

Intellectual Autobiography, 1922-1978 (Written in 1978)

I have divided this autobiography into three main parts: education, research, personal reflections. The second part on research is the longest and most substantial.

I. Education

I was born on March 17, 1922, in Tulsa, Oklahoma, and grew up as an only child; my half brother George was born in 1943 after I had entered the army. My grandfather, C. E. Suppes, moved to Oklahoma from Ohio in the early part of this century. He and my father were independent oil men, intelligent and competent in their business but not well educated. My mother died when I was four and a half; I was raised by my stepmother, who married my father before I was six. She also had not had much formal education, but her interest in self-improvement was strong, and she encouraged me in a variety of ways to pursue my intellectual interests, in spite of my father's ambition for me to follow him and his father in the oil business.

My interest in philosophy was generated early by my stepmother's devotion for more than a decade to the Christian Science of Mary Baker Eddy. From about the age of eight to fourteen years I attended Sunday school and church regularly and studied the works of Eddy as well as the Bible. The naive epistemological idealism of Eddy's writings stirred my interest, which turned to skepticism by the age of thirteen or so. I can remember rather intense discussions with fellow Sunday-school students about how we were supposed to reconcile, for example, the bacterial theory of disease with the purely mentalistic views of Eddy. No doubt our arguments were not at all sophisticated, but our instinct to distrust the flagrant conflicts with common sense and elementary science was sound.

I attended the public schools in Tulsa and was graduated from Tulsa Central High School in 1939. My public school education was more influential on my development than is often the case, mainly because I was a participant in what is known as the Tyler eight-year study of the Progressive Education Association. On the basis of examinations given in the sixth grade, able students were selected to participate in a six-year experiment of accelerated education. In many respects the most competitive and ablest classes I ever attended were those in high school. One of the important aspects of this special educational venture was the extended attempt to get us as young students to talk about a wide range of current events and everything else that interested us. As is often the case, this led into some unusual lines of effort. I can remember very well being chagrined at fourteen if I were not able to name the senators from every state in the union.

The high school courses in mathematics, chemistry and history were excellent, but physics and English were relatively mediocre. The English course was so dominated by the idiosyncrasies of the teacher we had for two years that we became experts in her life and tribulations rather than in the more essential matters of grammar and composition.

I began college at the University of Oklahoma in 1939 but, after the first year, found the intellectual life too pedestrian compared to the much more exciting high school years. In the second year I transferred to the University of Chicago, but, under the laissez-faire approach of Chicago, neglected my academic work so completely that my family insisted on my attending the University of Tulsa for my third year, where I majored in physics. At the beginning of my fourth year I was called up in the Army Reserves (this was 1942) and returned to the University of Chicago as a senior in uniform. I undertook there an intensive course in meteorology and received a BS degree from the University of Chicago in 1943. Knowledge of meteorology has stood me in good stead throughout the years in refuting arguments that attempt to draw some sharp distinction between the precision and perfection of the physical sciences and the vagueness and imprecision of the social sciences. Meteorology is in theory a part of physics, but in practice more like economics, especially in the handling of a vast flow of nonexperimental data.

From Chicago I was sent to the South Pacific for two years of duty in 1944 and 1945. After a short period of adjustment, I found the isolation and serenity of the Solomon Islands quite attractive. I occupied myself with swimming, poker, Aristotle, and a couple of correspondence courses in mathematics and French. After a year of living on a small island, occupied only by military troops, I was transferred to Guam, which seemed relatively civilized but less conducive to intellectual work.

I was discharged from the Army Air Force in 1946, and after some months of deciding what to do, changing my mind any number of times, and spending more than half a year working for my father in the oil fields near Artesia, New Mexico, I entered Columbia University as a graduate student in philosophy in January of 1947 and received a PhD in 1950.

As an undergraduate I moved too often to be strongly influenced by any one teacher. I do remember certain impressive individuals, including Richard McKeon at Chicago, who lectured on Aristotle, Norman Steenrod, who taught my first course in calculus, and Professor Tanner at the University of Tulsa, from whom I learned elementary Greek.

The situation was different in graduate school. I was influenced by Ernest Nagel more than by anyone else and I still relish the memory of the first lecture of his that I attended in 1947. I came to philosophy

without much background in the subject, since my undergraduate training was primarily in mathematics and physics. But Nagel's skeptical, patient and detailed analysis of F. H. Bradley and John Dewey, in the first course I took from him, won my attention and interest from the start. I have recorded these impressions in more detail in my account of Nagel's lectures on Dewey (1969e). In those days, interdisciplinary work was very much favored in the Department of Philosophy, and consequently I continued to learn a good deal more mathematics and physics. I still remember with great pleasure the beautiful and penetrating lectures of Samuel Eilenberg on topology and group theory, and I remember well the course in relativity theory I took from L. H. Thomas. Thomas was a terrible lecturer, in sharp contrast to the brilliance of Nagel or Eilenberg, but he knew so much about the subject and was in such total control of the theory and relevant data that it was impossible not to learn a great deal and to be enormously impressed by the organization and presentation of the lectures.

In those years Columbia was swarming with returning veterans, and in certain ways we veterans had our own ideas about what should be taught in graduate school. We organized in 1947 or 1948 an informal seminar on von Neumann and Morgenstern's theory of games, partly because we did not seem to be able to find a course on the subject, but also because Columbia's graduate school immediately after the war was so vastly overcrowded that the kind of individual attention graduate students now receive at Stanford, for example, was simply unheard of in almost all departments at Columbia. I felt myself extremely lucky to have as much personal contact as I did with Ernest Nagel. Friends of mine in the History Department saw their adviser only a few hours during their entire graduate-school career, including completion of the dissertation.

