Part Two

Practices, Methods, and Results of Normal Science and Action Science

Within the domain of normal science there are many different communities of inquiry, each with its own set of norms and practices. In this section, we will explore case examples from four such communities: experimental social psychology (Chapter Four); assessment research as it is used to study the fit between individual factors and organizational factors (Chapter Five); and both basic and applied ethnographic approaches as they are used to study educational contexts (Chapter Six).

In each case study we describe the researchers' approach to an important social problem: from obedience to authority, to satisfaction and productivity in the work place, to failure in the classroom. In doing so, we examine not only the methods and theoretical frameworks used but the operating assumptions that are embedded in them. We then explore the unintended consequences that result even when the researchers' thrust is implemented well, as is the case in each of these studies. Finally, we make suggestions on how some of these inner contradictions might be reduced, and we draw on case examples from action science to illustrate how these might be implemented.

After presenting the case studies, we then examine the different research approaches in light of the norms and rules that characterize the communities of inquiry in which the researchers work (Chapter Seven). We argue that researchers, like all practitioners, are bound by the rules and norms of their practice, and that these place researchers in dilemmas that cannot be resolved within their communities as they are now defined. Moreover, we try to show how the norms these researchers follow lead them to design their inquiry in such a way that the knowledge they produce is of limited use to practitioners. The basic researchers ask, "How does it happen to be?" And the applied researchers build on basic knowledge to pursue the question, "How do we achieve a given set of ends?" But both basic and applied researchers stop short of considering the practitioner's question, "How do I understand and act in real life contexts amidst all their complexity and dilemmas of value?"

In the final chapter to this section (Chapter Eight), we put forth the norms and rules that guide the practice of action science. Here we show how action science overlaps with and departs from more traditional norms of inquiry as it seeks to adapt these norms to the realm of social practice. We will show how this leads the action scientist to ask new questions, to use new methods, and to construct new solutions. We hope that these will suggest new ways of managing the dilemmas of practice—for both researcher and practitioner alike.

4

Beyond the Limitations of Normal Science: Comparing Laboratory Experiments and Action Experiments

Back in the 1960s in a couple of typical laboratory rooms at Yale, a research psychologist named Stanley Milgram conducted a series of experiments on obedience to authority (Milgram, 1974). Since then, many people have had a hard time figuring out whether we should be distressed more by what Milgram did or by what he found out. Few experiments, if any, have generated more controversy. On one side, proponents heralded it as the most important research in recent history; while on the other side, critics condemned it as the most unethical. While this debate still goes on, almost everyone agrees on one thing: Milgram's results were surprising, distressing, and important, because he found that ordinary people will obey authority even when it violates their most central values and leads them to harm others.

Milgram's experiment exemplifies social psychological inquiry into real-life dilemmas with nontrivial consequences. Experiments like Asch's (1952) on conformity, Latané and Darley's (1970) on the innocent bystander, and Zimbardo's on vandalism (1969) are all of the same genre. They look at actual

moral dilemmas faced in everyday life; they uncover inconsistencies between the moral reasoning we espouse and our actual behavior; and they explore the consequences of this for both the actors and those around them. Because this research tackles such important problems, we might expect that it would yield more useful results than research that takes on issues of little social relevance or consequence. To a certain extent this is true, but even this kind of research leaves us hanging on the edge of the dilemmas it raises, not knowing how to climb beyond them.

At present, pointing to solutions is not the normative role taken by scientists, so this last statement may seem either unfair or irrelevant. One might argue that science is a two-step process. Scientists generate data and build theory first, and then let practitioners or applied social scientists figure out how to solve the problems. But this strict division between research, theory, and practice may have serious consequences for each endeavor. We may be conducting research and creating theory that would be difficult or even impossible to use to solve problems under real-life conditions. And practitioners may be acting with tacit propositions about the world that are not easily falsifiable and therefore fraught with errors that go undetected. Since the left hand doesn't know what the right one is doing, both risk making elegant but irrelevant gestures.

In this chapter we use Milgram's experiment to consider this possibility not because of its deficiencies but because of its strengths. Milgram makes explicit connections between what is studied in the laboratory and what occurs in everyday life; he devises and tests out multiple variations and hypotheses in order to understand the processes and conditions underlying obedience; and he offers an elegant and comprehensive theoretical framework in which to explain the different sources, mechanisms, and processes of obedience. We thus choose it not only because it is representative of other inquiries but because, as Milgram's obedient subjects might say, "If doing what we are supposed to do gets us into trouble, then we really do face a problem."

Obedience to Authority

In the first few pages of Milgram's study we encounter two dilemmas. First, persons who abhor hurting others will do so if ordered to by an authority. Second, while obedience is essential for social organization, it is also responsible for phenomena like the mass murder of Jews during the Second World War. One thus might be left wondering: If obedience is at once necessary *and* potentially destructive, what can we do about it? If Milgram's formulation is correct, we can do very little. But it may be that in trying to explain only what exists while not exploring what might exist, he offers an incomplete understanding of what obedience is and how we might resolve these dilemmas. To consider this possibility, let us look at how Milgram studied the problem of obedience and the theory he developed to account for it. In looking at both, we will ask two interrelated questions: What kind of learning does his methodology and knowledge generate, and what kinds of solutions are most likely to follow.

Obedience in the Laboratory. Milgram points out that it is a bit of a leap from the extermination camps in the Germany of 1940 to the Yale labs in today's United States. But in setting up his experimental situation, he tried to maintain the essential features and dilemmas of obedience to authority by constructing a situation in which a person is told by an authority to act against someone else. Using the experimental situation itself, Milgram thus devised the now famous scenario of an experimenter who ordered a naive subject to administer increasing amounts of shock to a confederate subject each time the confederate made a mistake on a learning task. Then he asked of the situation that he had contrived: How far will the subject go before refusing to carry out the experimenter's order?

Of course the confederate "learner" was not actually shocked, but the naive subjects didn't know this. Instead the subjects fully believed that they were taking part in an experiment on the effect of punishment on learning and that they were administering increasingly dangerous levels of shock to an

actual learner. After all, the subjects had met the "learner"; they had seen him strapped into a chair in an adjoining room; they had watched as electrodes were attached to his wrists; they were placed behind an impressive shock generator with a range of volts from "slight shock" to "danger—severe shock"; and finally, to remove any doubts about what they were doing, they could hear the shouts and protests of the learner from the next room each time they administered a shock. Milgram thus set up a convincing scenario—so convincing that the subjects themselves were psychologically shocked to discover that it was all staged. As one subject exclaimed after being introduced to a nonplused "learner": "You're an actor, boy. You're marvelous! Oh, my God, what he [the experimenter] did to me. I'm exhausted. I didn't want to go on with it. You don't know what I went through here. A person like me hurting you, my God. I didn't want to do it to you. Forgive me, please. I can't get over this. My face is beet red. I wouldn't hurt a fly. I'm working with boys, trying to teach them, and I'm getting such marvelous results *without* punishment. I said to myself at the beginning, I don't feel you'll get anything by inflicting punishment" (Milgram, 1974, p. 82; italics his).

