1 Reasoning, Abstraction, and the Prejudices of 20th-Century Psychology

R. E. Nisbett
University of Michigan

Twentieth-century psychology has had a strong prejudice against abstraction, that is, against the view that the world is understood by means of rules that transcend the perception of a particular physical stimulus or the comprehension of a domain of related events. In the United States, the prejudice has been bound up with behaviorism and its successor positions. Behaviorists were determined to find the equivalent of the reflex arc in physiology—stimulus-response linkages that could be described with precision by a physical description of the stimulus, the response, and the conditions of their co-occurrence during learning. So complete was their dedication to such physical description that they felt confident that the study of animals could substitute for the study of humans in building a complete theory of behavior.

Early in the 19th-century, the behaviorist E. L. Thorndike performed a series of experiments that satisfied two generations of American psychologists that abstractions were not importantly involved in learning how to perform skilled tasks. He asked his subjects to perform a particular task for varying amounts of time (e.g., cancelling Os from a sentence, and then switched them to another task; cancelling adverbs from a sentence). He found that "transfer of training" effects were slight and unstable. Sometimes he found that performance of the first task enhanced the second, sometimes that it made it more difficult, and, often, that it had no effect at all. One would, of course, assume that performance on the second task would be improved if subjects learned something general from performance of the first task. Since they so often failed to show improved training, Thorndike inferred that people don't, in fact, learn much that is general when performing mental tasks. This meant that training was going to be very much a bottom-up affair, consisting of little more than slogging through countless stimulus-response associations.

This conclusion has suffused deeply into American psychology, cognitive science, and education. Newell (1980), based on some similar failed efforts to find training effects for reasoning tasks, has asserted that learned problem-solving skills generally are idiosyncratic to the task. Just as the earlier behaviorists took the evidence of weak transfer-of-training effects to buttress their case for the exclusive role of specific stimulus-response linkages, some modern cognitive scientists have used such evidence to support connectionism—the modern successor to behaviorism. To the connectionist, all learning is just a matter of adding strength values to an initially neutral "network" of highly specific elements. The connectionists have all the courage of the behaviorists' convictions—asserting that they can mimic the important details of learning and cognitive performance without the postulation of any rules whatever.

Other trends of modern psychology are opposed to abstract rules, though not necessarily to rules of all kinds. Devotees of case-based reasoning approaches to problem solving hold that people do little more when solving problems than perceive similarities between old and new problems and occasionally apply strategies of analogy construction. Biological and evolutionary theories of cognition are sympathetic to the notion that people operate using rules of limited generality, but these are usually assumed to be limited to relatively tight domains. Thus there are rules, even prewired rules, for language, or for physical causality, or for particular types of social relations, but these are limited to particular content domains and would never be used for understanding events outside those domains.

European psychology has never been so deeply antiabstractionist as American psychology. In fact, Jean Piaget, the European psychologist whose influence on world psychology has been greatest, explicitly endorsed the notion that there are abstract rules that guide thought and behavior. Piaget even thought that the very most abstract rules, those of formal logic, have their intuitive counterparts in the human cognitive repertoire. These rules, part of the equipment that Piaget called *propositional operations,* are used to acquire other, somewhat less abstract but still domain-independent rules in the course of development. These are the *formal operations,* which include the concept of proportionality, the notion that every action has a reaction, and what Piaget called the *probability schema,* but most people today would call the law of large numbers. Piaget believed that people possessed these highly abstract rules in a form in which they made contact with the most ordinary problems in everyday life. Indeed, many common problems could not be solved without the use of such rules, and it was the press of such requirements that pushed people toward their acquisition, by

means of inductive procedures operating within the logical constraints of the propositional operations.

Despite his endorsement of abstract rules for reasoning, Piaget was quite firm in his opinion that such rules could not be explicitly taught—certainly not by abstract or formal means. Such rules are the common equipment of every adult, and everyone acquires them by virtue of being the kind of organism that each human is and by virtue of living in the kind of world that each human does. Given our native equipment and the kind of experiences we are going to have, we perforce learn the abstract rules we are going to need for purposes of deductive and inductive inference. But such rules are learned only by an inductive process of self-discovery; the day of their acquisition cannot be hastened either by abstract instruction in the rule system or by a forced inductive march through specific problems. Nothing, not even an abstract rule, is learned by abstract or top-down training procedures. Nor is there any point in trying to fool Mother Nature by excessive drilling on concrete problems. No rule will be learned before the organism is ready for it and learning is inevitable once the organism is ready—so long as it is not kept in a closet.

