Chapter One

Eight Common Misconceptions about Psychology Papers

Misconception 1. *Writing the psychology paper is the most routine, least creative aspect of the scientific enterprise, requiring much time but little imagination.*

Many students lose interest in their research projects as soon as the time comes to write about them. Their interest is in planning for and making new discoveries, not in communicating their discoveries to others. A widely believed fallacy underlies their attitudes. The fallacy is that the discovery process ends when the communication process begins. Although the major purpose of writing a paper is to communicate your thoughts to others, another important purpose is to help you form and organize your thoughts.

Reporting your findings in writing requires you to commit yourself to those findings and to your interpretation of them, and opens you to criticism (as well as praise) from others. It is perhaps for this reason as much as any other that many students are reluctant to report their research. But the finality of a written report also serves as a powerful incentive to do your best thinking, and to continue thinking as you write your paper. It requires you to tie up loose ends that you might otherwise have left untied. As a result, reporting your findings presents just as much of a challenge as planning the research and analyses that led to those findings.

I have often thought I knew what I wanted to say, only to find that when the time came to say it, I was unable to. The reason for this, I believe, is that in thinking about a topic, we often allow ourselves conceptual gaps that we hardly know exist. When we attempt to com-

municate our thoughts, however, these gaps become obvious. Organizing and then writing down our thoughts enables us to discover what gaps have yet to be filled.

Misconception 2. *The important thing is what you say, not how you say it.*

As a college student, I was mystified to find that students who wrote well consistently received better grades on their compositions than did students who wrote poorly. Even in my own compositions, I found that the grades I received seemed less to reflect what I had to say than how I said it. At the time, I was unable to decide whether this pattern in grading resulted from the professors' warped value systems, or from their inability to penetrate the facade of written prose. Whereas their criteria for grading papers might be appropriate for an English course, these criteria seemed inappropriate for courses in subjects like psychology.

As a college professor, I have at last discovered the secret of the mysterious grading practices. The discovery came about in two stages, each one part of the initiation rites new college teachers must go through. The first stage occurred when I found myself having a large number of students' papers to read and very little time in which to read them. I was then sincerely grateful to students who wrote well because I could read their papers quickly and understand what they were saying. I did not have the time to puzzle through every cryptic remark in the poorly written papers, however, and I resented the authors' presenting their ideas in a way that did not enable me to understand or evaluate them properly. I also found myself with no desire to reward the authors for this state of affairs. If their ideas were good, they should have taken the time to explain them clearly.

The second stage of discovery occurred when I found myself with just a few seminar papers to read, and plenty of time in which to read them. Now, I thought, I can be fair both to students who write well and to those who do not. I was quickly disabused of this notion. I discovered that whereas it is usually easy to distinguish well-presented good ideas from well-presented bad ideas, it is often impossible to distinguish poorly presented good ideas from poorly presented bad ideas. The problem is that the professor's comprehension of what the student says is solely through the student's way of saying it. Professors can't read minds better than anyone else. If an idea is presented in a

sloppy, disorganized fashion, how is one to know whether this fashion of presentation reflects the quality of the idea or merely the quality of its presentation?

The question is not easily answered. In one case, I had talked to a student beforehand about what he was going to say, and I expected an outstanding paper on the basis of these conversations. During our conversations, certain details had not been clarified, but I expected these details to be clarified in the paper. Instead, the same ideas that had been inadequately explained in the conversations were inadequately explained in the paper as well. Either the student was unable to clarify these ideas for himself, or he was unable to clarify them for others. The outcome for the reader is the same: confusion and disappointment.

A comparable situation exists for researchers. One quickly notices that the best and most well-known psychologists are also among the best writers. Although there are exceptions, they are infrequent: Poorer writers have fewer readers. One reason for this fact is that poorly written articles are usually rejected by journal editors. Although journal editors are willing to make minor editorial changes in the articles they receive, they are usually unwilling to publish or rewrite poorly written articles. Even if a poorly written article is accepted and published, however, psychologists who receive a journal with 5 to 20 articles in it do not want to spend their limited time reading such an article. It is therefore important that you learn now how to present your ideas in a readable fashion.

Misconception 3. *Longer papers are better papers, and more papers are better yet.*

Until my first year of teaching, I believed that longer papers were better papers. Teachers had for years told me and my classmates that they didn't evaluate papers on the basis of length, but I viewed their remarks as a benign ruse designed to discourage length for its own sake. I changed my viewpoint when I started reading students' papers. Evaluating papers on both quality and quantity of ideas, I found little relation between either of these two criteria and the length of students' papers. Sometimes students wrote longer papers because they had more to say; other times they wrote longer papers because it took them several pages to say what could have been said in several sen-

tences. There is nothing wrong with length per se so long as length is not used as a substitute for tight organization and clear writing.