As this suggests, the subject was not only shocked to learn that the scenario was staged, she was also shocked to discover that while she didn't wish to hurt anyone, she had proceeded to the end of the board, where it read: "Danger—severe shock." Perhaps because of this, once the hoax was revealed, Milgram found that she and others tended to distance themselves from their responsibility for their actions. When asked why they continued, they explained that they did not believe the confederate subject was being hurt, they pointed out that an authority had required them to do so, they devalued the victim, or they cited some ideological justification for their actions. As one person put it, "In the interest of science, one goes through with it" (Milgram, 1974, p. 54).

Modifying this basic scenario, Milgram went on to examine the different mechanisms underlying obedience and to rule out alternative interpretations by varying experimental conditions. One set of variations focused on the variable of physical

distance, and it was found that the more distance between the subject and the authority and the less distance between the subject and the learner, the less likely it would be that the subjects would obey. A second set of variations examined aggression as an alternative interpretation of the results. It showed that left to their own devices, subjects did not choose to shock the learner, a result that ruled out the intention to harm as a plausible explanation. A third set, which asked individuals to predict how they themselves or others would handle the situation, found that almost everyone predicted that they and others would disobey, suggesting a deep belief in their good intentions and their capacity for carrying them out in a stressful situation —a belief not borne out by how individuals actually behaved once in the situation. Finally, another variation changed who ordered the shock and who received and opposed it. For instance, in one such variation, a person with the *status* of an authority (the experimenter) was put in the *position* of a learner and strapped to the chair where he *acted* to oppose the shock. But regardless of the position or the action taken, the moment that someone with the status of authority commanded a certain action, most subjects obeyed.

But not everyone. A small percentage disobeyed, unilaterally refusing to go further and putting an end to the experiment. When told to continue, these subjects opposed the authority's commands (Milgram, 1974, p. 51):

Mr. Rensaleer: Well, I won't [continue] —not with the man screaming to get out.

Experimenter: You have no other choice.

Mr. Rensaleer: I do have a choice. (Incredulous and indignant) Why don't I have a choice? I came here on my own free will. . . .

Why Obedience? For Milgram "the fact most urgently demanding explanation" was the "extreme willingness" of most adults to act against another individual and their own values on the command of an authority (1974, p. 5). To explain this fact, Milgram concluded that while such unquestioning obedience

may be "surprising and dismaying," it is nevertheless necessary
for social organization. He argued that "the basic reason why
[the suppression of local control] occurs is rooted not in indi-
vidual needs but in organizational needs. Hierarchical structures
can function only if they possess the quality of coherence, and
coherence can be attained only by the suppression of control at
the local level" (p. 131). Embedded in this explanation are two
interrelated assumptions. The first is that there are only two
ways of responding to an authority: either obey without ques-
tion or unilaterally disobey, as the behavior of the subjects in
his experiments in fact suggests. And the second follows from
the first. If our only two options are obedience or unilateral dis-
obedience, then it follows that social coherence would be con-
tingent on obedience at the local level and the exertion of uni-
lateral control at upper levels. Intuitively this makes much
sense. In everyday organizational life people ordinarily do obey;
and when they disobey, they usually do so unilaterally, often
jeopardizing organizational coherence and survival as a result.

So herein lies the dilemma. The obedience that organiza-
tional life requires can result in anything from the little "mur-
ders" of everyday life to the mass murders in extreme situa-
tions. Hence a paradox: The very conditions deemed necessary
for coherence can simultaneously create the incoherence of hu-
man suffering and slaughter. But suppose for a moment that it
was possible to invent an alternative form of authority relations,
one that does not yet exist but that might resolve this paradox.
If this were possible, it would show that existing authority ar-
rangements are not necessary, and it might aid us in resolving
the existing dilemmas of obedience to authority. But Milgram
did not submit this possibility to experimentation. In fact, the
methods of the laboratory rule out such exploration. As a re-
sult, the descriptions these methods yield necessarily contain
important limitations.

The Rules of the Laboratory. One rule of the laboratory
is reflected in Milgram's account of his experiment. He writes
that "the question arises as to whether there is any connection
between what we have studied in the laboratory and the forms
of obedience we so deplored in the Nazi epoch. . . . the differ-

ences ... are enormous, yet the difference may turn out to be relatively unimportant as long as certain essential features are retained. ... The question of generality, therefore, is not resolved by enumerating all the manifest differences between the psychological laboratory and other situations but by carefully constructing a situation that captures the essence of obedience" (1974, p. xii).

The logic here may be stated as follows: In order to reliably describe some phenomenon, one ought to retain its essential features and construct a situation that captures its essence. This is consistent with what Milgram actually did. The situations that he created replicated existing authority relationships, and he never called into question or tried to alter obedient responses under conditions conducive to them (closeness to authority and distance from learner). The basic assumption is that by replicating the world, we will generate a description of the phenomenon that can be generalized beyond the confines of the lab.

To a large extent we agree with this logic. But we believe that it simultaneously generates limits that most experimentalists are unaware of. One of the most important of these limits is that by not trying to alter what is, the experimentalist is unlikely to uncover the deep defensive structures that maintain existing social action and relationships. Such structures become evident only when existing social arrangements are threatened. It is at this point that our hidden defenses are mobilized and come to the surface, thereby opening up the possibility of inquiry into these structures and of transforming what is discovered. By replicating only what is, we can neither describe these defensive structures nor discover whether the social structures they maintain are in fact necessary. Only attempts at transforming what is can yield this knowledge.

But a second rule of the laboratory makes such attempts quite unlikely. As Milgram explains it: "Simplicity is the key to effective scientific inquiry. This is especially true in the case of subject matter with a psychological content. Psychological matter, by its nature, is difficult to get at and likely to have many more sides to it than appear at first glance. Complicated

procedures only get in the way of clear scrutiny of the phenomenon itself. To study obedience most simply, we must create a situation in which one person orders another person to perform an observable action and we must note when obedience to the imperative occurs and when it fails to occur" (1974, p. 13).