You will read Piaget in vain for any evidence for this extraordinarily influential theory of how rules for reasoning are learned. It was simply obvious to him that you cannot teach abstract rules of reasoning, and it became equally obvious to us largely because of his enormous prestige and persuasiveness. There are to be sure shreds of evidence available since Piaget's time that are consistent with his view: (a) solutions to the missionaries and cannibals problem do not generalize to formally identical problems; (b) accelerated learning of conservation of mass for clay does not seem to generalize to an understanding of conservation in general. And there remains Thorndike's work, showing that there can be little transfer of training even across tasks that require lower levels of cognitive skill than one would want to dignify with the term *reasoning*.

But what if Piaget and, even more so, the American psychological tradition, were mistaken about abstractions? What if you actually could teach people highly abstract rules of reasoning—and even do so by highly abstract and therefore efficient means? What if such instruction resulted in people being able to apply those rule systems potentially to the full range of problems in everyday life for which they are relevant? How would we think about the human mind then? How would we think about education?

This volume tries to answer those questions, coming up with some very surprising answers. Ten years ago, I held a version of the received views about reasoning. I was dubious that people had any abstract rules for reasoning and confident that even if they did, such rules could not be taught. Indeed, I had just completed 10 years of work that seemed to me to give substantial support to these views. I had worked on questions of reasoning about human social behavior, finding that people often violated the requirements of statistical, causal, and even logical rules of inference. This work was very much in the tradition of Kahneman and Tversky's research showing that people substitute simple judgmental heuristics for the more formal inferential rules that are necessary to solve the problems they gave their subjects. I believed not only that my subjects did not possess the necessary statistical rules, I believed that instruction in statistics resulted only in inserting a sterile set of formal rules that could make contact only with scientific problems or problems for which there existed some massive and probably ecologically uncommon cue triggering their use.

With Geoffrey Fong and David Krantz, I began what I thought would be a swift program of work establishing these points, namely that instruction in statistics does very little toward helping people to solve everyday problems that require a statistical solution. My very first attempt to look at this question showed me that I was wrong (or should have showed me—actually I didn't believe the implications at first). Kahneman and Tversky (1972) had developed a clever problem to show subjects' statistical weaknesses called the *maternity ward problem*. In this problem subjects are told that there is a town with two hospitals, one large and one small. At the large hospital, about 60 babies a day are born, and at the small hospital about 15. Subjects are then asked at which hospital they think there would be more days during the year in which 60% or more of the babies born would be boys. About one third of the undergraduates they studied believed it would be the larger hospital, about one third believed it would be the smaller, and about one third believed it would make no difference. The law of large numbers, of course, requires that it would be the smaller hospital, because deviant sample proportions are likely in inverse proportion to sample size. While teaching an upper level undergraduate class at the University of Michigan, I tried to duplicate these results in a classroom demonstration. To my surprise, most of the students got the problem right. I then asked students to indicate how much statistics they had had as well as their preferred answer. The results were clear-cut. The students who had had no statistics duplicated the pattern of the Kahneman and Tversky subjects, those who had had at least one course in statistics were unlikely to get the problem wrong.

Subsequent work showed there was no anomaly here. Problems that Kahneman and Tversky had looked at, as well as problems with more social content of the kind I had looked at, turned out to be highly influenced by statistical training. This was true even for problems without obvious statistical clues. For example, one problem we gave subjects asked why someone who had an excellent meal in a restaurant might be likely to complain about a less good meal the next time around. Untrained undergraduates almost always gave purely deterministic answers such as "maybe

restaurants change their chefs a lot." But subjects who had had many courses in statistics usually gave statistical answers, such as "there are probably more restaurants where you can get an excellent meal some of the time than there are restaurants where you can get an excellent meal all the time, so a person who gets an excellent meal the first time has to assume it's likely that the next one won't be."

We then began seeing whether you could teach such statistical reasoning in short training programs. We found to our surprise that even very brief interventions could produce fairly pronounced effects on the sorts of answers subjects would give to problems of the Kahneman and Tversky type. Indeed, purely abstract training, in which we defined the terms *sample, population, parameter,* and *variance,* and explained the relations among *variability, N,* and sample–parameter accuracy, even had an effect on solution of problems with purely social content. Moreover, training in a given domain, for example, training on problems concerning sports transferred fully to another domain, for example, problems concerning ability tests. In several studies, we found literally no advantage for the trained domain over other domains so long as testing was immediate. Such results are consistent with the view that people can operate with very abstract rules indeed, and that the techniques by which they learn them can be very abstract. Abstract improvements to the preexisting intuitive rule system are passed along to the full range of content domains where the rules are applicable, and improvements in a given domain are sufficiently abstracted so they can be applied immediately to a very different content domain.