Considering my relatively extensive research efforts in psychology from about 1955 onward, it is somewhat surprising that I took no work in psychology either as an undergraduate or as a graduate student, but there was a feature of my education that made it easier for me to pick up what I needed to know without prior systematic training. As an undergraduate I wandered about in several different fields, and because of the easygoing policy of the Department of Philosophy in those days at Columbia I spent a good deal of time in nonphilosophical courses. I thus developed early the habits of absorbing a wide variety of information and feeling at home in the problem of learning a subject in which I had not had much prior training or guidance.

Because of my background and interests in physics I wanted to write a dissertation about the philosophy of physics, in particular to give a modern axiomatic treatment of some branch of physics, but as I got deeper into the subject I realized that this was not the kind of dissertation that was considered appropriate in philosophy. The Department at that time was primarily historically oriented, and

Nagel advised me to consider a more informal approach to the foundations of physics. What we finally agreed on was a study of the concept of action at a distance, and: a good deal of the dissertation was devoted to an analytical study of this concept in the works of Descartes, Newton, Boscovich, and Kant. I was able to come closer to my original interest in a chapter on the special theory of relativity, but certainly what I had to say in that chapter was no contribution to the axiomatic foundations of the subject. I did find the historical work absorbing and have continued over the years to retain and, on occasion, profitably to use the knowledge about the history of philosophy and science I first systematically acquired as a graduate student at Columbia. The part of my dissertation I published, with only minor modifications, was the material on Descartes (1954a). But I have relatively recently used much more of the dissertation material in a long article on Aristotle's theory of matter (1974b), in which I also review the theories of matter of Descartes, Boscovich, and Kant.

Earlier, but still long after the dissertation was written, I used some of the material on Kant in a Festschrift article for John Herman Randall, Jr. (1967h). Randall's lectures in the history of philosophy were a great occasion at Columbia, and it was a pleasure to say something in detail about Kant's philosophy of science as an extension of Randall's interpretation in the Festschrift volume. Randall, like a lot of other philosophers mainly interested in the history of philosophy, did not have a strong analytical turn of mind, but his lectures were a memorable experience. As I said in the opening paragraph of my Festschrift article for Randall, "The wit, the literary quality, and the range of learning exhibited in these lectures were famous around Columbia long before the time of my own arrival. As a young scientist turned philosopher, the most important general thing I learned from Randall was not simply to read the great modern philosophers in terms of a close explication of text, but also to realize that they must be interpreted and considered against the background of the development of modern science." The contrast between Nagel's analytical and dialectical skill and Randall's sympathetic and impressionistic account of ideas was dramatic, but I can rightly claim that I learned a great deal from both of them although I was clearly more influenced by Nagel.

Another person whose intellectual habits influenced me at Columbia and who was of quite a different sort than any of the others I have mentioned was Paul Oskar Kristeller. He was interested in the fine details of historical scholarship, famous of course for his work in Renaissance philosophy, but I was more influenced by the seminar on Kant that he and Randall gave. Kristeller's meticulous insistence on a textual basis for every interpretation suggested in the seminar, and his decisive way of handling the text, were a model of scholarship that I admired and learned from, even though Kristeller always modestly insisted that he was in no sense a Kantian specialist.

I received my PhD from Columbia in June of 1950 but my education scarcely stopped there. Although I began teaching upon arrival at Stanford in the fall of 1950, where I have been ever since, almost as soon as I could I also began to think about research in the philosophy of science. My problem was that I did not really know much about how to do any serious research and writing, since my graduate education did not involve the personal supervision of research of the kind so familiar to graduate students today. I was neither forced nor encouraged to produce early in my graduate career a publishable paper nor to become acquainted with the "how-to-do-it" aspects of research.

I was, however, full of energy and brimming over with ideas. I thrashed around for a few months but fortunately I soon became acquainted with J. C. C. McKinsey, a logician who had recently joined the Department of Philosophy at Stanford. McKinsey served as my postdoctoral tutor. It was from him that I learned the set-theoretical methods that have been my stock in trade for much of my career. It was not, however, just set-theoretical methods as such that McKinsey taught me but also a passion for clarity that was unparalleled and had no precedent in my own prior education. I remember well his careful red-penciling of the initial draft I gave him of my first article on the theory of measurement (1951a). McKinsey was just completing a book on the theory of games and consequently also had some of the interests that were of importance to me as my own interests in the social sciences began to blossom.

But there were other people, at Stanford and Berkeley who continued my education in those early years at Stanford. With McKinsey's encouragement I attended Alfred Tarski's seminar in Berkeley. McKinsey always claimed that he had learned everything he knew from Tarski. This was not true, but after attending Tarski's seminar I understood why he liked to say this. Tarski was a ruthless taskmaster. I think he probably got the most out of his students of anyone I have known. His seminar provided perhaps the best example I know of vicarious learning. A student who made a poorly prepared seminar presentation was so ruthlessly and mercilessly questioned that the other students did not need any hints about the state of preparation they should achieve before making their own presentations. Tarski, as one of the great examples of the Polish school of logic, was unwilling to go forward on a single point unless every thing covered thus far was completely clear—in particular, unless it was apparent to him that the set-theoretical framework within which the discourse was operating could be made totally explicit. It was from McKinsey and Tarski that I learned about the axiomatic method and what it means to give a set-theoretical analysis of a subject.

McKinsey died in 1953, and our collaborative work on the axiomatic foundations of empirical sciences was brought to a sudden stop. Another learning experience that influenced much of my later work was the summer research position I had early in the fifties for a

couple of years with David Blackwell and M. A. Girshick while they were writing their influential book *Theory of Games and Statistical Decisions* (1954). As I remark later, I did not learn as much from them as I should have, but what I did absorb played an important role in a number of my early papers.

