

Master's Series on Field Research

A series of interviews with major figures in field research conducted in the early 1980s by
Peter Blanck

Transcript of an interview with Robert Rosenthal



Burton Blatt Institute
SYRACUSE UNIVERSITY

Peter Blanck: 12/9/81. Bob, I'd like you to start by telling us some of the things that led you to your field research experience, specifically, "Pygmalion in the Classroom." What were some of the critical issues you faced - some of the ethical problems you faced in that study?

Robert Rosenthal: OK. We really got into the Pygmalion research via the experimenter expectancy research, so I've got to back up a little bit and sort of lead us into it. What happened was that I had done a doctoral dissertation at UCLA that I badly screwed up. And, one of the results that came out of that was the possibility that an investigator could quite unintentionally bias the results of his research. That somehow the expectations of the experimenter could come to serve as a self-fulfilling prophecy. Well, we did some experiments at the University of North Dakota that showed that. What we did was to create different groups of experimenters with different expectations for how their research subjects should respond, and these subjects responded pretty much as they had been expected to. So I wrote a paper in a journal and ended the paper that had described these studies by saying that if in a psychological experiment, experimenters get the results that they expect to get, then maybe, also in the classroom pupils will respond as they're expected to respond. A school principal by the name of Lenore Jacobson, South San Francisco, California, read this journal article - it was in the American Scientist - it was not in a journal that I thought school principals ordinarily read. But, she wrote me a letter and asked me whether I really meant it, whether I thought that it was a worthwhile thing to do, and whether was I planning to do the experiment? Well, I guess I wasn't planning, specifically, to do that experiment, but Lenore Jacobson offered me the use of her school. So - we did it there. So, the Pygmalion experiment really turned out to be just a further extension, or a further replication of experiments that we'd been doing with human and animal subjects in laboratory settings up until then.

Peter Blanck: And was that the first, your first experience with field research methods?

Robert Rosenthal: I guess it was really the first time I'd gone out into the field to do research. Most of my work before then had been with human and animal subjects in laboratory settings, first at UCLA, then at the University of North Dakota and then at Harvard University.

Peter Blanck: And what were the most striking and interesting differences between the type of field research you did, or field research in general, and experimental lab work that you had been doing, maybe as an experimentalist, you can tell some of the things and critical problems you face that were unfamiliar with in the field setting.

Robert Rosenthal: Well the technical problems are the problems that we all know about. That is, in Don Campbell's terminology we have the limited external validity in the lab setting and...

Peter Blanck: Going out into the field for the first time, what were some of the more interesting stories you have to tell about your first field research experience, and how did that differ from your previous laboratory training? And, more generally, how do you think

- what are the types of problems field researchers face as opposed to more experimental laboratory approaches?

Robert Rosenthal: Well, I guess the first thing that I hadn't been prepared for was the extreme importance of establishing relationships with real people in the real world. If you're doing lab research in the university setting, everything is sort of geared up for you - there are subject pools available or there are lists that you get your subjects to sign up on - but if you have, let's pause - collect yourself.

Other: All right.

Peter Blanck: OK. I'm sorry, OK.

Other: Is it a research joke...?

Peter Blanck: No - it's just poor laughing control. And, the question being, how are your lab experiences basically different than field experiences, and generally for practitioners in both fields?

Robert Rosenthal: It seems easier in many ways to do lab experiments. It's easier to do them because everything is sort of geared up to do them. If you're going to do research in the field, you sort of have to go out into the world and make contacts with school principals, with school superintendents with hospital administrators and so forth. And if you can manage that, you can do the research, and if you can't manage it, you can't get the research done.

Peter Blanck: OK. I wanted to ask you, Bob, about some of the ethical issues you face in field research, specifically in a study like "Pygmalion." If you could tell us more specifically, first of all, some of the ethical problems you face doing a study like Pygmalion, and, more generally, how do you think ethical issues are - different force in the field as opposed to the lab?

Robert Rosenthal: I think that, in the lab, subjects in psychological experiments have a certain degree of self protection built in. They already know that psychologists don't tell the truth, the whole truth, and nothing but the truth, and they're kind of geared up to not believe too completely the things that they're told. Out in the world, people may not even know you're a psychologist, and, if they do, they don't know that you're likely to withhold information from them, or even to lie to them. I guess that the ethical issue that I saw primarily in our Pygmalion experiment was the fact that, in order to do the experiment, we had to lie to the teachers. That is, we had to tell the teachers that certain of their children had scored on a test in such a way that they would show unusual intellectual blooming, or spurting, in the coming academic year when that really wasn't true. We would just pick those kids' names at random. We picked their names out of a hat just to see whether the expectations, the favorable expectations the teachers had for children would actually improve the children's intellectual functioning. As it turned out, it did. So, I didn't see very many ethical problems in terms of the children involved, since they came out smarter

than they started. But I certainly did see an ethical problem with having to lie to the teachers.