It is important to note that the sort of instructional effects we discovered in these studies are by no means limited to laboratory or academic settings. When subjects are contacted outside of such settings (e.g., in the context of an opinion poll), the trained subjects answer questions differently from the untrained. In one study, male college students who had either just begun or just finished their first statistics course were asked to participate in a poll on opinions about sports. After answering a number of questions about the National Collegiate Athletic Association rules and National Basketball Association salaries, they were reminded that the top batters in both baseball leagues typically have averages of .450 or higher at the end of the first two weeks of play, yet no one has ever finished the season with such a high average. The students were asked to explain why they thought this was the case. The students just beginning statistics nearly always responded with purely deterministic answers such as "the pitchers make the necessary adjustments." The students who had taken the course were twice as likely as novices to give a statistical answer, such as "two weeks isn't a very long time, so you get some atypically high (and low) averages; no one really has the ability to hit .450 over the long haul."

Over the next 10 years, I pursued the implications of these findings on trainability with different sets of colleagues who were expert in particular rule systems, including the self-selection concept critical to control procedures in the social sciences, "pragmatic reasoning schemas" for contractual relations such as permission and obligation, rules for assessing causality, and the cost-benefit rules of microeconomic theory. The generalizations below hold for these rule systems taken as a group. All of the generalizations have been tested on at least three different abstract rule systems; none are contradicted by any evidence I am aware of.

1. People have intuitive versions of these formal rule systems that they apply to at least some problems in everyday life. We know this because they solve problems that require use of the rule systems, because they articulate the rule systems in justifying their solutions, and because instruction in the rule systems increases the correct solution of the problems.
2. People at a given level of education, prior to formal instruction in a particular rule system, differ in the degree to which they understand the rule system and are able to apply it to solve concrete problems. Such individual differences are associated with verbal intelligence.
3. Formal education beyond secondary school produces dramatic differences in people's use of different rule systems. It is no exaggeration to say that people who have substantial knowledge of statistics, or of economics, view the world very differently from those who do not. All sorts of mundane problems are understood differently by people with differing levels of education in the relevant rule system.
4. The rule systems are embodied at a level of abstraction equal to that posited by Piaget for the so-called formal operations. The absence of domain specificity is a striking observation across training studies, as is the ability of investigators to "insert" the rules by purely formal and abstract instructional means.
5. Despite their abstract nature, the rules are not applied across all domains equally. The same student who has no trouble applying statistical rules to the behavior of random generating devices, such as dice, may apply statistical rules rarely or never to problems with social content. Problems differ a great deal in how transparent they are with respect to a rule system necessary for their solution.
6. A consequence of people's differential ability to apply rules in different domains is that training in coding a given domain in terms of the rule can have dramatic effects – making it possible for people to apply a rule they already have to a new domain where previously it was unlikely for them to use the rule.

The upshot of these findings is that modern cognitive science and modern educational theory must accommodate themselves to the existence of abstract inferential rules. Psychological theories that hold that there are no rules, or no domain-independent rules, for problem solving, are not tenable in the light of the work presented in this book. Educational positions that emphasize self-discovery and maturation must make room for the generalization that abstract techniques of instruction can be very powerful. Psychological and educational positions, as well as philosophical positions, that assume a universal adult competence with respect to reasoning must give way to the recognition that adult inferential competence is highly variable and highly dependent on educational history.

The rest of the volume presents work making these points in detail. Some of the chapters have been published before and some were written especially for this volume. (Reference styles are not consistent across the different papers. We chose to leave the previously-published material in the same form in which it appeared originally, minus their abstracts.)

Part I documents the existence of abstract, intuitive, and statistical rules. Chapter 2, by Nisbett, Krantz, Jepson and Kunda, shows that people without formal training in statistics solve problems using the law of large numbers and actually articulate the rule in justification of their answers. This chapter also shows that the presence of various cues about the partially random nature of the events in a given problem can dramatically affect the likelihood that people will apply the law of large numbers to the problem.

Chapter 3, by Thagard and Nisbett, proposes a solution to Hume's riddle of induction, namely "Why is a single instance, in some cases, sufficient for a complete induction, whereas in others myriads of concurring instances, without a single exception known or presumed, go such a very little way towards establishing a universal proposition?" The solution lies in the law of large numbers, coupled with real world knowledge about the variability of kinds of objects with respect to kinds of properties. A "single instance is sufficient for a complete induction" when we take it for granted that objects of the kind in question are invariant with respect to properties of the kind observed. For example, observing the color of a sample of a new chemical element leaves us in little doubt about the color of future samples. Myriads of concurring instances do not convince when we take it for granted that the kind of object is highly variable with respect to the kind of property observed. For example, observing a bird in the rain forest that is green does not convince us that the next bird we see of the same type will be green because we do not assume invariability of color for bird types.

Chapter 4 directly attacks the notion assumed by many philosophers and psychologists that there is a single human inferential competence. Some philosophers, notably Jonathan Cohen, have argued that empirical demonstrations of human inferential error are logically impossible since "Ordinary