In later years I have learned a great deal from many persons who are mentioned below. Rather arbitrarily I have terminated this section on my education as of about 1954, but it seems important to mention those who have taught me much of what I know about psychology because, although my allegiance to philosophy had continued strong throughout all these years, for a considerable period I published more papers in psychology than in philosophy. Much of this work was due to the influence of William K. Estes, R. Duncan Luce, and Richard C. Atkinson. My joint work with each of them is discussed later at appropriate points.

II. Research

I have grouped the discussion of my research under seven headings: foundations of physics; theory of measurement; decision theory, foundations of probability, and causality; foundations of psychology; philosophy of language; education and computers; and philosophy and science. Where details warrant further organization, appropriate subheadings are also introduced. There is some arbitrariness in the classification, but I do not think any serious confusions will arise. The number of topics on which I have worked is large and it would be foolish to claim that I have contributed in a fundamental way to all of them. On the other hand, each of the main headings represents an area that has been the focus of my interests at some period over the past several decades.

Foundations of Physics

As already mentioned, my doctoral dissertation lay within the philosophy of physics. In particular, I studied the problem of action at a distance as it had occurred in 17th- and 18th-century physics and philosophy, especially in the writings of Descartes, Newton, Boscovich, and Kant. The final chapter dealt with the problem in the special theory of relativity. Working on it strengthened my earlier desire to give an axiomatic formulation of classical mechanics in the spirit of modern mathematics rather than 'physical' axiomatizations common in physics. Serious joint work on this project began soon after my arrival at Stanford, in collaboration with J. C. C. McKinsey, and is represented in four papers we wrote on the foundations of physics prior to McKinsey's death in 1953 (1953a, 1953b, 1953c also with A. C. Sugar, and 1955b). Shortly thereafter I wrote with Herman Rubin a similar paper (1954c) on the axiomatic foundations of relativistic particle mechanics. It is a long and very complicated piece of work that has not been read, I suspect, by very many people. My main interests soon turned to research on decision theory and related problems in the foundations of psychology but I have

continued through the years to have an interest in physics, and an irregularly spaced sequence of papers has reflected that interest.

Special Relativity

In 1957 I organized jointly with Leon Henkin and Alfred Tarski a symposium on the axiomatic method with special reference to geometry and physics. Working with Henkin and Tarski on the organization of this symposium was an exhilarating experience. My own paper (1959a) was concerned with the derivation of the Lorentz transformations of special relativity from an explicit but minimal set of assumptions. Essentially, the aim of the paper was to give an elementary derivation of the Lorentz transformations, without any assumptions of continuity or linearity, from a single axiom concerning invariance of the relativistic distance between any two space-time points connected by an inertial path or, put another way, from the assumption of invariance of segments of inertial paths. Important new results on the derivation of the Lorentz transformations have since been published by Walter Noll (1964), E. C. Zeeman (1964), and others.

The elegant aspect of Noll's paper is that he axiomatizes Minkowskian chronometry using coordinate-free methods. More importantly, Zeeman shows that it is not necessary to assume invariance of time-like intervals as I did, but that it is sufficient to assume the preservation of order, that is, the relativistic partial ordering of one point being after another is sufficient. Like many simple and beautiful ideas, it is surprising that this did not occur to someone sooner. The key to the results is already present in the early work of Robb (1936), which shows that the binary relation of being after is a sufficient conceptual basis for the kinematical theory of special relativity.

In the final section of my 1957 paper (1959a) I discuss the possibility of introducing a relation of signaling in order to fix the direction of time. It is obvious that this can be done very directly in an *ad hoc* fashion. What is needed, however, is some natural approach that is fully satisfying from an intuitive and a conceptual standpoint. In his article, Noll makes some remarks about this, and he raises the question of whether his approach solves the problem I raised. Essentially, Noll introduces a directed signal relation that is asymmetric, and of course if we postulate that the numerical representation must preserve the direction of signals passing from earlier to later events, the direction of time is guaranteed. I find this approach unsatisfactory since this is an arbitrary stipulation in the definition of isomorphism, and we get just as good an isomorphism from a structural standpoint if the direction in time is reversed. At the time it did not seem feasible to give such an analysis, but recently I have been rethinking and, more importantly, learning a great deal more about optics. It now seems to me that natural qualitative postulates differentiating signals being received and being sent should be feasible, although such postulates must go beyond the

purely kinematic aspects of the special theory of relativity to include some substantive assumptions, even if of a general nature, about physical optics.

After a gap of some years, my interest in classical physics (in which I include special relativity) was revived while working on Chapter 10 of *Foundations of Measurement* (1971a). This chapter is concerned with dimensional analysis and numerical laws. I think I made some contributions to the chapter, but almost surely I learned more than I contributed, especially from Duncan Luce. The rather technical material in this chapter on the algebra of physical quantities and on dimensional analysis, including the question of why numerical laws are dimensionally invariant, has not been much picked up by philosophers of science, but I think that in due time it will be.

Quantum Mechanics

Most of the effort that I have put in on the foundations of physics since 1960 has been devoted to quantum mechanics, and this continues to be a current active intellectual interest. Almost everything that I have written about quantum mechanics has been intertwined with questions related to the foundations of probability, especially as to how probabilistic concepts are used in quantum mechanics. My first paper on the subject (1961c) was concerned with the absence of a joint distribution of position and momentum in many standard cases. I shall not enter into the technical details of the argument here, but I do want to convey the basic philosophical point that I continue to find the real puzzle of quantum mechanics. Not the move away from classical determinism, but the ways in which the standard versions seem to lie outside the almost universal methodology of modern probability theory and mathematical statistics. For me it is in this arena that the real puzzles of quantum mechanics are to be found. I am philosophically willing to violate classical physical principles without too many qualms, but when it comes to moving away from the broad conceptual and formal framework of modern probability theory I am at once uneasy. My historical view of the situation is that if probability theory had been developed to anything like its current sophisticated state at the time the basic work on quantum mechanics was done in the twenties, then a very different sort of theory would have been formulated.