Rather than writing longer papers, some people have taken the other route of writing more papers. Why say in one paper what can be said in two for twice the credit? This kind of mentality meets the needs of people who count publications, but not of those who read publications. An integrated series of related experiments will have more impact if published as a single, tightly knit package than if published as a string of hastily written articles, none of them of much interest in itself.

Misconception 4. *The main purpose of a psychology paper is the presentation of facts, whether newly established (as in reports of experiments) or well established (as in literature reviews).*

A common misconception among the general public is that the goal of science is the accumulation of facts. This misconception is fostered by popular scientific writing that emphasizes scientific findings, which may be easy to describe, at the expense of explanations of these findings, which may be both diverse and difficult to describe. Diverse explanations, however, are the hallmark of science.

Students in introductory psychology courses are prone to this misconception, and it carries over into their writing. I could cite numerous examples of this carry-over, but one in particular comes to mind. I received some years ago a beautifully written paper reviewing the literature on the testing of infant intelligence. This was one case, however, in which flowing prose was insufficient to obtain a high grade. The paper was flawed in two respects. First, the author made no effort to interrelate the various attempts to measure infant intelligence. Each attempt was described as though it had been made in isolation, even though the various attempts to measure infant intelligence have drawn upon each other. Second, the evaluative part of the paper consisted of a single sentence in which the author stated that it is still too early to draw final conclusions regarding the relative success of the various infant intelligence tests. This sentence is literally true: It is too early to draw *final* conclusions. But it will be too early to draw final conclusions as long as new data about the tests continue to be collected. Because data will continue to be collected for the foreseeable future, and because the tests date back to the early part of the 20th century,

it now seems appropriate to draw at least tentative conclusions. In writing a psychology paper, you must commit yourself to a point of view, even if you may change your mind later on. If the evidence on an issue is scanty, by all means say so. But draw at least tentative conclusions so that the reader knows how you evaluate what evidence is available.

Your paper should be guided by your ideas and your point of view. Facts are presented in service of ideas: to help elucidate, support, or rewrite these ideas. They provide a test against which the validity of ideas can be measured. You should therefore select the facts that help clarify or test your point of view and omit facts that are irrelevant. In being selective, however, you must not select only those facts that support your position. Scientists demand that scientific reporting be scrupulously honest. Without such honesty, scientific communication would collapse. Cite the relevant facts, therefore, regardless of whose point of view they support.

Misconception 5. *The distinction between scientific writing, on the one hand, and advertising or propaganda, on the other, is that the purpose of scientific writing is to inform whereas the purpose of advertising or propaganda is to persuade.*

Successful advertising or propaganda need only persuade. Successful scientific writing must both inform and persuade. Students often believe that a successful piece of scientific writing need only inform the reader of the scientists' data and their interpretation of the data. The reader is then left to decide whether the theory provides a plausible account of these (and possibly other) data. This conception of scientific writing is incorrect.

When a scientist writes a paper, he or she has a product to sell. The product is his set of ideas about why certain phenomena exist. Occasionally, it is the only product on the market, and he need only convince the consumer to buy any product at all. Whether or not the scientist is successful will depend in part upon how persuasive he is, and in part upon how much the product is needed. No advertising campaign is likely to sell flowers that are guaranteed not to germinate, nor an explanation of why people don't normally stand on their heads rather than their feet. In most cases, however, there is an already established demand for the product. Because competing salespersons are

trying to corner the market, the scientist must persuade the consumer not just to buy any product, but to buy his product.

One of the most common mistakes students make is to sell the wrong product: They misjudge the contribution of their work. I recently received a paper that was full of good, original ideas. The presentation of these ideas, and of other people's as well, was unusually lucid. The only major problem with the paper was that the discussion of the original ideas was condensed into one paragraph buried inconspicuously in the middle of the paper, whereas the discussion of the other people's ideas spanned about 10 pages, starting on page 1. The contribution of this paper should have been in its new perspective on an old problem. But the author had deemphasized this potentially significant contribution in favor of a relatively unimportant one: providing a well-written but unexciting review of other people's perspectives. The hurried reader will usually take the author's emphasis at face value. In this case, the reader might conclude that the paper did not have much of an original contribution to make.