Peter Blanck: And how did you as a psychologist cope with that problem? Did you follow up and - go back and tell the teachers the truth, or did you give the pupils who might not have received extra experience or extra teaching methods that extra bonus, or - what sort of...?

Robert Rosenthal: Well, we have to distinguish, I think, between those kinds of field research in which there are potential benefits only or where there are potential benefits and potential harms. In our study, as far as the children were concerned, I don't think there were any potential harms. There were some potential benefits for a subset of the children, but there was no evidence to suggest that the children whose names had not been given to the teachers would in any way suffer as a result. Indeed, we did some analyses that showed that the children in the classrooms of the kids who'd made the greatest gains from having been designated as intellectual bloomers, those control group kids also gained the most. So, the more benefit there had been to these special children, the ones that we designated as potential bloomers, the more benefits there were in intellectual gains also for the other children in the classroom. The real ethical problem, I think, in that study did come from having to lie to the teachers. When we went to the teachers two years later, to debrief them and to explain what the experiment had been all about, they certainly didn't seem very upset. Some of them did wonder why we had spent all this time and money to prove the obvious. They sort of said, "We know that teacher expectations are important. Why didn't you just come and ask us? We could have told you that."

Peter Blanck: But that in itself is an interesting point - to prove the obvious, you just said. Do you think that a lot of field research seems to prove the obvious?

Robert Rosenthal: That's right.

Peter Blanck: What sort of issues, do you think determine a good field research program as compared to a sort of mundane one. I mean, to prove the obvious, might, on one hand, seem uninteresting, and, on the other hand, in your instance, it's very interesting. It had a tremendous impact on educational systems.

Robert Rosenthal: Well, I think your question is a hard one to answer because it calls for a prescription of how we should do field research, and I'm not sure I know how to answer that for - for anybody else, but I can give you my own biases about it. First, let me say that I think a lot of field research and a lot of lab research winds up proving what everybody already knows. The point is that people already "know" so many things that aren't true. And, so, it's very much worth while to do the field and/or lab research to show that what people believe to be true either is or isn't true. That is, just because it's a widely held belief doesn't mean that it has been well established from a scientific point of view. Specifically, how to do good field research - I don't know. I sort of feel best about field research that is experimental. That is, where you have an independent variable that you can experimentally manipulate. It's not always possible in field research to do that. In the

Pygmalion experiment, it was. That is, we randomly chose certain children to be in the experimental condition, randomly chose others to be the control condition and you can draw very strong causal inferences if you can do that. And you can draw those strong causal inferences if you're out in the field or if you're back at the lab. If you don't do random assignment of subjects, if you don't have experimental control over your independent variable, you really aren't in a strong a position to draw causal inferences whether you're out in the field or back in the lab. I think that there's a kind of a feeling that somehow if you're out in the field, it's OK because it's real world - it's OK to be a little more sloppy. And I guess I don't believe that. In medical research - a lot of biomedical work, for example, has almost institutionalized resistance to doing well controlled randomized experiments. Most of the research doesn't tend to be of that kind, but it could be. But people somehow feel that if it's in the real world, it can be sloppy. There's no real reason for that.

Peter Blanck: Of the methodological practices and programs that are available to the typical lab experimenter, what methodological approaches or combinations of approaches do you think are most relevant? Of course, it's determined by your study, in part, but maybe using Pygmalion as a specific example, ten or fifteen years later, what types of cross measures, or cross techniques, might you have used?

Robert Rosenthal: I guess I'm not entirely clear. Do you mean - could you say a little more about? -

Peter Blanck: Sure - When we think of field research, typically, at least, in this area, we think of case analysis, or interviewing. Most field researchers are not typically aware of many methodological techniques, whether statistical or randomization procedures available to most lab researchers

Robert Rosenthal: I see what you mean. Yeah, to me, the major distinction between the field and lab is that the lab is in a much more hermetically sealed environment, as it were. You as the investigator have control of a lot of variables. You can control the lighting, you can control the temperature, you can control the instructions that you give your subjects, and so on. And every subject is pretty much in the same kind of situation. In a field setting, you are where the people are to a greater extent. You're in a hospital, you're in a clinic, you're in a business organization, you're in a school. It's the hurly-burly of everyday life. You're probably going to get greater variability. You're going to get greater variability because grade school children are more variable than Harvard students. You're going to get greater variability because you may not be able to do all of your interviews in exactly the same kind of room. You may have different people doing the interviews. If you're studying psychotherapy, you're going to have many different psychotherapists, not just one or two data collectors or experimenters. So, you're building more variance into the system, but the basic principles of good experimental design are not one bit different for field research than they are for lab research. I think it's a mistake to think that because it's field research, anything goes, that it's OK to be sloppy if it's just field research. In some ways, you have to be a better experimentalist to do good field

research because you have to get out there and compensate for the lack of control, maybe by picking more powerful independent variables.