It is worth recording a couple of impressions about this because they indicate the kind of changes that can take place in one's attitudes as the years go by. Initially I was much impressed by the mathematical formulation of foundations given by von Neumann in his classical work and, later, by Mackey (1963), whose book has also become classical in its own way. No doubt I was originally struck by the mathematical clarity and sophistication of this work, but in later years I have become dissatisfied with the unsatisfactory conceptual basis from a probabilistic standpoint of the way in which the theory is formulated. I shall give here just two examples to indicate the nature of my conceptual dissatisfaction. Von Neumann stresses that we can take the expectation of the sum of any two operators, even

though they are conjugate, that is, do not commute. But once this is said, the natural question is to ask about the underlying probability space that justifies the exact probabilistic meaning of the expectation. A similar question arises with respect to Mackey's treatment. Mackey takes as fundamental the concept of the probability that a measurement in a given state of an observable will lead to a given value. This seems innocent enough, but when the fundamental postulates of the theory are stated in these terms, what seems missing from what one would expect in a standard causal physical theory is any clarity about the relation between observables. The axioms he gives would seem to concentrate too deeply on the relatively simple properties of the probability of a given measurement on a given observable and not enough on the causal dependencies between observables. (It is important to remember that I am not really making a technical argument here but trying to give the intuitions back of arguments that I think can be formalized.)

A detailed analysis of the kinds of requirements that a satisfactory probabilistic theory of quantum mechanical phenomena should have is laid out for the simple case of the harmonic oscillator in a relatively recent paper by Zanotti and me (1976m). As I write this autobiography I am struggling to develop in a much deeper way than I have previously a thoroughgoing probabilistic interpretation of interference phenomena in quantum mechanics, especially as produced by the classical two-slit experiment.

Until recently I thought that the most important philosophical problems of quantum mechanics could be analyzed in the nonrelativistic framework characteristic of the early development of the subject. An effort to understand the place of probability concepts in relativistic quantum mechanics has abruptly changed my mind. The meshing of probability and relativity seems to be badly worked out for even the most elementary problems. One sign of the difficulty is the superficiality of the development of probabilistic concepts in relativistic quantum mechanics. An extensive search of the literature, for example, has not revealed a single discussion of particular distributions. The multivariate normal distribution is invariant under linear transformations, and the Lorentz transformations are linear, but the proper space-time hyperplane on which the distribution is defined needs to be correctly chosen to be Lorentz invariant as well. As far as I know, discussion of these 'first' questions does not yet exist, which I find very surprising.

Although I continue to be fascinated by the conceptual problems of quantum mechanics and I think of it almost as a responsibility of any philosopher of science with wide interests to know a good many of the details of what is surely the most important scientific theory of the 20th century, I find that my own work here is less satisfying to me than other areas I discuss later, just because I do not anticipate making a scientific contribution to the subject. In the case of the theory of measurement or psychology, for example, my contributions

have as much a scientific as a philosophical flavor, and I find that this suits my temperament better. Independent of whether the scientific contribution is of greater or less significance than the philosophical one.

Theory of Measurement

In my first published article (1951a) I gave a set of independent axioms for extensive quantities in the tradition of earlier work by Hölder and Nagel. My contribution was primarily to weaken the assumptions of Hölder axioms and also to prove that both the axioms and the concepts used were independent. Looking around for other topics in measurement, and returning to the earlier interest in the theory of games and utility theory, it soon became apparent that there were more outstanding problems of measurement in psychology than in physics. One of my first efforts in this direction was a joint article with my student Muriel Winet (1955 d). We gave an axiomatization of utility based on the notion of utility differences. The idea of considering such utility differences is a very old one in the literature, but an explicit and adequate set of axioms had not previously appeared. In 1956 I published two other articles which fell between decision theory and measurement theory. One was on the role of subjective probability and utility in decision making. In this article (1956b) I used the results of the joint work with Winet to provide an axiomatization alternative to that given by Savage in his book *Foundations of Statistics* (1954). And in the second article, my colleague Donald Davidson and I gave a finitistic axiomatization of subjective probability and utility (1956c).

Shortly after this I began to think more generally about the foundational aspects of theories of measurement and was fortunate to have as a collaborator the logician and mathematician Dana Scott, who was at that time a graduate student in mathematics. (Scott is also one of the Berkeley-Stanford persons from whom I learned a great deal, beginning when he was an undergraduate in a course on the philosophy of science I taught at Berkeley in 1952, along with Richard Montague. What a pair to have in such a course!) Scott and I tried to give a general framework for theories of measurement and to obtain some specific results about axiomatization. This article was published in 1958, a year or so after it was written. The framework that Scott and I set up has, I think, been of use in the literature, and probably the article with him has been the most important article in the theory of measurement that I have written, although the chapter in the *Handbook of Mathematical Psychology*, written with J. L. Zinnes and published in 1963, has perhaps been more influential, especially in psychology.

My most important recent effort has been the extensive collaboration with David Krantz, Duncan Luce and Amos Tversky in the writing of our two-volume treatise *Foundations of Measurement*, the first volume of which appeared in 1971. At the time of writing this autobiography, we are hard at work on Volume II. My present

feeling is that when Volume II is published I shall be happy to let the theory of measurement lie fallow for several years. It is, however, an area of peculiar fascination for a methodologist and philosopher of science like myself. The solution of any one problem seems immediately to generate in a natural way several more new problems. The theory nicely combines a demand for formally correct and explicit results with the continual pursuit of analyses that are pertinent to experimental or empirical procedures in a variety of sciences but especially in psychology, where the controversy about the independent measurability of psychological concepts has been long and intense. The theory of measurement provides an excellent example of an area in which real progress has been made in the foundations of psychology. In earlier decades psychologists accepted the mistaken beliefs of physicists like Norman Campbell that fundamental measurement in psychology was impossible. Although Campbell had some wise things to say about experimental methods in physics, he seemed to have only a rather dim grasp of elementary formal methods, and his work in measurement suffered accordingly. Moreover, he did not even have the rudimentary scholarship to be aware of the important earlier work of Helmholtz, Hölder, and others.