We agree that simplicity is desirable and that complicated procedures can obscure a phenomenon. But the way simplicity is achieved in the laboratory can simultaneously confound phenomena and keep us from probing them fully. To simplify the question of obedience, Milgram focused on a limited number of variables, a priori defined and identified by him; and he had to conceal the experimental hypothesis and manipulation to avoid confounding the results. In so doing, he exerted a high degree of conceptual and situational control over the research context and the interactions that took place in it (see Cassell, 1982, for a discussion of different forms of research control). The problem is that this control, intended to minimize the effect of the experimenter and extraneous factors, can itself exert causal influence that threatens validity. Recall the subject with the beet-red face. She exclaimed, "Oh, my God, what he [the experimenter] did to me. I'm exhausted." Reactions like these suggest that the experimental hoax led subjects not only to feel distressed by their own actions but to feel that the experimenter's actions had set them up, leaving them publicly exposed and caught at violating their own values in a situation not of their own making. In short, they may have felt that someone had just pulled a fast one—which someone had—thereby influencing how the subjects accounted for their actions in the postexperimental interview. If so, then the actual experimenter may also have contributed to, and may in part have accounted for, the defensiveness and disclaimers of responsibility reflected in the subjects' self-reports.

Equally important from an action science perspective is that this form of control creates an additional limit on the kinds of inquiry possible. According to Milgram, obedience is highly learned in the course of socialization at home, at school, and at work. Over and over again we come across and internalize the

axiom: "Do what the man in charge says" (1974, p. 138). Milgram continues, "Just as we internalize grammatical rules, and can thus both understand and produce new sentences, so we internalize axiomatic rules of social life which enable us to fulfill social requirements in novel situations" (p. 138).

We agree with this formulation. But suppose we now wished to discover whether these rules are alterable? First, our research suggests that altering rule-governed behavior requires commitment on the part of participants, and this in turn requires that they share control over the situation. Consequently, the variables participants identify, the meanings they impose on them, and the hypotheses they wish to test all must move to the foreground of the inquiry and can no longer be considered extraneous or be controlled unilaterally by the experimenter. Second, we have found that changing rule-governed behavior requires providing a kind of primary socialization process in which new routines for everyday life can be learned (Berger and Luckmann, 1966). This means that multiple opportunities must be provided for interrupting old rules, for experimenting with and reflecting on new ones, and for continuously reflecting on what's being learned. A few experimental manipulations would not be sufficient, and the unilateral control of the laboratory would be impossible.

Yet even if this control were possible, from our point of view it would not be desirable. Such control is more apt to foster than to alter obedient responses and diminished responsibility for one's actions. Two of the requisite conditions for obedience are the legitimacy and unilateral nature of authority relationships (Harmon, 1981; Milgram, 1974). If we wish to alter the nature and legitimacy of these relationships, we cannot do it in a context that is legitimately characterized by them. As Zuniga (1975) argued, "The epistemological strategy invariably generates the sociological one. . . . Scientists elevate the laboratory setting to the rank of ideal research paradigm. The problem is deciding whether it is the ideal representation of the power relations of a democratic society" (p. 110).

A third rule further limits our ability to alter what we describe. Theories developed in the social sciences are primarily

meant to explain causal relationships, and social scientists avoid judging what is from a moral standpoint. Milgram is no exception. His intention was to understand, not to judge or to prescribe. The latter is the task of the normative realm, best pursued by practitioners or moral philosophers. Milgram adhered to this prevailing view: "The dilemma inherent in obedience to authority is ancient, as old as the story of Abraham. What the present study does is to give the dilemma contemporary form by treating it as subject matter for experimental inquiry and with the aim of understanding rather than judging it from a moral standpoint" (1974, p. xi).

This view makes good sense, given that most individuals judge the subjects in Milgram's experiment to be either sadistic or morally reprehensible. As we saw earlier, Milgram's data suggest that neither is the case. Unless commanded to do so, his subjects never chose to hurt the learner, and they were extremely distressed by their own behavior. Such judgments are thus inaccurate. But most important, they may lead individuals to underestimate their own potential to act similarly under similar circumstances. For this reason, an emphasis on understanding over moral evaluation is welcome.

Even so, we ought not to abandon moral evaluation altogether. Unless we take the stance that subjects truly had no choice, we might still evaluate the consequences of the choices they did make and the reasoning that led to their choices. In Milgram's attempt to steer clear of moral evaluation, he took a different perspective, emphasizing the causal impact of situational factors. In so doing he avoided a negative judgment, but he may have inadvertently lent credibility to the belief that we have no choice but to obey, a belief that may reinforce obedience, when we do not know if this is the case. We have not yet tested whether we might alter obedient responses in the face of situational factors that at present are conducive to obedience.

Normative theory does not require the moralizing character of Milgram's observers. One kind of normative stance may help us both describe the world and plot our course out of the dilemmas that we encounter in it. For instance, critical theory

argues that it may be possible to use participants' own principles as the basis for moral critiques (Geuss, 1981). In Milgram's experiment, subjects faced competing principles or requirements, thereby generating the dilemma of obedience. They wished to simultaneously satisfy the principles of loyalty to authority and of responsibility to others—hence the stress they experienced both during and after the experiment. It therefore was not a simple case of not putting one's principles into practice but a case of conflicting principles that in the eyes of the subjects demanded a choice among polar actions: either obey or disobey. To the extent that subjects took responsibility for their actions, they evaluated their own actions negatively. Their dilemma was that they knew of no way to act that would allow them to resolve their conflict. But what if we were able to invent an alternative that would make this possible? Such an alternative would not only tell us something important about the nature of our social world (existing arrangements are not necessary), but it would also tell us how we might reconstruct that world to make it more livable.

What We Learn. Milgram's experiment yields knowledge that might be used at the levels of insight, action, and structural change. At each level the learning is important yet insufficient for resolving the core dilemma of obedience that Milgram describes.

• Insight as inoculation. Milgram's most important finding was that even well-intentioned people will harm others if ordered to do so by an authority. Yet most of us are unaware of this potential, showing that we have "little insight into the web of forces that operate in a real situation" (Milgram, 1974, p. 30). This suggests that, given the right circumstances, the best of us are capable of acting against others and our values and that we are unaware of this. Since forewarned is forearmed, we might now be on the lookout for such circumstances. And like some of Milgram's subjects, we might now try to deal more effectively with the value conflicts we encounter in these circumstances.

But such insight alone cannot solve the dilemma of obedience. Once we encounter a value conflict, we still must

find some way to negotiate it more effectively. Yet if Milgram is right, this would require interrupting years of socialization to "do what the man in charge says," our predisposition to turn our sense of responsibility over to an authority, and organizational norms and structures that foster both. Moreover, it is not enough just to counter these norms and disobey; it is necessary to find actions that pose true alternatives to them, that is, actions that can resolve the dilemma rather than oscillate between its two poles.