Peter Blanck: You've done a lot of work on bias in research, experimenter bias, for example. Typically, a field researcher is his own researcher and it's very hard to remove yourself because a large part of getting access in the field is as the main researcher, developing a relationship, as you say, with these people. What sort of things do you think typically bias researchers who have this very close established relationship with their - with their participants in the field?

Robert Rosenthal: I think that's a good point. I think it may be harder in many field contacts to maintain blindness and double blindness in your experimental designs. By that, I mean what the psychopharmacologists and the biomedical researchers have been finding out for many years. They haven't recognized sufficiently - and, that is, to a great extent, even such drugs as morphine worked better when the doctor administering the morphine believes it to be morphine. A Harvard Professor of Anesthesiology showed that some - some years back. But, I think that the key thing is to stay blind as long as possible or become blind as early as possible on the experimental condition that your subjects are in. Now, if you're doing a psychotherapy study - and I regard that as a kind of field research - if you're doing it out in hospitals and clinics, the therapist can't be blind to whether his or her patient is getting psychotherapy or not. There is no way to keep the therapist blind, but, in a way, you could regard the therapist as just another subject, and it would be at least possible, if you're the data collector and you do the testing of the patient, it would certainly be a good idea, then, if you were blind as the data collector to whether that patient had been in the actual therapy or in the control therapy condition, or, if you're the person doing the testing of the school children who have been exposed to a teacher with high or ordinary expectations, it would certainly be well if you kept yourself blind to whether the particular kids had been in the experimental or the control condition.

Peter Blanck: If you were going into the field today, knowing what you know, ten years down the road from your first major study in the field, how would you change that study - particularly "Pygmalion"? What sort of things would you do differently, even in establishing a relationship with Lenore Jacobson?

Robert Rosenthal: Well, Lenore Jacobson is a one in a million type person. To begin with, she established the relationship with me and sort of got me into this research part of - I'd never done educational research, so I certainly owe going into the field and going into the schools entirely to Lenore Jacobson. Because she was so great, there's probably nothing that I would need to do differently. There are things that I would have done differently, maybe, in the write-up of the Pygmalion experiment. That's because I'm older and wiser now, but I think the research itself - I would do just about the same way. But it was partly out of blind luck the first time that is turned out so well - luck, that is, of hooking up with Lenore Jacobson.

Peter Blanck: Have you done any field research since then? Or work in formal organizations - industrial...?

Robert Rosenthal: I've done - or I have been involved in - research in hospital settings. I've been a research consultant in psychotherapy projects. I've worked in settings with spinal cord injured patients. I've worked in my own research in psychotherapy mediation variables. That is how therapists might communicate expectations to their pupils. That's some of the research, in fact, that you and I are now doing together.

Peter Blanck: I want to switch gears a little bit and tap some of your expertise on statistical and experimental design approaches, specifically related to field research. And, as your writing has shown recently, there is a growing body of literature suggesting meta-analytic techniques for describing and analyzing data. I wonder if you could talk a little bit about that and maybe about some of the approaches within that realm that are available to the field researcher, and where you see that as sort of helping field research, quantifying and qualifying data in the future?

Robert Rosenthal: The idea of summarizing large bodies of research is a very fundamental one. It's always been a complaint of the social behavioral sciences that we cumulate our knowledge so poorly, compared to physics, or compared to chemistry. Somebody does an experiment, everybody jumps on the bandwagon, does the replications, and nails it down very quickly. But in psychological research, it seems almost as though every doctoral dissertation, every journal article starts from scratch. Sure, in an introductory paragraph, they pay some homage to some other researchers who have published similar types of things. But, basically, nothing seems to be known. So, it's a field that cumulates its knowledge poorly. Various techniques of meta-analysis, as Gene Glass has called it, have been developed for summarizing large bodies of research, and one basic way of approaching that is to think of every experiment, every study that's been done bearing on a particular research question as an "N" of one - as a sample size of one - and do summarize all the data that are available to bear on the research question. So, if you are doing psychotherapy research, or meta-analyzing psychotherapy research as Gene Glass did - he summarized with his colleagues five hundred studies, control studies, of psychotherapy to see what on the average was the size of the effect of psychotherapy. Don Rubin and I have done analogous meta-analyses on the effects of interpersonal expectations for some 345 studies, and we find, similarly, that there are some substantial sized effects averaging over these hundreds of different experiments. I don't think that field research is any different from lab research in that regard. The importance, I think, of the field research is that it adds a new dimension, it adds a dimension of generality that you could never have if you did all your studies in the lab. The lab is a wonderful place to begin to nail down your procedures and to get as clear a look as you can at what may be going on. But, until you take it out in the field to cross-validate it or to replicate it, you really are coming up short. So, it's very important, I think, to go into the field with a replication. But the statistical techniques are identical, whether you're doing it in the field, or doing it in the lab.