The work of a number of people over the past several decades has led to a relatively sophisticated view of the foundations of measurement in psychology, and it seems unlikely that any substantial retreat from this solid foundation will take place in the future. I am somewhat disappointed that the theory of measurement has not been of greater interest to a wider range of philosophers of science. In many ways it is a natural topic for the philosophy of science because it does not require extensive incursions into the particular technical aspects of any one science but raises methodological issues that are common to many different disciplines. On the other hand, by now the subject has become an almost autonomous technical discipline, and it takes some effort to stay abreast of the extensive research literature.

Although important contributions to the theory of measurement have already appeared since we published Volume I of *Foundations of Measurement*, I do think it will remain as a substantial reference work in the subject for several years. What is perhaps most important is that we were able to do a fairly good job of unifying a variety of past results and thereby providing a general framework for future development of the theory.

Having mentioned the seminars of Tarski earlier, I cannot forbear mentioning that perhaps the best seminar, from my own personal standpoint, that I ever participated in was an intensive one on measurement held jointly between Berkeley and Stanford more than ten years ago when Duncan Luce was spending a year at the Center for Advanced Study in the Behavioral Sciences at Stanford. In addition to Luce and me, active participants were Ernest Adams, who is now Professor of Philosophy at Berkeley and was in the fifties my first PhD student, and Fred Roberts, who was at that time a graduate

student in mathematics at Stanford and is now a member of the Department of Mathematics at Rutgers University. William Craig also participated on occasion and had penetrating things to say even though he was not as deeply immersed in the subject as were the rest of us. Our intensive discussions would often last well beyond the normal two hours, and it would not be easy to summarize all that I learned in the course of the year.

There is also a pedagogical point about the theory of measurement, related to what I have said just above about measurement in the philosophy of science, that I want to mention. The mathematics required for elementary examples in the theory of measurement is not demanding, and yet significant and precise results in the form of representation theorems can be obtained. I gave several such examples in my textbook in logic (1957a) and also in my paper 'Finite Equal-interval Measurement Structures' (1972d). I continue to proselytize for the theory of measurement as an excellent source of precise but elementary methodology to introduce students to systematic philosophy of science.

Decision Theory, Foundations of Probability, and Causality

Decision Theory

It is not easy to disentangle measurement theory and decision theory because the measurement of subjective probability and utility has been such a central part of decision theory. The separation that I make will therefore be somewhat arbitrary. My really serious interest in psychology began with experimental research on decision theory in collaboration with my philosophical colleague Donald Davidson and a graduate student in psychology at that time, Sidney Siegel. Davidson and I had begun collaborative work with McKimsey in 1953 on the theory of value and also on utility theory. We continued this work after McKimsey's death, and it is reflected in Davidson, McKimsey, and Suppes (1955a) and in the joint article with Davidson (1956c) on the finitistic axiomatization of subjective probability and utility, already mentioned. The article on the measurement of utility based on utility differences, with Muriel Winet, was also part of this effort.

Sometime during the year 1954, Davidson and I undertook, with the collaboration of Siegel, an experimental investigation of the measurement of utility and subjective probability. Our objective was to provide an explicit methodology for separating the measurement of the two and at the same time to obtain conceptually interesting results about the character of individual utility and probability functions. This was my first experimental work and consequently in a genuine sense my first real introduction to psychology. The earlier papers on the foundations of decision theory concerned with formal problems of measurement were a natural and simple extension of my work in the axiomatic foundations of physics. Undertaking experimental work was quite another matter. I can still remember our

many quandaries in deciding how to begin, and seeking the advice of several people, especially our colleagues in the Department of Psychology at Stanford.

I continued a program of experimentation in decision theory as exemplified in the joint work with Halsey Royden and Karol Walsh (1959i) and the development of a nonlinear model for the experimental measurement of utility with Walsh (1959j). This interest continued into the sixties with an article (1960g) on open problems in the foundations of subjective probability. Then in 1961 I drew upon my interest in learning theory to try to create a behavioristic foundation for utility theory (1961a), and I also made an attempt in that same year to explain the relevance of decision theory to philosophy (1961b).

The most important effort in this period was the writing with Duncan Luce of a long chapter, 'Preference, Utility and Subjective Probability' (1965i), for Volume III of the *Handbook of Mathematical Psychology*. The organization of a large amount of material and the extensive interaction with Luce in the writing of this chapter taught me a great deal that I did not know about the subject, and I think the chapter itself has been useful for other people. It is also worth mentioning that large parts of the joint effort with Krantz, Luce and Tversky in writing our two-volume treatise on the foundations of measurement have been concerned with decision theory.

In the latter part of the sixties I wrote several articles in the foundations of decision theory, oriented more toward philosophy than psychology. Three of the articles appeared in a book on inductive logic edited jointly with Jaakko Hintikka, my part-time philosophical colleague at Stanford for many years.

One article dealt with probabilistic inference and the concept of total evidence (1966j). Here I advanced the argument that under a Bayesian conception of belief and decision there was no additional problem of total evidence, contrary to the view held by Carnap and also Hempel. According to this Bayesian view, which I continue to believe is essentially right on this matter, if a person is asked for the probability of an event at a given time, it will follow from the conditions of coherence on all of his beliefs at that time that the probability he assigns to the event automatically takes into account the total evidence that he believes has relevance to the occurrence of the event. The way in which total evidence is brought in is simple and straightforward; it is just a consequence of the elementary theorem on total probability.