At the opposite extreme, it is possible to dwell so heavily on the contribution of your paper that the contribution is actually muted. I learned this lesson the hard way. A colleague and I wrote a paper intended (a) to compare different measures of a psychological construct called subjective organization, and (b) to demonstrate that one of these measures is superior to all the rest (Sternberg & Tulving, 1977). We compared the measures on a number of different criteria. One measure proved to be superior to the rest on every one of these criteria. Despite my colleague's warnings, I explicitly called attention to this fact several times in the paper. Leaving nothing to chance, I pointed out the inescapable conclusion that one measure is better than all the rest, and therefore should be the measure of choice.

We submitted the paper for publication, and several months later received two scathing reviews. We were attacked for making what both reviewers believed to be exorbitant claims. According to the reviewers, we had by no means developed an open-and-shut case in favor of the measure we claimed was best. I thought that the arguments made by the reviewers were weak and in some cases plainly incorrect. I was so annoyed with the whole affair that I let the paper sit on my shelf for about a year. Rereading the paper and the reviews a year later, I still

believed the reviewers were on the wrong track. My colleague and I decided to tone down our claims for our preferred measure, however, while retaining the same basic line of argument. We resubmitted the paper, and this time received a very favorable review. We achieved much more effective results by understating our case than we had by overstating it, an outcome my colleague (but not I) had anticipated from the start. My subsequent experiences have confirmed repeatedly that in psychology papers, a soft selling technique is more successful than a hard selling technique. By using the latter, you invite a reaction against you as salesperson that is likely to hurt the sale of your product. I can recall numerous occasions on which I refused to buy a product because I detested a pushy salesperson. In writing the first draft of the paper on measures of subjective organization, I unwittingly occupied the role of the pushy salesperson, and I received what should have been a predictable response.

Misconception 6. *A good way to gain acceptance of your theory is by refuting someone else's theory.*

A surprisingly common ploy in scientific papers, even some published in prestigious journals, is to resort to explanation by default. Whereas students may not know better, professionals should. The investigator describes two (or more) theories of the well-known XYZ phenomenon. She then presents devastating evidence against all theories except one. She concludes on the basis of this evidence that this one theory is correct.

This indirect method of proof is compelling only when the two (or more) alternatives are (a) mutually exclusive and (b) exhaustive. Mutually exclusive alternatives are ones in which one outcome precludes the other(s). If a coin lands heads, for example, it cannot at the same time land tails. Exhaustive alternatives are ones that include all possible outcomes. A flip of a coin can result in heads or tails, but nothing else.

The ploy described above has been used in some (but by no means all) research studying sources of differences between groups in intelligence test scores. A study would be presented in which obtained differences in test scores could not be attributable to environmental factors. The author would conclude on this basis that the differences must be due to hereditary factors. These alternatives, however, are neither mutually exclusive nor exhaustive. First, it is possible – indeed,

probable – that both heredity and environment influence intelligence test scores. Second, a further source of influence upon intelligence test scores is the interaction between heredity and environment – the effect produced by their joint influence. As an example, certain genes for intelligence may manifest themselves only under favorable environmental conditions.

One other disadvantage of the indirect method of proof bears mention. Criticism of other people's theories often gains one more opponents than it does converts to one's own theory. This was another lesson I learned the hard way. I once wrote a paper that had two major goals: (a) to show that my theory of a phenomenon was correct; (b) to show that someone else's theory of the phenomenon was incorrect. I presented what I believed was strong evidence in favor of my theory and in opposition to the other person's theory. I submitted the paper to a journal, and it was rejected. The main reviewer of the paper, predictably enough, was the other theorist. It is a common practice to send papers attacking Theory X to Theorist X, with the editor then using her judgment as to whether the review is a fair one. The reviewer criticized not the positive aspect of my paper, but its negative aspect. He argued that our theories actually dealt with somewhat different aspects of the phenomenon under investigation, so that there was no need to attack his theory in the process of supporting my own.

In retrospect, I think the reviewer probably had some valid points; I also think he overreacted. In papers I've reviewed that attack my work, I've probably overreacted as well. Scientists have a reputation among the general public for being objective seekers and impartial evaluators of the truth. I think this reputation is generally deserved, but only when it comes to each other's work. When it comes to their own work, scientists lose their objectivity. When a scientist is attacked, he or she behaves in much the same manner as anyone else under attack. When someone lunges at you with a fist flying toward your face, you don't stop to reflect upon the various considerations that may have led your opponent to attack you. You counterattack. Because scientists are personally so involved in their work, they often treat an attack on their work as a personal attack, even if there is no rational basis for treating it as such. The result can be a personal confrontation in which scientific issues are placed on the back burner.