Peter Blanck: OK. All right. I have a couple further questions about research methodology. Specifically, your impressions on how we can better get a handle on issues of reliability and validity with regard to the typical field research experimenter, maybe giving the typical field researcher a better understanding of why these issues are important.

Robert Rosenthal: I think here's another example where too often though we hold to very high standards of research when we do it in the lab. We somehow think that once you get into the field, anything goes. We all know about the reliability and validity of our measuring instruments when we use them in the lab, but there's no reason why the instruments that we use out in the field shouldn't be subject to the same kind of rigorous analysis as to their reliability and validity. I refer, for example, to the interview. The interview seems to be sort of a non-method, or a non-technique, because anybody just sits there, or stands there, and talks to somebody. But it needn't be that way at all. One can assess the reliability of an interview, or of a particular interviewer, and one can assess the validity of an interviewer or of a particular type of interview. Holtz and Laborski, in their field study of psychiatrists in preparation give good a example of treating the interview just like any other measuring instrument, like a Wechsler-Bellevue - or a Bender-Gestalt, or a Rorschach. Those are all instruments subject to considerations of validity and reliability, and so is the interview. And, I think that we should probably do more to try to standardize interviews, or to calibrate interviewers.

Peter Blanck: Could you be a little more specific, in that, how would you - give an example of it, one way to calibrate an interviewer?

Robert Rosenthal: One way is to calibrate interviewers - it's a relatively expensive way. But field research often is - would be to have - a number of the same people interviewed by a number of the same interviewers. So, in some kind of a counterbalanced design, you might have several business people interviewed by several interviewers, so that you would have two or three people doing interviews with the same group of interviewees, primarily, initially, just to establish the reliability and/or validity of the interviewers. You might find, for example, that the kind of information obtained by one of the interviewers is consistently different from the kind that is obtained from another.

Peter Blanck: And, I guess that would apply for observation techniques - even simpler, you just sit down a couple of observers and have them watch the same....

Robert Rosenthal: Exactly. Exactly, if you're out in the field watching naturally flowing behavior, and if you are trying to code gross body movements, or if you are trying to code psycholinguistic aspects of speech, you'd need more than one observer make the codings in order that you could establish the reliability. Are you coding something on which observers can agree? I mean, the reliability doesn't have to be super great, and I think it would be a mistake to think that we should follow the edicts of our undergraduate psychometric textbooks that have decent reliability - you have to have reliability coefficients of .8 and .9. I think that's absurd. Very often, depending on the purpose, you can do very well with reliabilities of .2 and .3. But, when that happens, you need to know about it in order to compensate for it. One way, of course, to compensate for it is to have more interviewers, or more coders, or more observers, or more judges. An interesting bit of history is that Gordon Allport, my colleague for many years here at Harvard, and one of the founders of the non-verbal communication business in the 1930's, got out of that business because his psychometric friends told him he wasn't getting reliability that was good enough. If they hadn't given him that bad advice, he might have discovered, indeed he probably would have discovered, a lot of the things that we are finding out, thirty, forty,

almost fifty years later, things that Gordon would surely have discovered if he'd, in a sense, been allowed to play a little bit more, and not worry so much about the reliability. Reliability, if it's not zero, can always be brought up by adding observations, whether it's more judges, or more codings, or more observations, or the like.

Peter Blanck: And if you could speak to that issue with regard to validity a bit - how would that...?

Robert Rosenthal: With validity it's sometimes very hard because you don't know whether you should use other judges as the criterion for whether you have validity. Let's take the field research case of psychiatric contacts. You might use a psychiatrist, a clinical psychologist, a psychiatric social worker, a psychiatric nurse as judges, maybe to make ratings of psychopathology. What would you use then as the criterion of how disturbed the patient really is? Well, you probably have to use something like the pooled judgment of the bunch of other psychiatrists, or psychiatric nurses, or psychiatric social workers, or clinical psychologists, because there is no better criterion than the pooled judgment of experts. Yet, any one of those experts would have their validity or reliability determined pretty much in the same way by correlating that a person with either another individual or with the mean of all the other individuals making the same kind of judgment. So, sometimes it's very hard. There are times when it's easier to make distinctions if, for example, you're developing a test for salespersons. If you can get good agreement among interviewers in their predictions of who'll be a good salesperson, then that's establishing the reliability of the interviewer. There, however, you might have a much better criterion than the pooled judgment of the interviewers and the subsequent sales figures. And, if it turns out that one interviewer consistently predicts more accurately than the others, the sales figures five months later by the people that he or she's been interviewing, then we would have to say that that interviewer is a more valid interviewer.