• Disobedience as an alternative action. In the extreme situations that Milgram had in mind, it makes good sense to unilaterally disobey. But in everyday life we encounter less extreme but nevertheless thorny dilemmas of obedience. We are asked to discipline a subordinate against our own judgment; we are told to support a policy we believe is discriminatory; we are warned that if we wish to get ahead, we have to conduct quick and dirty studies that violate our own standards. Each time we go along, we learn how to give up a little bit of ourselves, while telling ourselves: "That's how things work" or "I had no choice—it's publish or perish." There is much validity to these responses. It is indeed how things work, norms for getting ahead do exist, and individuals may not know how to manage these realities without continuing to perpetuate them. So in small increments we also learn how to give up some of our sense of personal responsibility, making larger increments more tolerable and more likely.

These situations can serve as practice grounds for more extreme ones. Yet under these conditions, disobedience may create more problems than it solves. As Milgram suggests, if we each act unilaterally on the basis of our own discretion alone, we may do more to promote social chaos than personal responsibility. The view that obedience is necessary would soon be reaffirmed, and pleas for law and order would reassert themselves. If so, then the unilateral disobedience that is necessary in extreme situations may also be conducive to creating them. To prevent the extreme conditions might require alternatives to existing responses that normal science cannot at present discover.

• Structural change as prevention. Milgram's results do suggest ways we might alter situations or structures, so that we

might at least modify the problem of obedience. We saw earlier in Milgram's experiment that the less distance between subject and victim and the more distance from an authority, the less likely subjects were to shock the learner. Embedded in this finding is the idea of creating structures that increase our distance from authorities. Conceivably individuals at the local level could make judgments and render decisions based on information immediately available to them at that level. If adopted, this solution would transfer a good deal of authority to the local level.

Within existing hierarchical structures, it is unclear what would lead authorities to design their organizations in ways that would diminish their control. But suppose they did. To ensure coordination, individuals would still need to negotiate and resolve the conflicting judgments apt to arise from significantly different informational bases. If polarization or working at cross purposes is to be avoided, then individuals must still find some way of *acting* within these structures that would allow them to resolve conflicting views and values. Therefore, even if these structures were adopted, effective implementation would depend on initiating new actions. Structural change alone would not be sufficient.

To genuinely resolve the dilemma of obedience, it is not sufficient to simply oscillate between the two requirements that comprise it: personal responsibility and social coherence. But in studying obedience, Milgram assumed the necessity of what exists, and consequently kept his experiments within existing arrangements and responses. As a result his experiments could not yield knowledge that might help individuals break out of this oscillation between the conflicting demands that are the basis of the dilemma. We never learn of alternatives that might better manage it, and we do not discover the deep structures that maintain it. In fact, by replicating only what is, we may learn more about how to produce obedience than how to solve the dilemma it yields.

Beyond Obedience

But what if we now wished to consider the question of whether obedience is in fact necessary? As suggested previously,

this would require that we create an alternative universe in two respects: a universe that departs from prevailing experimental contexts *and* that diverges from the real-life contexts in which obedience is the norm. In our view such contexts are consistent with Models I and O-I. In seminars designed to help students learn to act consistently with Models II and O-II, we seek to transform existing norms such as obedience and the unilateral control of the person in charge over subjects. Since our subjects are usually programmed with Model I competencies and therefore create O-I learning contexts, these seminars are a fertile domain for experiments designed to at once explore and transform the status quo.

What follows is a description of one such experiment that focused on obedience and how we might better understand and move beyond it. More specifically, it examined an incident in which a participant (a student) was faced with a request made by an authority (the instructor). In doing what the person in charge said, the participant went against his own view of what was right, and he came to regard himself as no longer responsible for his actions. Recall that these are the features thought by Milgram to capture the essence of obedience to authority. But in the action science experiment, the aim was to probe and to reconstruct those essential aspects of our social world that simultaneously yield desirable and undesirable effects. We therefore have an opportunity to see what an experiment with these aims looks like and to consider the knowledge it yields. The incident of obedience itself occurred within a larger action experiment designed to help participants see their actions and the risks of experimentation differently. It came to be called the *passivity experiment,* and it was set in motion by the instructor after he discovered a pattern of withdrawal that he thought was undermining the students' learning process. It took place toward the beginning of the seminar's second semester; we enter as the experiment began, so that we can set the context for our inquiry into obedience.

Passivity Experiment. At the beginning of the semester's third session, the instructor opened the class by saying that he wished to initiate an experiment. He then went on to make two

inferences about the group, publicly testing to see if the group would confirm them:

> Okay, another one of my experiments. What I'd like to do is start by making two attributions about this class, which I'd like to test out, if I may.
>
> One, since our time is scarce, there is an issue of justice. Most of you believe you should not take more air time than however you measure your fair share. Is there anyone who'd disagree with that attribution? [People confirm.]
>
> And another one was that Paul, when he began, had the equivalent of what many of you might have felt was a fair share, regardless of whether you agreed with the way he began or not. [Most agreed; some said he took more than his fair share.]
>
> I then said: "Who would like to go first?" Utter silence. [Paul] looked at me; I looked at him. I looked around three or four times. Paul looked around. He finally took over.
>
> I want to know: How come? What's the dynamic here that says the guy who's already had enough air time is now asked to even use more?

The instructor opened the experiment by publicly identifying a puzzling process in the group and testing to see if the inferences that comprised it were accurate. Once others confirmed them, he then inquired into the puzzle, and a range of new data came forth about how individuals saw participation in the class. To summarize, they said the following kinds of things:

> "I feel I should say something smart or the instructor will attack me."
> "I feel I'm hiding out and easing in."
> "I'm waiting for someone to make a mistake to see what the instructor will do."
> "I'm confused about where to begin."

"I felt my intervention had to be perfect, and I had no model of perfect."

"I don't want to appear stupid."

"I feel on the spot, like I'm breaking ice."

"I feel the need to warm up before making an intervention."

"I have a sense of impending embarrassment."

As these data suggest, individuals had been making and holding a range of inferences about themselves, the instructor, and what it meant to participate. In response, the instructor tried to determine what had led to these reactions, inquiring into what had happened and how the students had gone from these data to their conclusions. As he did so, he checked to see to what extent reactions were shared or differed. Through such a process participants came to learn that the instructor had unintentionally contributed to the first person's fear of being attacked, and they saw how the class had reinforced the instructor's actions by not confronting the instructor at the time.

Later on in the experiment the instructor began to build on these early discoveries by developing a parallel between how participants had responded to him and how they had responded to an authority in a case discussed the previous week. In that case a director and her staff had become polarized over who was responsible for the internal problems in a counseling program, each unilateraally blaming the other and deflecting attempts to examine his or her own responsibility. Nevertheless, of the total responses, 65 percent of the students held the authority responsible, calling the director closed, defensive, and blind to her impact on her staff. From these data and those from the class, the instructor began to speculate out loud that many group members might hold those in positions of authority more responsible than they held themselves or subordinates, leading them to react to those in power with a combination of hostility and passivity. On the one hand, they might not approve of what the authorities do, and they might privately make negative attributions about them (he'll attack me, she's controlling). On the other hand, however, they might regard authorities as primarily

in control and responsible for these problems, thus going along and remaining unaware of how their passivity contributes to the problem.