A second article in the volume (1966e) set forth a Bayesian approach to the paradoxes of confirmation made famous by Hempel many years ago. I will not outline my solution here but much of the philosophical literature on the paradoxes of confirmation has taken

insufficient account of the natural Bayesian solution, at least so I continue to think. A third article in the volume (1966f) dealt with concept formation and Bayesian decisions. Here I attempted to set forth the close relations between formal aspects of the psychology of concept formation and the theory of Bayesian decisions. I now think that the ideas I set forth here are the least interesting and the most transitory of those occurring in the three articles. The general idea of value in this article concerns the relation expressed between concept formation and the classical problem of induction. For those restricted settings in which no new concepts are needed but for which an induction about properties is required, a Bayesian approach is sound and can meet most, if not all, of the conceptual problems about induction that I regard as serious. On the other hand, a Bayesian viewpoint toward induction does not provide a general solution because it does not incorporate a theory of concept formation. Genuinely new inductive knowledge about the world requires not only a framework of inductive inference of the sort well worked out in the contemporary Bayesian literature but also a theory about how new concepts are to be generated and how their applicability is to be dealt with. This large and significant aspect of the general problem of induction seems to me still to be in a quite unsatisfactory state. In my own thinking, the problem of induction and the concept of rationality are closely tied together, and as I point out in the article on probabilistic inference mentioned above, the Bayesian approach still provides a very thin view of rationality, because the methods for changing belief as reflected in the introduction of new concepts or in the focus of attention are not at all adequately handled. The outlines of any future theory that will deal in even a partially satisfactory way with the central problem of concept formation are not at all visible, and it may even be that the hope for a theory that approaches completeness is mistaken.

Distributive Justice

For a variety of reasons, the literature on decision theory has been intertwined with the literature on social choice theory for a very long period, but the focus of the two literatures is rather different and I have certainly had more to say about decision theory than about the normative problems of social choice or distributive justice. To a large extent, this is an accident of where I have happened to have had some ideas to develop and not a matter of a priori choice. I have published two papers on distributive justice (1966i, 1977a). The main results about justice in the first one, which were stated only for two persons, were nicely generalized by Amartya Sen (1970). The other paper, which was just recently published, looks for arguments to defend unequal distributions of income. I am as suspicious of simplistic arguments that lead to a uniform distribution of income as I am of the use of the principle of indifference in the theory of beliefs to justify a uniform prior distribution. The arguments are too simple and practices in the real world are too different. A classical economic argument to justify inequality of income is productivity, but in all societies and economic subgroups throughout the world differences

in income cannot be justified purely by claims about productivity. Perhaps the most universal principle also at work is one of seniority. Given the ubiquitous character of the preferential status arising from seniority in the form of income and other rewards, it is surprising how little conceptual effort seems to have been addressed to the formulation of principles that justify such universal practices. I do not pretend to have the answer but I believe that a proper analysis will lead deeper into psychological principles of satisfaction than has been the case with most principles of justice that have been advanced. I take it as a psychological fact that privileges of seniority will continue even in the utopia of tomorrow and I conjecture that the general psychological basis of seniority considerations is the felt need for change. A wide range of investigations demonstrate the desirable nature of change itself as a feature of biological life (not just of humans) that has not been deeply enough recognized in standard theories of justice or of the good and the beautiful.

Foundations of Probability

The ancient Greek view was that time is cyclic rather than linear in character. I hold the same view about my own pattern of research. One of my more recent articles (1974 g) is concerned with approximations yielding upper and lower probabilities in the measurement of partial belief. The formal theory of such upper and lower probabilities in qualitative terms is very similar to the framework for extensive quantities developed in my first paper in 1951. In retrospect, it is hard to understand why I did not see the simple qualitative analysis given in the 1974 paper at the time I posed a rather similar problem in the 1951 paper. The intuitive idea is completely simple and straightforward: A set of 'perfect' standard scales is introduced, and then the measurement of any other event or object (event in the case of probability, object in the case of mass) is made using standard scales just as we do in the ordinary use of an equal-arm balance. This is not the only occasion in which I have either not seen an obvious and simple approach to a subject until years later, or have in fact missed it entirely until it was done by someone else.

On the other hand, what would appear to be the rather trivial problem of generalizing this same approach to expectations or expected utility immediately encounters difficulties. The source of the difficulty is that in the case of expectations we move from the relatively simple properties of subadditive and superadditive upper and lower measures to multiplicative problems as in the characteristic expression for expected utility in which utilities and probabilities are multiplied and then added. The multiplicative generalization does not work well. It is easy to give a simple counterexample to straightforward generalization of the results for upper and lower probabilities, and this is done in Suppes (1975a). I have continued to try to understand better the many puzzles generated by the theory of upper and lower probabilities, in joint research with Mario Zanotti (1977j).

Partly as a by-product of our extensive discussions of the qualitative theory of upper and lower probabilities, Zanolini and I (1976n) used results in the theory of extensive measurement to obtain what I think are rather elegant necessary and sufficient conditions for the existence of a probability measure that strictly agrees with a qualitative ordering of probability judgments. I shall not try to describe the exact results here but mention the device used that is of some general conceptual interest.

Over the years there have been a large number of papers by many different individuals on these matters. Essentially all of them have formulated conditions in terms of events, with the underlying structure being that of the Boolean algebra of events and the ordering relation being a binary relation of one event being at least as probable as another. The conditions have turned out not to be simple. The important aspect of the paper with Zanolini is our recognition that events are the wrong objects to order. To each event there is a corresponding indicator function for that event, with the indicator function having the value one when a possible outcome lies in the event and the outcome zero otherwise—as is apparent, in this standard formulation events are sets of possible outcomes, that is, sets of points in the probability space. We obtain what Zanolini and I have baptized as extended indicator functions by closing the set of indicator functions under the operation of functional addition. Using results already known in the theory of extensive measurement it is then easy to give quite simple necessary and sufficient axioms on the ordering of extended indicator functions to obtain a numerical probability representation.