Peter Blanck: It sounds like, also, that if we as field researchers employed some of these techniques, it would help us in designing the questionnaires which we are more and more using these days with large organizations. Maybe you could tell us a little about that?

Robert Rosenthal: Well, I think the design of questionnaires is if not quite - a science, certainly a highly developed art and technology. People at the Institute for Survey Research at the University of Michigan and at analogously elegant research centers are very expert and very experienced in drawing up questions that are likely not to be misleading, that are not so likely to be misunderstood, and it takes a lot of practice, and it takes a lot of trying them out, that is, the particular test items on populations like those you're finally going to administer the questions to, before you have a decent questionnaire. Just asking a bunch of questions - doesn't make for a scientific instrument.

Peter Blanck: I'd like to ask you, just simply, what is your definition of field research?

Robert Rosenthal: I probably can't really give you a good definition of field research, or, if I could, it might be a list of things. But I guess I can start out by telling you something it isn't. The one thing it isn't is just anything done in the field. Just because it's done out in

the field doesn't make it research, doesn't make it science. Like any other good science or good research, it has to be carefully planned, there has to be a goal, you have to know what you're looking for, to some extent, though, at the same time, being open to surprises, and, you have to know how to collect data, you have to know how to do interviews, you have to know how to do questionnaires, you have to know how to administer tests, and you have to know how to evaluate all of these instruments as to their reliability and their validity. So, while I'm having a hard time really defining field research, I guess I can say what it isn't, and it isn't just any old thing done out in the field.

I guess a distinctive feature of field research is that it is done with people, or in contexts that most people would regard as somehow more real life than college sophomores in a lab at Harvard or in a lab at Ohio State, or in a lab at UCLA. It's somehow seen as closer to where the real action is, and I think that is a distinguishing feature of field research. Not all field research need be experimental, but I guess my private plug would be to make much more of it experimental than currently is the case. I think that, with many studies, people say, "Well, it's in the field so we have to settle for second best." I think that's not justified. I think we often could do really elegant experimental work with more noise in the system but without quite the threat to being able to draw decent causal inferences. So, my vote would be for more strong causal inference in the field by doing more randomized experiments.

Peter Blanck: Would you say that, in your definition, randomization is the key to experimental research?

Robert Rosenthal: It's the key. It's absolutely the key. At least half the time when people tell you they can't do the randomization, it's because they haven't thought about it hard enough, or because it seems to them to be harder. But it's much easier when the study is all done, and you're going to write it up, it's so much nicer to be able to draw the strong causal inference, than to have to settle even for some fairly decent alternative; they're just not the same thing. Lovely as some of the newer structural equation modeling procedures are, lovely as things like cross-lag panel analyses are, they simply don't allow you the same kind of leverage to draw causal inference. So, if somebody says, "I know how to make organizations produce more," they're not going to convince me of that by a path diagram, no matter how elegantly drawn, or how well estimated the data coefficients are. If somebody could think of a way to do an experiment, and someone almost always can, then that's what I'd want to see, and that's what would persuade me, whether it's organizational research, or organizational productivity, or research on the outcome of psychotherapy, or research trials on psychopharmacological agents. It's the same basic idea.

Peter Blanck: I think the basic idea in your discussion relies on your phrase "to draw causal inference." Actually, I would speculate that maybe half or more of the field research is not intended to draw causal inference. For instance, you're looking at CFOs of CEOs. You're just looking at it trying to describe the way they go about their duties. That's different than saying CEOs who come from a different socioeconomic status are more likely to do better in certain types of organizations

Robert Rosenthal: You're right.

Peter Blanck: But I think that your point is well taken, in that most researchers are not aware of that - they're not aware of the issue of drawing causality.

Robert Rosenthal: I think that's a very important point. It just happens that a lot of the things that I'm interested in, I want to draw causal inferences about, but I think you're quite right that many people don't. One research project we did with Dupont, for example, had to do with measuring sensitivity to non-verbal cues by higher level and lower level executives, and it turned out that the higher level executives were more sensitive to non-verbal cues than the lower-level executives. That's just a relation kind of statement. No causal inference is really appropriate in that case. But I could see where even in that case of a correlation type of result, someone might want to draw a causal inference in the back door. For example, they might say, well, "Let's hurry up and set up one of these management training courses that makes everyone more sensitive to non-verbal cues, since our top management people are more sensitive to non-verbal cues than the ones slightly lower down. Then, if we make people more sensitive, maybe that will make them better. But, of course, that's a *non sequitur* - that doesn't follow unless we do the experimental analysis that allows us to draw the causal inference. We really have no way of knowing whether being more sensitive to non-verbal cues actually helps in being a higher level executive or whether it hurts. All we know is that the higher level executives were slightly better at decoding non-verbal cues, and it may be due to whatever makes them higher level executives, also makes them better non-verbal decoders. But we don't know if the non-verbal skill makes them better executives.