An Individual Confronts an Authority. Throughout this inquiry the instructor played an active role in bringing to the surface and examining these issues and the dilemmas they posed for everyone's attempts to learn. But unlike an experimenter in a laboratory, the instructor made his hypotheses public and enlisted the participants' help in inquiring into them. It was the instructor's view that such moves would serve to explore participants' passivity while simultaneously helping them to break out of it. But one participant, George, took a different view. He thought these moves might inadvertently exacerbate the very passivity they were intended to remedy. He thus confronted the instructor, calling his experiment into question:

George: Can I say something about process? I'd like to tie together my reactions to what just happened to the dilemma you raised about how can you help us become less passive. It seems to me that the way you shifted gears, and even overrode what I think was going to be an objection by Donna, to continue what you call "the experiment," you were exercising unilateral control.

Instructor: Yes.

George: I feel that in a more subtle way you've been doing that throughout the entire class in that you put on the agenda the issue of our passivity, thereby displacing other things that might have been on our agenda, like the case that we prepared.

Instructor: Right.

George: My sense is that to the extent that you do that, you will continue to find us behaving in passive ways, because we then learn that the behavioral cues for what happens in here come from you.

Instructor: Let's check it out with others.

In this interchange a student confronts authority by as-

serting that the instructor is acting unilaterally in directing the attention of the class members to their passivity and displacing the case they had prepared. On the basis of this inference, he then goes on to predict that the instructor will foster the very passivity he wishes to eliminate. If so, the instructor would be acting contrary to his intentions, and he might inadvertently be hindering the ability of the class to learn.

To discover whether this was the case, the instructor first asked for others' reactions instead of asserting a view of his own. In this way data could be generated on how others saw his actions and the experiment, and the predisposition to fall in line with the view of an authority might be minimized. As Milgram's data also suggest, once a peer confronts an authority, it is less likely that a group will automatically continue to do what the person in charge says, and it is more likely that they will start to question the authority (1974, pp. 116-121). In an action science experiment, the instructor therefore encourages this continued confrontation by eliciting such reactions from others.

An Authority Confronts a Student. Even so, the instructor does not automatically accept the views of participants as valid. Instead, when another student, Paul, agrees with George's critique, the instructor calls Paul's view into question in the same way he has invited the students to question his:

Paul: I think [what George says] makes a lot of sense. I think it would have been preferable if you had raised the issue of our passivity as one we wanted to discuss or not discuss.

Instructor: How do you raise the issue of passivity with people you attribute to be passive?

And he then turned to George to see if he could produce such an alternative:

Instructor: Could you give a for instance, George, of how I might have done this?

George: One way [*pauses*]. I think it's a dilemma, so let me say that.

Instructor: You said that.

George: You're in a dilemma because you have knowledge to impart to us that may empower us, namely, the perception that we are passive. In order to impart that knowledge it may be necessary to exercise unilateral control. That's possible. I think the suggestion [from Paul] that you check with us whether we felt the need to discuss [it was a good one].

Instructor: Could you produce that?

In this excerpt, the instructor inquires into the alternative put forth by Paul by asking that it be produced. He thus neither accepts nor rejects the competing claim automatically, but rather asks that it be illustrated, so that he and the group might judge for themselves whether it solves the passivity dilemma. As this suggests, breaking out of passivity is not sufficient in the instructor's view. It is also necessary to submit the assertions of the students to the same scrutiny that he submits his. By taking this line, he implicitly rejects the notion that either he or his students have the right to unilaterally impose their views on the other.

An Individual Obeys. But from George's point of view, the instructor's request was problematic, since he was unaware of a way out of the dilemma and felt he could not say so. As he revealed the next week, he privately reacted by thinking, "Yikes. What do I say?" while publicly restating that it was a dilemma. So when the instructor again asked him to produce an alternative, George did as requested, role playing the following intervention:

George: I have a sense that there's an unspoken norm in here that people should take no more than their fair share of the time. Check that out, which you did. I further sense that you think that Paul took at least that amount of time last week. Check that out. For me that points out the dilemma that I think that many people had an opportunity to speak, no one did, and therefore Paul was in a sense stepping into a vacuum that was created by—I attribute—your passivity. My sense is that poses a problem for your learning, it poses a problem for

this class. I think it would be worthwhile spending some time at
the beginning of this class discussing that issue. Do others of
you share that perception?

Instructor: [*To the class*] What's your view?

Miscellaneous voices: Sure. Yeah.

The alternative given here begins by making and testing
inferences in much the same way the instructor did. But it de-
parts at the end by explicitly asking whether the group thinks
passivity is a worthy issue to pursue. As he did previously, the
instructor responded to the alternative by first eliciting others'
views and then stating and inviting reactions to his own. He
then put forth how he saw his own actions in relation to the
group:

> My view was that you didn't have a choice
> about Donna's case. You didn't have a choice
> about the short case. You didn't have a choice
> about the third long case. Okay?
> I said, "This is not set in concrete." But no-
> body said, "I don't want to do any of the other
> two." I got assent, "Go ahead, if you're doing
> something that's useful, we'll go along with it. If
> it's not, we'll tell you so." So my view is that I was
> staying within the assent I got from this group.

The instructor confirms that he did not give the kind of
choice that George suggested he should have given, yet he im-
plies that it was not necessary. He assumed that he continued to
enjoy the assent of the group, since his agenda was not set in
concrete and no one told him to stop. The instructor and the
students thus have a different notion of choice and of who is
responsible for generating it. From the students' point of view
they ought to be explicitly asked whether or not they wish to
pursue an issue. From the interventionist's perspective, how-
ever, there was a standing invitation to confront whatever was
problematic, and this invitation was continuously reinforced. In

his view, to add to this might have been to give the kind of choice that entraps students, because it could lead them to rely on such invitations before confronting authorities. This alternative might make the passivity dilemma even worse.

Instead of acting as he did, the interventionist might have uncritically accepted the alternative of the students because, in asserting it, they had broken out of their passivity. But to do so would be to implicitly encourage the alternative that "anything goes according to the discretion of those at local levels"; and, as Milgram argued, such a norm would undermine any possibility for coordinating conflicting views, and it would surely be unrealistic in real-life contexts. So for the instructor in this case, the experiment had just begun. In this view the moves taken by the students to overcome their passivity were insufficient for solving the dilemma. They were still based on the a priori premise that all authorities are unilaterally responsible for and in control of a process that subordinates require permission to affect. As long as this framing of the situation remained, it was unlikely that the students would be able to negotiate their conflicting views with authorities beyond the classroom. So instead of accepting their alternative, the instructor continues to ask that they put forth and illustrate their views. And as they do so, he continues to critique them, to say where his own views diverge from theirs, and to encourage others to critique his views.