Recently we have found correspondingly simple necessary and sufficient qualitative axioms for conditional probability. The qualitative formulations of this theory beginning with the early work of B. O. Koopman (1940a, 1940b) have been especially complex. We have been able drastically to simplify the axioms by using not only extended indicator functions, but the restriction of such functions to a given event to express conditionalization. In the ordinary logic of events, when we have a conditional probability $P(A|B)$, there is no conditional event $A|B$, and thus it is not possible to define operations on conditional or restricted events. However, if we replace the event A by its indicator function A_c , then $A_c|B$ is just the indicator function restricted to the set B , and we can express in a simple and natural way the operation of functional addition of two such partial functions having the same domain. The analysis of conditional probability requires considerably more deviation from the theory of extensive measurement than does the unconditional case: for example, addition as just indicated is partial rather than total. More importantly, a way has to be found to express the conceptual content of the theorem on total probability. The solution to this problem is the most interesting aspect of the axiomatization.

The move from events to extended indicator functions is especially interesting philosophically, because the choice of the right objects to consider in formulating a given theory is, more often than I originally thought, mistaken in first efforts and, as these first efforts become crystallized and familiar, difficult to move away from.

Apart from technical matters of formulation and axiomatic niceties, there are, it seems to me, three fundamental concepts underlying probability theory. One is the addition of probabilities for mutually exclusive events, the second is the concept of independence of events or random variables, and the third is the concept of randomness. I have not said much here about either independence or randomness. A conceptually adequate formulation of the foundations of probability should deal with both of these concepts in a transparent and intuitively satisfactory way. For any serious applications there is a fourth notion of equal importance. This is the notion of conditionalization, or the appropriate conceptual method for absorbing new information and changing the given probabilities. I have ideas, some of which are surely wrong, about how to deal with these matters and hope to be able to spend time on them in the future. However, rather than try to sketch what is still quite premature, I want to end with some general comments about the foundations of probability and decision theory.

It has been remarked by many people that logic is now becoming a mathematical subject and that philosophers are no longer main contributors to the subject. Because of its more advanced mathematical character this has really been true of probability from the beginning. The great contributions to the foundations of probability have been made by mathematicians—de Moivre, Laplace, von Mises, and Kolmogorov come quickly to mind. Although there is a tradition of these matters in philosophy—and here one thinks of Reichenbach and Carnap—it is still certainly true that philosophers have not had a strong influence on the mainstream of probability theory, even in the formulation of its foundations. On the other hand, I strongly believe in the proposition that there is important and significant work in the foundations of probability that is more likely to be done by philosophers than by anyone else. The various interpretations of foundations, ranging from the subjective view of the classical period through the relative frequency theory of the first part of this century to propensity and other views of late, have probably been discussed more thoroughly and more carefully by philosophers than by anyone else. I see no reason to think that this tradition will come to an end. The closely related problems of decision theory are just beginning to receive equal attention from philosophers after their rapid development by mathematical statisticians in the two decades after World War II.

It is important for philosophers to be familiar with and to know the formal and technical developments by mathematicians and statisticians. It is unfortunate that there has been a tendency for

philosophers to pursue their own special formalisms that do not relate well to the mainstream of work. Such formalisms tend to be, from a mathematical and conceptual standpoint, too elementary to come to grips with complex problems of applications or to offer sufficient complexity of structure to handle the main problems of interest to those pursuing technical issues. What seems to me to be the right role for philosophers in these matters is to be able to comment on and to use the concepts that are actively developed in most cases by others. I do not see as viable a special philosophical approach to probability, and my views on this matter are consonant with what I think about other issues in the philosophy of science or in philosophy generally.

Causality

Because my own approach to causality is probabilistic in character, I have included it in this section. It is hard to think of a philosophical topic that has received more attention historically than that of causality. It has already become clear to me that what I have had to say (1970a) has got to be extended, revised, and deepened, in order to meet objections that have been made by other people and to account for a variety of phenomena that I did not consider in any detail. Causality is one of those concepts that plays a major role in a variety of scientific disciplines and that can be clarified and enriched by extensive philosophical analysis. On some subjects of a probabilistic kind I find it hard to imagine how I, or another philosopher, could improve in a substantial way on what has been said with clarity and precision by probabilists and statisticians—the concept of a stochastic process is a good example. This is not true of the concept of causality. A good many statisticians use the concept in various ways in their research and writing, and the concept has been a matter of controversy both in the physical sciences and in the social sciences over the past several decades. There is a major place in these discussions for philosophical analyses of causality that join issue firmly and squarely with this extensive scientific literature.

A recent article by Woods and Walton (1977) emphasizes a point that is something of a minor scandal in philosophy. This is the absence of clear and definite elementary principles for accepting or rejecting a causal relation. The teaching of elementary logic depends upon extensive use of material implication and other truth-functional sentential connectives, in much the same way that beginning students of physics are taught Newtonian and not relativistic mechanics. We unfortunately do not at the present time have the same tradition in philosophy about a range of concepts that lie outside of formal logic. Causality is perhaps the prime example. I mention the point as a matter of pedagogy but in fact it is a matter of philosophy proper, because there has not been sufficient development or agreement about the developments that have taken place to provide a set of transparent systematic concepts that can be used in introductory teaching.