Peter Blanck: Yeah. I think part of that problem of the willingness to jump to conclusions like that lies in part in the type of graduate training that's different in the psychology department as opposed to a business school. Maybe you can describe what you think would be an optimal program, a methodological program, that could teach experimental and field research methods, and that could also be of use to both field researchers at the Business School and experimental researchers in the psychology department.

Robert Rosenthal: Well, I think many of the tasks of people in management, or people in the business school, are very similar to the tasks of psychologists. Their end goal may be somewhat different, in that it may be practical, in the one case, and more theoretical in the other. But, I think they all have in common that they want to make accurate statements. And, I think that's where, as a discipline, psychology has the most to offer. Even, being - maximally critical of what psychology has achieved so far, and I think we've achieved a fair amount, but not an awful lot. The one thing, I think we really can contribute, whether it's in business schools, or law schools, or medical schools, is research methodological training, including, but certainly not restricted to, the quantitative material. There are biostatisticians and there are decision theorists, and very strong quantitative-type scholars already in business school. So, when I talk about methodology, I mean much more than just quantitative analysis of data, or the design of experiments. I mean instrument selection, instrument development, questions of reliability, of validity, pushing across notions of the sort that have been so beautifully developed by Don Campbell about

triangulation using multi-method approaches. Specifically, looking at whether different instruments are getting at sufficiently different kinds of results. Not relying on just one approach, not doing just interviews, not doing just questionnaires, but using all the techniques together in a balanced way, and seeing if they all point to the same kind of conclusion.

Peter Blanck: Are there any techniques that generally fall together, or is it just up to the researcher which instruments or methodologies that they want to use, realizing that certain of them may be appropriate for experiments. Do certain experimental techniques naturally fall together, and actually, are statistically better to use together, and methodologically experimental design-wise are better to use together?

Robert Rosenthal: That's such a hard question. Could you say a little more about it?

Peter Blanck: If somebody was interested in drawing causal inference about whether CEOs who went to more prestigious universities end up doing better relative to other CEOs in their organization, and if the main instrument or the main methodology that the researcher employed was interviewing, would there be a natural sort of covalidation technique available to the researcher? For instance, expert ratings from people outside - the field of the CEL as opposed to - ratings inside the field - as opposed to observations of his daily tasks.

Robert Rosenthal: I guess that for the most part, I would think the more different approaches that you can think of, the better off you are. Some are going to be more expensive than others. Some are going to be more feasible than others. But, in general, there's probably no one instrument or set of instruments that are best. I think that there's a place for many of the different techniques. The question that you ask is an interesting one. That would be a very difficult one to institute appropriate controls for. I mean does it mean that the same kinds of characteristics that get you into a really good business school are the same kinds of characteristics that get you into high level positions in business and industry, or does it mean that, no matter who it would be that would go a very prestigious business school, if you're a graduate of a prestigious business school, that gives you such a leg up that you're going to go further regardless of what your actual prior attributes would be. One could do, through particular statistical procedures, multiple regression procedures, for example, that could help you sidle up to that question. I have my doubts about how firmly one could answer questions of that sort, but one could certainly do better than just guessing, or one could certainly do better than just stating the relationship that CEOs of big companies are more likely to have arisen from more prestigious business schools, for example, or Ivy League Undergraduate Schools.

Peter Blanck: A more political question now in light of Reagonomics and the obvious cuts in research funds, although it's obviously less relevant to business schools and more professional schools, which are less dependent upon government sources. But, as a psychologist and coming from a department and departments that are generally heavily reliant on government for funding, what sort of issues do you think that it would be beneficial in the eyes of the government for psychologists to address? Specifically,

basically, I am saying, do you think that because of all these cuts, the eye-catchers in the field are going to become more apparent research, or do you think people are going to stick to more hardcore lab work?

Robert Rosenthal: There's going to be less and less money, it would appear, at least in the near run, for hard core lab work, whether that hard core lab work is in general experimental psychology, or social psychology, or even personality psychology. There'll probably be more money out in the real world and that should be added impetus for collaboration with law schools, with business schools, with medical schools. I think that over the years seen an increase in that kind of collaboration, and one will probably see more of it, partly brought on by these economic exigencies, but partly also because they make good sense. A lot of the research that I've been interested in doing, I couldn't do without access to psychiatric settings, for example. And, so, for a long time, even before the money crunch, I've been working on and off with various hospitals, and that makes it easier to get certain kinds of research done working with patients. I've worked with alcoholics over the years, for example. If you have, if you have access to professional schools, you have access to real life problems, whether they're business problems or medical problems, or legal problems and a more natural opportunity to do worthwhile field research, but with the high standards that you carry with you when you go over to the business school, or the law school, or the medical school, or the school of public health.