An Individual and an Authority Confront Each Other. In response, George did continue to confront the instructor, at one point saying that while he recognized the validity of what the instructor had done, he thought that it put him in a dilemma:

George: My dilemma is this. It's certainly valid to do [what you did]. I think that for me this may well help to attain the goal that you espouse, to make us less passive. But I think the dilemma is one of form and content. To the extent that you define the agenda, namely, our passivity, and you essentially assert your view of it, you put yourself in the position that you did very early on in the first class, saying, "You're incompetent. I'm here telling you that, because I have a broader view, more experience, and am more competent."

Although the espoused message is, "I want you to become more competent," I think the experience for me is one of being in a situation defined by you in which you are judging me. And so although I think I intellectually get something from it which may empower me (if I'm able to practice it), the experiential learning of this class is one of continuing to be passive, and your choosing to put me in that role.

Instructor: That's helpful. If that's how others of you are experiencing it, it's important for me to know, because I don't think it is possible for me to design experiments that will reduce your passivity. I don't think that doing what you produced would have [altered the group's passivity]. . . . I don't believe that your not having control over the first three sessions means you're going to be passive, especially if part of the strategy on my part is to confront the passivity. I think you believe it will make you more passive.

George: Yes.

Instructor: One way is to test it and see what happens in the interaction as we're going.

George claims that the instructor's actions are responsible for his experience of passivity. In his view the instructor has chosen to put him in a passive role, and has unilaterally defined the agenda and asserted his view of the problem. In neither case does George illustrate or test his assertions, and in fact the data are that the instructor did publicly test his view of the problem and others confirmed it (see the initial intervention). George therefore acts in the very way he says the authority ought not to act. He "essentially asserts [his] view of it," he "puts [himself] in the position [of] saying, 'You're incompetent,' " and he puts the instructor in a "situation defined by [him] in which [he is] judging [the instructor]." This is the essence of social injustice. George requires of another what he does not require of himself under similar conditions, in this instance that they share control over the situation. Moreover, George's version of what occurred is inaccurate, and it is based on his own private experience, thereby making his view of truth nonnegotiable and rendering social justice unattainable.

In contrast, the instructor continues to confront George and to defend his own position, while keeping his views open to refutation. He states his views explicitly: "I don't think that doing what you produced would have done anything." He recognizes that there are competing views, and he suggests that they submit their different views to test: "One way is to test it and see what happens." This test is of a public nature. It asks that they search for data that everyone in the group can verify, and it consequently allows others to share control over the way the situation is defined. In so doing the instructor demonstrates an alternative to passively doing what he says or unilaterally going against his views. His actions imply the norm: Submit competing claims to public test and critique. This norm provides an alternative to the "chief axiom" cited by Milgram as "do what the man in charge says" (1974, p. 138).

Interrupting Existing Norms and Defenses. As George's assertions were critiqued in this way, existing norms and defenses began to emerge. At one point George defended his view by explaining that his intervention was designed to do what the man in charge said. As he put it:

> I tried to design an intervention of the kind that I think you [the instructor] wanted. I think my instincts about how to handle this would be a little different altogether.

Like Milgram's subjects, George explains that he designed his actions, not to suit his beliefs but to suit the authority's request. But unlike Milgram's experimenter, the experimenter in this case sought not only to understand such responses but to probe them in a way that would bring to the surface and work through the defenses that maintain them. So just as he called into question the students' passivity, he now questioned the tacit rules embedded in the way George acted:

> But there are two problems with your presentation. One is that the foundation of it is what you feel. [And the] other is that I've said, "Would you produce this?" And you produced something

that you thought I wanted, not something that you
would do. That says to me: If that's your psycho-
logical state and if you now put that in your gener-
alization, then I think that people like you would
in fact not learn from behavior like mine.

If you're so programmed to do for me what
you think I want, even if it's not yours, then I'm
not sure I'm ready to trust your feelings, unless
you are ready to say, "These are the feelings of a
person who when asked to produce something pro-
duces something that he thinks the professor
wants, even when it's not what he wants."

The instructor suggests that there are two rules embedded
in George's method of putting forth his position. The first is
that he bases his views on his feelings alone, and the second is
that he designs what he does to fit what the authority wants.
He then builds on this by predicting that such rules will make it
difficult for George to be trusted or to learn. The instructor's
moves are thus aimed at extracting the tacit rules embedded in
participants' actions and at showing how they necessarily lead
to consequences that they themselves cannot accept. To the ex-
tent that individuals confirm this logic, they usually abandon
their position and reconsider the norms or rules embedded in
their actions. But ordinarily they do not confirm such logic
without first defending their own, and such was the case with
George. In response to the instructor's critique, George mo-
bilized several lines of defense, each one deflecting his respon-
sibility for the actions and outcomes that the instructor had
described. Yet each time George brought forth a new line of
defense, the instructor rendered his new position unacceptable
by George's own standards.

To illustrate, George's assertions are followed by the in-
structor's moves to critique them. Notice that each time the
instructor points out a gap in George's position, George
switches to a new one and the instructor follows by pointing
out new gaps:

George: I felt you were asking me to model Paul's suggestion. I wasn't endorsing it.

Instructor: What does it take to get you to never be put in a situation to endorse something that isn't you?

In his initial defense, George explains that he simply modeled Paul's suggestion but did not endorse it, as if this made it acceptable. But when the interventionist questions George's willingness to endorse something that is not his, George evokes a second line of defense in which he asserts that he had no choice but to do so:

George: You asked me to produce Paul's suggestion, not my own. If you had asked me to do my own, I wouldn't have acted as I did.

Instructor: You could have said, "Sorry, I'm not going to produce what he said. I'll produce what I want to produce."

Here George defends his actions by saying that he had no choice, but the instructor points out how George might have refused, thus disconfirming George's assertion. Once George sees this, he abandons this second line of defense and takes up the defense that he was only doing what is socially acceptable:

George: I think I can account for what I did as just being courteous enough to give a response, rather than a program to do what you want.

Instructor: If you know me, that's about the most discourteous thing you can do. The most courteous is to say, "You're asking me the wrong question."

By appealing to social mores, George evokes yet a third defense. But by saying that he regarded it as an insult, the instructor again makes George's defense problematic, since he shows how George's actions yielded the opposite of what he had intended.