In conclusion, it is wise to stress the positive contribution of your

paper. This does not mean that you should forgo criticizing other theories. Such criticism may be essential to your point. If it is, keep in mind my earlier admonition that understatement is a more effective means of persuasion than is overstatement. Avoid statements that can be interpreted as contentious but lacking in substance. And if you publish your paper, don't expect the investigator you criticize to congratulate you on your cogent refutation of her work.

Misconception 7. *Negative results that fail to support the researcher's hypothesis are every bit as valuable as positive results that do support the researcher's hypothesis.*

Because science is a fair game, the scientist wins some and loses some. Students often believe that the only honest course of action is for the scientist to report his losses as well as his wins. To do otherwise would seem to present a false picture of both the scientist and the state of nature.

After reading a diverse sampling of journal articles, the student is bound to arrive at one of two conclusions – either scientists have uncannily sound intuitions about the way experiments will turn out or they maintain closets full of unsuccessful and unreported experiments. Although scientists usually have at least fairly sound intuitions about how experiments will turn out, the state of the journals is more a reflection of well-stocked closets than of unerring intuitions.

Scientists' failures to report failures are attributable not to their dishonesty, but to the frequent uninterpretability of negative results. Suppose, for example, that an investigator predicts that giving children rewards after learning will increase their learning. The investigator conducts an experiment with two groups. In one group, children receive rewards after learning; in the other group, they do not receive rewards. The investigator finds no difference in learning between groups. What can she conclude? Unfortunately, not much. Whereas a significant difference between groups would have provided good evidence that rewards can facilitate learning, absence of a significant difference could be explained in a number of ways, most of them uninteresting. Consider three such uninteresting explanations:

1. The reward used in the study did not prove a powerful enough incentive. If the reward, for example, was a peanut, then children's cravings

for a single peanut might not have been strong enough to increase their efforts to learn.

2. The sample of children might not have been large enough. It is a well-known rule of statistics that if any treatment effect exists at all, then it can be discovered if one's sample is large enough. A small effect may be detectable only with a relatively large sample. If there were only three children in each group, then the investigator might have failed to detect the effect of the reward.
3. The measure of learning might have been inadequate. Suppose, for example, that the task was to learn the set of multiplication facts for one-digit numbers, and that the measure of learning was a single multiplication fact. This measure probably would have been inadequate to detect learning in either group, and hence a difference in learning between groups.

Under two sets of circumstances, negative results can be of interest:

1. An investigator repeatedly fails to replicate someone else's results. Suppose someone reports that subjects who stand on their heads for 30 seconds prior to taking a test of visual-motor coordination perform better on the test than do control subjects who do not stand on their heads. Another investigator, suspicious of this result, tries to replicate it with two groups of subjects, and fails. Realizing that his failure to replicate the result may be due to sampling fluctuations, the investigator tests two more groups of subjects, and again finds no significant difference between groups. At this point, he feels ready to report the result. Whereas one failure to replicate a result is not informative, repeated failures to replicate can be informative. The number of failures needed depends in large part upon the strength of prior evidence in support of the result in question. Two failures are probably more than adequate for the "headstand hypothesis," whereas a great many failures would be needed to overthrow a more well-established result, such as that under normal circumstances learning increases with practice.
2. A significant result vanishes when a methodological weakness is corrected. Suppose that the experimenter who wrote the "headstand" paper knew which subjects had stood on their heads, and which had not. This aspect of the methodology suggests a possible bias in the experimenter's scoring of the coordination test (especially if the experimenter is Public Relations Director of the American Association for

the Advancement of Acrobatics). A worthwhile methodological refinement would be to conduct the experiment under circumstances in which the experimenter does not know which subjects stood on their heads and which did not. A negative result would be of interest in this case, because it would suggest that the significant difference between groups in the first experiment was due to experimenter bias.

Misconception 8. *The logical development of ideas in a psychology paper reflects the historical development of ideas in the psychologist's head.*

If one were to take journal articles at face value, one would conclude that scientific results come in neat, attractively wrapped packages. One need only go through a uniform series of well-defined steps in order to ensure delivery of such packages:

1. The scientist starts with some clever ideas about a phenomenon, which she explains in the introduction to the paper. These ideas are carefully formulated before the scientist has collected any data, and the data merely serve to confirm (or in rare cases disconfirm) their validity.
2. The scientist tests these ideas by carefully choosing variables that can be manipulated in a controlled experiment. The scientist's deep understanding of the phenomenon under observation and of scientific method enables her to choose the correct variables and experimental manipulation on her very first attempt, which she describes in the *Method* section of the paper.
3. The scientist performs the experiment, presenting in the *Results* section of her paper the outcomes of data analyses scrupulously planned in advance.
4. The scientist finally reflects upon the broader implications of the results, presenting her reflections in the *Discussion* section of the paper.