Peter Blanck: Is it true more generally that you or people in general learn research by doing research and, no matter what kind of graduate training you get, if you don't do any research as a graduate student, you might as well start from scratch who you get out in the real world?

Robert Rosenthal: Absolutely. Though I might qualify the very last sentence that you made. I think it's absolutely the case that you learn the most about doing research by doing research, and that you can't really regard yourself as a researcher no matter how exquisitely you've been trained to do it, until you've done it and done it a lot. I think that you do have a head start going into the research setting having had good training, versus not having had good training. So that if you take two people, equally able, to begin with, and turn them both loose into their first research project, the one who's done the appropriate background reading and has had the appropriate courses in research methodology will certainly have a head start over the one that doesn't. But neither one of them is going to be a great researcher until they've done a fair amount of research.

Peter Blanck: Is that how you learned about research?

Robert Rosenthal: Doing it. That's the only way to learn, by doing it.

Peter Blanck: OK. Do we have time for...?

Other: Rolling. Yes.

Peter Blanck: OK, Bob. I wanted to ask you, since you've done a lot of work on experimenter bias, if you could define bias for us, and maybe give us an illustration of how it profoundly can affect an experiment?

Robert Rosenthal: OK. There is actually a fair number of different ways in which experimenters can unintentionally foul up the results of their research, but the one that I've studied the most has been the one that has to do with the experimenter's hypothesis. His or her hypothesis can come to serve as a self-fulfilling prophecy. So, they get the results they expect to get not necessarily because they been so clever in anticipating what nature is going to say, but because they have treated their research subjects in a different way, in accordance with their expectations. For example, if someone is doing an experiment with the Rorschach, if you tell half the Rorschach examiners that their responders are going to see a lot of human movement responses, those examiners will get more human movement responses from their research subjects than will those who've been led to expect more animal movement responses, who, in turn, will get more animal movement responses. There've been just scores of studies using human subjects to demonstrate that this occurs with an alarming frequency. Some of the most interesting and compelling examples that have come, not from work with human subjects, but from work with animal subjects. In one early experiment that I did with Kermit Fode at the University of North Dakota, we had a group of rats that we labeled arbitrarily as maze-bright, and maze-dull. And, we told the experimenters that we had imported these rats from special breeding grounds, where they had been bred for maze brightness or maze dullness. And, what we really did was use a table of random numbers to arbitrarily label half of the rats as maze-bright and half of the rats as maze-dull. We put them in a maze to see how fast they could learn, and it turned out that those who had been arbitrarily labeled as maze-bright, actually learned the maze faster than did those who had been labeled maze-dull, presumably because of the differential handling patterns of the experimenter. Each time the rat had run the maze, the experimenter would have to pick up the rat in his or her hand, and start the rat over for the next trial. We think that the way in which the rat was picked up communicated to the rat how the experimenter felt about the rat. And that the expectation for the rat's brightness or dullness was actually communicated through pressure to the rat. It's not very farfetched to think that the same kind of handling pattern operates in the classroom, where the handling may not be quite so concrete, it may be more symbolic and more abstract, but teachers' handling of pupils has certainly been shown to be a self-fulfilling prophecy in the classroom as well.

Peter Blanck: And moving out of the lab a bit and into more mediating factors, as you call them, of these expectancy effects, I know you've done a lot of work on how non-verbal cues may mediate interpersonal expectancy effects, and expectancy effects in particular. What sort of things are going on in that relationship, and what sort of things does a field researcher have to be aware of? These subtle cues may actually affect their interaction?

Robert Rosenthal: Well, the research shows that many different kinds of cues, non-verbal cues in particular, can serve as the mediators of these unintentional, interpersonal self-fulfilling prophecies. So, the field researcher has to be aware that his or her tone of voice, his or her body movement patterns, his or her facial expressions, can communicate a lot to

the person with whom they are in interaction. And, that this communication can often come to be of the self-fulfilling sort. That if you are doing an interview with people who have done very well in business, or done very poorly in business, and you as the interviewer know that, you may communicate your feeling about them as very successful or as very unsuccessful people, and make them appear in the interview to be even more successful or more unsuccessful, not because that's how they normally would interact with people, but because of your very specific expectation that they're going to be terrific on one hand, and/or awful on the other.