This quick switching of defensive moves requires a good deal of fancy footwork, yet it should also suggest the consistency of George's actions. That is, even though he switches positions, a consistent logic runs through them. In each case he is in one way or another asserting either that he was not responsible for what he did or that what he did was the right thing to do. His first defensive position was that "it's not mine"; here, he simply tried to disown his intervention. When the instructor then showed how this position itself was problematic, he switched to the defensive position of "no choice," attributing responsibility to the instructor. And when the instructor disconfirmed this, he then switched to the position that he was "doing the right thing," simply being courteous, thus asserting that his actions were not as problematic as the instructor believed. But as we saw, the instructor pointed out that if that was his intention, it had backfired. With his defensive moves blocked by his own criteria, George begins to look inward for his own responsibility but only in glimmers. He reflects: "I didn't want to answer the question but somehow felt required to, perhaps because of some tacit theory of learning." But as illustrated in the dialogue that follows, it is not until his peers start to identify gaps in his position that he begins a process of actively reconsidering his actions.

Peers Confront a Peer. After much discussion, George continued to take the stance that, regardless of his alternatives, his initial critique had validity. At this point his peers began to question whether his assertion could stand up in light of what had actually happened during that class:

George: [Summarizing] : I had a criticism of what you did. And I think that I want to maintain that criticism regardless of the viability of any of the alternatives that I can come up with. The criticism was "people don't become less passive by your telling them that they are passive."

Tim: Maybe some people do, some people don't.

David: Yeah, I guess that's my reaction. The data for the rest of this meeting, at least the way I've seen it, is that a number of

people have in fact become less passive as the meeting was worn on.

George: Not, I don't think, because of what the instructor did.

David: How could you possibly separate that out?

George: Well, by looking at when they became less passive.

Joe: What would you attribute that to, if not to the instructor?

George: Can I ask you [David] first when you think that happened?

David: It's puzzling to me, George, because my response actually when the instructor accused us of being passive was not to entrench. But was kind of, along Model I lines, "Oh, yeah? If you're going to point the finger and tell me I'm passive, I'll show you." So it would be a correct strategy for provoking someone like me to get involved.

Vince: Except that if you have to wait for the other person to provoke. The warning for you, Dave, has to be to see what happened, and say: Why did I have to wait for the instructor to provoke me, because if the next time someone has to provoke you, then you haven't learned very much.

David: Well, consider this. What was most uncomfortable was a certain recognition. He asked me to confirm something: "You wanted to cover your ass and bow out and withdraw from the group." And I had to confirm that. Once I confirmed it, I didn't like it. It wasn't something I was going to accept in myself. So, that was a challenge for me to involve myself. Once it was surfaced and discussed, it became more difficult for me to be passive.

This interchange starts out with George reiterating his claim, but this time a peer confronts him as if to say: "You might be right, you might be wrong, but we don't know yet." From here they go on to inquire into each other's views and to describe their own reactions in ways suggestive of the alternative norm put forth by the instructor: Submit competing claims

to critique and test. For example, once Tim opened the possibility that George might be wrong, David came in and cited some data to suggest that this might indeed be the case: "A number of people have in fact become less passive as the meeting has worn on." Then when George asserts that this is not attributable to what the instructor did, others do not counter with an assertion of their own but ask how he arrived at that view: How could you separate out the influence of the instructor, and to what would you attribute the decrease in passivity?

When George does not answer but inquires further into David's position, David does not resist the inquiry. Rather he reports how he himself reacted to the instructor by feeling, "Oh yeah? If you're going to point the finger and tell me I'm passive, I'll show you." He thus describes at a low level of inference the way in which the instructor's actions affected how he felt, only then concluding that it was a "correct strategy" for someone like him. Once David's reactions were made public, they could be readily critiqued, and another participant, Vince, did so by pointing to their implications: David might have to wait to be provoked into action. Rather than deny this possibility or switch positions, David remains open in the sense that he stays with his position and describes it further, thus defending it in a way that provides others continued access to what it was based on, so that they might accept or reject it independently.

Thus, defensive maneuvers and the predisposition to automatically go along are not as evident in this protocol. Participants provide others access to their reactions rather than cover them up. They not only put forth their views but describe the data on which they are based. And most important, they do not prematurely accept or reject either position but critique and inquire into the competing views being considered. In so doing they follow the same norms of action science that the instructor follows. But because they are peers, they can provide data that the instructor cannot. Moreover, because of George's stance toward authority, their views may carry a weight that the instructor's do not.

An Individual Reflects on His Own Responsibility. In the week that followed, George listened to a tape recording of the

session and considered what he heard in light of the questions posed in the experiment. As he did so, he began to see to what extent he had "laid a trap" for himself. He saw that, rather than acknowledge that he knew of no alternative, he had chosen to cover this up by producing one that approximated Paul's suggestion, thereby setting in motion the rest of the process. And as he listened further, he also saw ways in which the instructor had contributed to his defensiveness. He wrote down what he had discovered and returned to the class the following week to reflect publicly on his new understanding of what had occurred. First, he described the reactions that he had kept private the previous week and how this had led him to lay a trap for himself. He then identified ways in which the instructor had not fully explored his own position and how this had contributed to, rather than interrupted, his defensiveness. And, finally, when asked what the instructor should do differently, he proposed an alternative of his own. In this session, then, he focused on his own responsibility as well as the instructor's, and he made his own views confrontable as he confronted others. By doing so he broke out of the responses that had reinforced his unawareness the previous week, and he helped the instructor to do the same.

Results. But what are we to make of the outcome of this experience? The aim of an action science experiment is to describe and to transform those aspects of our social world that present us with blind spots, dilemmas, and constraints of which we are unaware. In this particular case the instructor wished to discover and to unfreeze what had caused the students' passivity, that is, what had caused them to do what the person in charge said, even when it went against their beliefs, and then to hold the authority responsible for what occurred. Implicit in such aims are the criteria by which we evaluate the results of this experiment. We thus ask: Do participants and instructor show in their actions and reflections a new way of seeing themselves and others that will enable them to transform this passivity?

The data on what happened with George and his peers suggest that the experimental hypotheses were affirmed. As a

consequence of the instructor's interventions, individuals came to see themselves in a new way. David recognized his passivity for the first time, while George came to see how he had "laid a trap" for himself. Neither discovery was acceptable to them, but both ended up taking responsibility for his behavior. They confirmed the instructor's attributions, and they did not think that what they discovered was solely a product of what the instructor said or did, despite George's initial efforts to make them do so. Once they were aware of their responsibility for these results, the unacceptability of what they saw provided the impetus for new actions. David became more active, and George made his views more vulnerable to testing. We therefore see evidence of new understandings in their actions, as they not only report new insights but act consistently with them.

