

constructions we use in science and everyday life, but they apply widely enough, I claim, to be surprisingly useful in psychological theory, in the interpretation of psychological experiments, and in guiding and evaluating causal inferences in psychology. More than useful, essential. This book consists of illustrations of this thesis in studies of adult judgement, in developmental psychology, in cognitive neuropsychology, and in psychometrics and social psychology. The illustrations are both positive and negative. I suggest that theories—the *theory theory*—of how cognition develops in infants and children can be fruitfully elaborated by supposing that one of the main tasks of children is to learn the causal structure of the world, and that what is to be learned, and how it could possibly be learned, is illuminated by causal Bayes nets. I show, or at least claim to show, that the results of well-known experiments in adult judgement have been seriously misinterpreted, and that well-known theories of adult causal judgement entail a range of unrecognized, and so untested (but testable), predictions. These suggestions lead to proposals for psychological experiments that have not been done and interesting projects in the development of heuristic learning procedures, including the formation of categories. They include a sketch of an approach to the frame problem, which I understand as the task of specifying feasible algorithmic procedures by which factors relevant to actions and plans can be isolated from irrelevant factors.

In cognitive neuropsychology, so-called “box and arrow diagrams” are causal hypotheses, in fact graphical causal models with extremal (that is, 0 or 1) probabilities. Using the graphical representation and elementary computational learning theory, I investigate two methodological questions that have been disputed at length in the neuropsychological literature: What inferences can be reliably made from the study of the deficits of brain damaged individuals? What inferences can be reliably made if such data are aggregated in various ways to yield “group data”? Bayes nets are a species of neural nets, which are in turn a species of graphical causal models. Reintroducing probabilities, I use graphical causal models to begin to address arguments in cognitive neuropsychology based on the behavior of lesioned neural networks representing brain damage. Can any mathematically possible combination of normal and brain-damaged behaviors be explained by postulating some

neural network representing the normal brain (or normal functional architecture) and lesioning it appropriately?

Many of the statistical models used in psychometrics and in social psychology are graphical causal models in disguise. That recognition, combined with what we know about inference to graphical causal models from data, ought, I argue, to change radically the technology that social psychologists use to represent causal hypotheses, and the methods by which they argue for them. There is a dismaying unity between social statistics and psychological interpretations of adult judgements of causation. The fallacies of statistical methods popular in social statistics—regression and factor analysis—become, in the experiments of some cognitive psychologists, the norms against which the judgements of experimental subjects are assessed. Subjects make judgements that are sometimes, perhaps often, normatively correct, and on the basis of those judgements, psychologists, or some of them, claim to have discovered features of human irrationality so fundamental that, were these psychologists correct, our capacity to get around in the world would become quite mysterious. A central fallacy in regression is the very same error that occurs in many psychological interpretations of experiments on adult human judgement, the topic of much of part 2 of this book. All that changes from cognitive psychology to social psychology is the mathematical clothing of the fallacy. Psychologists are the victims in this practice, the victims of confusions about causal inference buried in the statistical methods they have borrowed from other disciplines. I hope to persuade them that there are better lenders. The points are illustrated with a discussion of the use of data and causal hypotheses in *The Bell Curve*, surely the most notorious work of social science in the last decade, but they equally apply to many less famous efforts in social psychology.

Two features of this book require brief apology. First, while many experiments are proposed, and many are analyzed, this book contains no details of any original experiments. My scientific training is in chemistry and chemical physics, not psychology, and I was never very good even in the chemistry laboratory. I have the greatest admiration for the ingenuity and wealth of tacit knowledge that enables good psychological experimenters to get clean results. Second, psychological papers on

causal judgement often begin with reviews of philosophical metaphysics about causation, generally citing Hume and Kant. There is none of that here, and in the usual sense of the moderns, but not of the ancients, there is no philosophy. I do not propose a philosophical *analysis* of causation, and except as they present experimental ambiguities, I am not much concerned with how people express themselves when making causal judgements or when offering causal explanations. My concern in developmental psychology, for example, is less with how children come to generate explicit causal explanations and more with how they come to be able to predict and control their environment. Even so, a number of distinctions about kinds of causal knowledge necessarily emerge in the discussion.

So far as possible, I have suppressed formalism and mathematical details. Undoubtedly, many readers will think not nearly enough. While I have provided a brief summary of causal Bayes nets and their properties in chapter 3, for the most part I have tried to motivate the essential ideas through examples from the psychological literature. Especially in parts III and IV, substantive discussion is interrupted by historical digressions. I believe that historical considerations are an essential tool for understanding the fundamental problems that motivated the introduction of technical methods and for seeing how contemporary techniques have either solved or evaded those problems. I also think that historical perspective sometimes diminishes the sense of contemporary originality, and rightly so.

I

---

Developmental Psychology and Discovery