mind that his findings can be used for the systematic manipulation of the population, in ways and for purposes that would produce a net constriction in freedom of choice. In deciding whether or not to proceed with his research, he must try to make some estimate of the probabilities of different uses of his research product, in the light of existing social forces. If he expects restrictive uses to outweigh constructive ones, he would be bound to refrain from engaging in this research. If, on balance, he decides to undertake the research—and there are, of course, many good reasons for doing so—then he must continue to remain alert to its manipulative potential, and must constantly review his decision, particularly as his research emphases shift or as social conditions change.

Researchers in this area also have a special responsibility to be actively concerned with the ways in which the knowledge they produce is used by various agencies in their society. Eternal vigilance to the possibilities of manipulation is, of course, the duty of every citizen. But, as producers of knowledge about manipulation, social scientists are in a position similar to that of the many nuclear physicists who feel a *special* sense of responsibility for the ways in which their knowledge is being used.

Earlier, I suggested that research on attitude change may not only increase the knowledge of the potential manipulator, but also help in finding ways to counteract manipulative forces. So far, research along these lines has been rather limited. If investigators of attitude change and related problems are to mitigate the manipulative potential of their research, they will have to focus more deliber-

ately and more actively on this other line of work. Thus, in order to build some protection against manipulation into the social structure, we will have to extend our research on processes of resistance to control, and make a special effort to communicate relevant findings to the public. Such an emphasis will contribute to the development of antidotes against manipulation at the same time that research is contributing to the development of knowledge about manipulation itself. From a scientific point of view, such work will be highly germane to the study of attitude change, since it represents an exploration of its limiting conditions.

In order to promote the enhancement of freedom of choice as a positive goal, research will have to focus on the conditions favoring a person's ability to exercise choice and to maximize his individual values. Admittedly, this is a rather value-laden way of stating a problem for basic research. However, if we want our science to contribute to the liberation of man rather than to his dehumanization, this is the kind of problem to which we will have to turn our attention.