At the same time the experiment yields a rich description of the participants' passivity. By calling into question their passive responses, the instructor discovered that students tacitly framed authorities as significantly more in control and responsible than they themselves were. Moreover, they assumed a priori that this view was correct. With the authority's role thus framed, we could see how their own role would become that of passive recipients of the authority's actions. Later on, as the experiment unfolded further and George responded to the instructor's critique, we then uncovered some of the tacit rules that both followed from and helped to maintain a passive framing of one's role—rules such as "Do what the person in charge says" and "Base your assertions on your private experiences alone." Finally once these rules themselves were brought into the open and challenged, defenses that might otherwise have remained hidden were mobilized, as George deftly moved from one defensive position to another: "It's not mine," "I had no choice," and "I was just doing what is socially acceptable." Yet equally important from an action science point of view is the discovery that individuals can begin to enact a very different norm, one that asks them to submit competing claims to public test and inquiry. By the end of the experiment more participants had come to take an active role in critiquing and inquiring into competing views, no longer simply accepting or rejecting them.

This experiment at once brought to light the deep struc-

tures maintaining obedience and sought to enact a norm that might render obedience unnecessary for social coherence. The results of the experiment suggest that it might be possible to enact very different authority relationships from those assumed by Milgram (1974) to be necessary. If so, it means that the dilemma posed by the conflicting requirements of obedience and social cohesion might be better managed. But to explore this possibility requires not only that individuals discover their existing responses but that they try out fundamentally new ones, and this in turn requires that the researcher create conditions that go beyond the norms of mainstream science. He has to create an experimental universe that adheres to fundamentally different norms from those found either in the laboratory or in the work place.

In What Sense Was This an Experiment?

The first purpose of experiments is to induce individuals to act in a predicted manner or to choose not to do so. In this sense the episode just described was an experiment because the instructor acted in ways that he predicted would help to reduce obedience or would permit individuals to choose not to alter their behavior.

The second purpose of experiments is to produce empirical generalizations that remain valid beyond the experimental context. The strategy is to formulate hypotheses that, if not disconfirmed, become empirical generalizations. The hypotheses of this experiment may be formulated as follows: Under Model II conditions it is possible to reduce individuals' Model I automatic predispositions to unquestioned obedience in such a way that the participants will report that their sense of order and governance within the seminar are at least not harmed, and probably strengthened.

This hypothesis requires that we can establish that at least three conditions occurred. The first of these is to show that the participants (read subjects in the experiment) did act with unquestioned obedience toward the instructor when such obedience was not required by him. The transcripts do show, we suggest, that the participants did act in such a way and that

many of them admitted doing so. Moreover, some not only admitted acting obediently but held the instructor at least partially responsible for their behavior.

The instructor explored the possibility that he had contributed to the students' passivity because the experimental manipulation required that he not induce obedience in the participants. And, after critical inquiry, the group came to hold themselves responsible for their obedient actions.

The second condition is to show that the instructor behaved consistently with Model II. In this particular instance this required the instructor to confront and inquire into obedient responses, while exploring the possibility that through his actions he might be unknowingly fostering such responses.

The third condition is to see whether the participants' obedient responses were reduced, while not simultaneously reducing the order and governance of the group. This task makes this experiment more complex and difficult than Milgram's. In most experiments it is assumed that subjects have the skills to produce the responses required by the experiment. For example, Milgram's experiment depended on the subjects' being able to understand orders, observe the individuals in pain, and carry out whatever decision they chose: to shock or not to shock. These skills are so obviously held by most adults that Milgram, correctly so, took them for granted. Experiments are not usually designed so that they require the subjects to use skills they do not have.

In our experiment to reduce obedience, the opposite was the case. The participants were skilled in producing, not reducing, obedience. If the experimental manipulation (creating Model II conditions) was to succeed in reducing obedience, the participants would have to acquire abilities that they did not have. These included the abilities to become aware of their automatic actions toward obedience and to learn action skills that they did not yet have, such as confronting authority in ways that ensured personal responsibility and the continued self-governance of the group.

Moreover, becoming aware of their automatic actions and their lack of skills to create the behavioral worlds they value tends to be threatening to individuals. This means that the participants in our experiment would also have to become aware of

the defenses that they used when they were threatened. In other words, in our experiment we could not take for granted that the participants would have the skills to reduce obedience. They would need help to learn these skills and to create a systemic culture that reinforces them.

It is possible to assess the degree to which these requirements have been met, and we have tried to demonstrate this here in three ways. First, we can analyze the transcript to assess the degree to which the instructor acted consistently with Model II, as well as the degree to which he was willing to explore the times when he may not have been doing so.

Second, it is possible to analyze the students' initial reactions to the instructor's attempts to assess the students' unawareness of their obedient responses, as well as their defensiveness toward becoming more aware and reducing such responses.

And, third, it is possible to follow the interactions over the course of the experiment to show how some students eventually moved toward reducing their unquestioning obedience, as well as how the group norms about obedience changed.

If we wished to do so, we could continue this type of analysis in future sessions during the seminar, because there will be other opportunities to test what individuals learn. Unquestioning obedience is not likely to be reduced significantly in one episode. We could expand these studies to assess the external validity of these findings by observing the students in other sessions where they face authorities or where they themselves may take the role of an instructor who is trying to help others reduce their predisposition toward unquestioning obedience.

Another requirement embedded in this condition is to show that group norms conducive to an orderly society, self-governance, and double-loop learning have been maintained while obedience has been reduced, so that we do not have the kind of disintegration of the group's social fabric predicted by Milgram.

We could determine whether this requirement has been met in three ways. First, we could interview the participants to obtain their views on these matters. Second and more important, we could observe how the participants acted to maintain their group processes, the quality of the governance that they

created, and the amount and quality of double-loop learning that occurred. It is not difficult to obtain information on all these points if tape recordings are made, as they were in this case. Observers can be added if they are available, but we have found that observers are not nearly so demanding as participants whose competence depends on the results of the experiment. A third resource to gather such data are the participants themselves. They are centrally concerned with the answers to these questions and how each contributed to them. For example, students can be broken down into smaller groups. They can listen and analyze the tapes. They can subject their own analysis to validation by others. They will probably be especially careful to develop valid data because such data are the basis for deciding how well they are learning, what they must overcome to gain the skills that they seek, and what kind of group culture is required if the group is to be a viable context for learning. In short, the data required for their learning will be consistent with the data required to answer the three questions that we asked about the internal and external validity of the experiment.

There are two other questions that may be raised regarding the internal validity of the experiment. First, might the passivity and unquestioned obedience have been an act put on by the students? This seems highly unlikely since they were trying to learn to overcome these kinds of behavior. Moreover, their initial defensive reactions suggest that they were not playing games. Also, the fact that some learned faster than others and then helped others to learn suggests that it was unlikely that they had conspired ahead of time to create the passivity experiment. The reactions in the follow-up experiment support these inferences. Second, how do we know that the changes would not have occurred without the Model II interventions (the experimental manipulation)? As shown elsewhere (Argyris, 1982), individuals are not able to produce actions consistent with Model II even though they understand the model, want to act consistently with it, have watched several groups attempt to do so, and have made several attempts themselves. Learning Model II requires supervised practice.