Peter Blanck: Yeah. That seems a tremendous problem for the field researcher to avoid, given that he has to develop his personal access, and, given that, he obviously knows the idea that he's looking for. Aside from this awareness of these subtle types of cues, are there any types of things that field researchers can use? You mentioned before - using different interviewers. Are there any specific things - other things that interviewers can be aware of, or field researchers can be aware of - to prevent bias?

Robert Rosenthal: I think it might be useful to distinguish the place which we wouldn't mind the bias so much, and where we would mind it a lot. And that's in the very preliminary stages where you are first formulating your hypotheses, where you as a behavioral researcher or as a student of management or as a student of CEOs want to generate some hypotheses. I don't think you ought to be worried about bias. I think you ought to go and do very open ended interviews, use questionnaires in addition if you like, but think of that very much as hypothesis generating. Then, when it comes to the hypothesis testing time, after you've formalized the propositions, after you've formalized your hunches or hypotheses, your theories about how things work, that's the point at which, I think, you have to institute these controls against your own biases and expectations, and it may be at that point that you replicate with other chief executive officers, with other investigators, that you may want to employ, with whom you don't come fully clean - that is, you don't tell them everything that's in your mind about what you expect to find, you just tell them the kind of interviewing data that you want to collect, and send them out as professional data collectors - basically, highly trained ones to be sure, but without really access to your particular hypotheses. They don't need to know that you're comparing conversational styles of CEOs who have gone to Ivy League undergraduate schools, and most prestigious business schools, for example, if that's what you're interested in. So, in the hypothesis generating stage, I think you can afford to be very relaxed about these methodological pitfalls. It's only if you want to make claims, it's only if you want to ascribe some generality, it's only when you think you've tested the hypothesis, rather than simply suggested the hypothesis, that I think you really have to be extremely careful.

Peter Blanck: I want to talk briefly now about another aspect about the research process, in general, and that's presentation of your results. We've moved through hypothesis generating ideas, through analysis, and now we're at the stage where you have all this large amount of data, and a lot of it is very qualitative, especially if you're working in the field. How do you find the best way to present what's important to put in, what's important to leave out, what's the level of detail you think that is necessary?

Robert Rosenthal: I think that's - going to depend on the audience that you have in mind, and I can see the very same research project presented in very different ways for different audiences. For example, you might have - a very technical, statistically oriented discussion in a professional journal, like the *Journal of Applied Psychology*, of some business research program that you've undertaken. You might then go into one of the business journals with some of quantitative material covered, but leaving out a lot of the details. If it involves educational issues, like training with an industry, you might want to make that kind of information available to the educational community, but again in a very different way. I think of some research, for example, that could be published in *The Journal of Educational Psychology* which is sort of the prestige journal of the American Psychological Association for educational type research - that wouldn't be read by classroom teachers, or school principals. There are other journals that would be read there. You might want to see that the same information is made available to them, but you'd have to rewrite the article completely. You wouldn't want to have all those chi-squares, and linear contrasts, and all these fancy kinds of things that are fine for a journal of educational psychology, in a journal like *The Reading Teacher*, because the average reading teacher who reads that wants to know what's been found. For example, that boys have reading difficulties more than girls only when they are being taught to read by teachers who believe that boys have more trouble reading than girls, which is a result that actually was reported in the educational literature by Polardi, and has been replicated by others. So, the outlet is extremely important, and I think, far from it being a questionable, dubious practice to publish multiply, to some extent, I think it's almost an obligation to do so because in one way or another, even if you don't have a National Science Foundation Grant or a National Institute of Health Grant or a National Institute of Education Grant for your research, one way or another if you're an academic researcher, you're being supported by the taxpayer, even if you're at a private school. So, I think there's a certain obligation to go beyond reporting the results just to your peers, just to the prestige journals in your field that earn you yourself as an investigator the most brownie points. I think you should make those data available to other people to consumers in the field who can make some practical use of the knowledge, even though it doesn't do you any particular good to publish in an educational journal if you're a psychologist.

Peter Blanck: One final question. What do you find fun about doing research, or field research, in particular, or fun in the field setting, and satisfying as a person, that you're making a contribution to science in general?

Robert Rosenthal: What an interesting question. I'm blessed: there's almost no part of the research process that I don't find fun. Probably the part that I find least fun is the writing. I really don't like to write. I do a lot of it, but I really don't like it. But I love to plan the research. I guess I don't myself spend an awful lot of time collecting the data, though I have done that, even fairly recently. I love to analyze data, and to think about how to extract more information from the same set of numbers. I like to think about how to present the material, and I like to work with colleagues who will do most of the writing. And I enjoy teaching about research, so that many of my collaborators are people who are

younger than I and are sort of getting into the research trade plying business. Basically, though, I think I enjoy it all.

Peter Blanck: OK. I think that's a good ending.