I doubt that 1% of the papers published in scientific journals developed in a way even remotely resembling the outline sketched above. Yet the large majority of published papers are written as though they had developed in this way, or in some way closely resembling it. Let us reconsider the series of steps:

1. Before carrying out an experiment, one usually has only a vague and tentative idea of what the outcome will be, if only because there are so many possible outcomes that one can scarcely even enumerate them all. One's ideas develop along with the experiment.

2. One sometimes finds oneself performing the right experimental manipulation on the wrong variables, or the wrong experimental manipulation on the right variables. In order to avoid wasting large amounts of time and money, scientists frequently conduct small-scale pilot experiments that test the feasibility of the experiment as designed. Adjustments in method can then be made in preparation for the full-scale experiment, or the experiment can be scrapped altogether.
3. Major data analyses are usually planned in advance. Indeed, it is necessary to do this planning in order to assure that the design of the experiment permits one to analyze the data in the most advantageous way. Minor data analyses are frequently decided upon after the data have been collected. Often the results of a planned data analysis will suggest a subsequent unplanned one. Only fools fail to go where the data lead them. One of the most valuable skills scientists can have is a knack for getting the most out of their data. A given set of data can be analyzed in an infinite number of ways, some of them more revealing than others. The scientist must select a small number of ways that are likely to yield maximum payoff.
4. Ideas for the *Discussion* section of a paper usually start forming at the same time the experiment does, not merely after it has been completed. The reason for this fact is simple. Unless the experiment has at least some potentially broad and interesting implications, or unless it can lead to some sensible next step in research, it is probably not worth doing.

Why does the picture of research presented by journal articles correspond so poorly to the actual state of affairs? There are at least three reasons:

1. Journals operate under severe space limitations. A large percentage of articles submitted to the journals must be rejected for lack of space. In some journals, more than 90% of submitted articles are rejected. Those articles that are accepted must be as concise as possible. An "autobiographical" form of presentation, describing all one's false starts and initial misjudgments, consumes a great deal of space. This space is more profitably devoted to other articles.
2. An autobiographical account of an experiment tends to be of more interest to oneself than to one's colleagues. An associate recounted to me the way in which he learned this lesson. He submitted a 20-page theoretical article to one of the most prestigious psychological journals. He spent the first 19 pages of the article describing how he had come to his

conclusions after a lengthy series of false starts; he presented his final conclusions on the 20th page. The article was rejected, not because the final conclusions were wrong, but because the editor believed that there was only one publishable page in the article – the last. The editor was interested in the psychologist's conclusions, but not in the lengthy soul searching the psychologist had done to arrive at them.

3. The object of description in a scientific report is a phenomenon and its explanation, not the reporter of the phenomenon and explanation. The focus of the report must reflect this fact. A graduate student and I once completed an experiment investigating the development of reasoning skills in children at the second-, fourth-, and sixth-grade levels. Children were presented with reasoning problems, which they were then asked to solve. Because the experiment involved a considerable investment in time and money, we decided to pretest our reasoning problems on some colleagues' children. Our original plan had been to use number of problems correctly solved as the dependent measure. We discovered, however, that even the youngest children made almost no errors on the problems once they fully understood the task. We therefore changed our dependent measure when we did the full-scale experiment, using response time to solve problems correctly instead of numbers of problems correctly solved. Had someone else planned this experiment, she might have realized immediately that the problems were too easy to use number correct as the dependent measure; or she might have stumbled longer than we did until the discovery that the problems were too easy. A description of this trial-and-error process is slightly informative about the development of the investigator's intuitions, but it is uninformative about the object of the investigation, in our case, the development of reasoning in children. The scientifically informative statement is that the problems were of a level of difficulty that made response time an appropriate dependent measure.

There is often a fine line between the omission of autobiographical information and the omission of critical details. If a hypothesis is post hoc, then one is obliged to indicate this fact.

In sum, the steps one follows in planning and carrying out research do not neatly correspond to the successive sections of the psychology paper. In the next two chapters, we will consider the steps in carrying out library and experimental research and how to describe them in the psychology paper.