There are one or two systematic points about causality I would like to comment on here without entering into technical details. The first is the objection to my characterization of causality in terms of probability. A standard remark about this characterization is that all kinds of spurious relations will satisfy the definition of *prima facie* cause. According to my formulation, an event A is a *prima facie* cause of event B if A occurs earlier than B and the conditional probability of B given A is greater than the unconditional probability of B alone. It is properly pointed out that many kinds of events are connected in the sense of this definition by a *prima facie* causal relation, for example, the lowering of the barometer and the rain that follows, and yet we all immediately reject the falling of the barometer as a *prima facie* cause of the rain. I see the situation here as no different from that which applies to logical inference. The machinery is set up to be indifferent to our intuitive facts about the world, so that we can make logical inferences that seem silly. Standard examples are easy to give and are familiar to everybody. The same point is not as easily accepted about causality, but it is my claim that this is a virtue and not a defect of a general theory of causality. It should be universally applicable; intuitive counterexamples simply reflect the fact that the formal theory is indifferent as to what intuitive knowledge or substantive theory is being called upon.

Moreover, the full formal theory has appropriate devices for eliminating falling barometers as causes of rain. The standard notion to use here is that of a spurious cause. Showing that other events account for the change in conditional probability when the barometer is not present or broken provides the intuitive evidence we all accept for the absence of a causal relation between falling barometers and rainfall. The second point, related to this one, is that the notion of spurious cause itself and the closely related one of genuine cause must be relativized to a particular conceptual framework. This is made especially clear when one wants to prove a theorem about causality within the framework of a particular scientific theory. In my 1970 monograph I did not make the relativization to a particular framework an explicit part of the definitions. It is obvious how this can be done and perhaps in many cases it should be done. I do think that the insistence on relativizing the analysis of cause to a particular conceptual framework is a point on which to make a stand. Absolutists, who think they know the full truth and do not need such relativization, have the burden of providing forceful examples. I know of no interesting ones myself. I take this point as no different than the point that the systematic formal concept of truth is relative to a model and not in any sense appropriate to reality taken straight. There is another and more interesting point raised in conversations on various occasions by Nancy Cartwright, Paul Holland, and others. It is that the full notion of causality requires a sense of experimental manipulation. There are many ways of formulating the idea. Holland likes to say that, from a statistical standpoint, without random assignment of individuals to experimental groups an unimpeachable

causal inference cannot be made. My most immediate reply is that ordinary talk and much scientific experience as well does not in any sense satisfy these conditions of experimental design, that is, the causal claims that are made in ordinary talk or in much of science have not arisen from well-designed experiments but from quite different circumstances—in fact, from circumstances in which no experiments have taken place and in many cases are not possible. The great classical example is celestial mechanics. From the time of the Babylonians to the present, we have seen a variety of causal theories to account for the motion of the planets, the moon, and the stars. In the case of some terrestrial phenomena that are not themselves directly subject to experiment but for which an analysis can be built up in terms of experimental data, we are faced with a rather more complicated decision about what we regard as proper extrapolation from experiment. In fact, one underhanded way to meet the objections raised by Cartwright and Holland is to point out that the use of scientific theories outside the experimental domain and the power of the application of science depend upon sustaining causal claims in nonexperimental settings. Are we to conduct experiments on extendability in order to establish a justification of using the results of experiments in nonexperimental settings? It would not be difficult to set up a straw man of infinite regress by literal pursuit of this line of thought. My own view is that, rather than claiming that only in experimental settings can we really make proper causal claims, we should formulate theorems that are applicable to experimental settings but not to others. It seems to me one kind of theorem we might want to insist upon is that for experiments whose theory of design is adequate we should expect to be able to prove within a framework of explicit probabilistic concepts that all prima facie causes are genuine. We would not expect such a theorem to hold in general in nonexperimental settings.

Kreisel has pointed out to me that the general theory of causality is unlikely to be of much scientific significance once specific scientific theories are considered. Indeed, the interest of such theories is to provide a testing ground for the correctness of the general notions. On the other hand, not only ordinary talk but much highly empirical scientific work does not depend on a well-defined theoretical framework, and for these cases the general theory of causality can provide useful analytic concepts.

Foundations of Psychology

I have already remarked on my earliest experimental work in psychology in connection with the test of various concepts and axioms of decision theory. I shall not refer further to that work in this section. Because of my extensive work in psychology over the past two decades, I have organized my remarks under four headings: learning theory, mathematical concept formation in children, psycholinguistics, and behaviorism.

Learning theory

Either in my last months as a graduate student at Columbia or shortly after my arrival at Stanford in the fall of 1950—I cannot remember which—I developed my first interest in learning theory. As might easily be surmised, it began with trying to understand the various works of Clark Hull, not only the *Principles of Behavior* (1943) but also the relatively unreadable work written earlier in collaboration with the Yale logician Frederick Fitch and others (Hull, Hovland, Ross, Hall, Perkins, & Fitch, 1940). Part of my interest was stimulated by some very bright graduate students in psychology who attended my lectures in the philosophy of science. Probably the most influential was Frank Restle. I was a member of his dissertation committee, but I am sure I learned more psychology from him than he learned from me. My serious interest in learning theory began, however, in 1955 when I was a Fellow at the Center for Advanced Study in the Behavioral Sciences. Restle was there, but even more important for my future interests was the presence of William K. Estes. In his own and very different way, Estes has the kind of intellectual clarity I so much admired in McKinsey and Tarski. We began talking seriously about the foundations of stimulus sampling theory, which really began with Estes' s classical paper (1950). It became apparent to me quite soon that stimulus sampling theory was from a conceptual and mathematical standpoint much more viable and robust than the Hullian theory of learning. No really interesting mathematical derivations of experimentally testable results could be made from Hull's axioms. The great virtue of stimulus sampling theory was that with variation of experimental conditions new experimental predictions could be derived in an honest way without the introduction of *ad hoc* parameters and with the hope of detailed experimental test.

Perhaps it will be useful to say something more about the contrast between Hull's theory and stimulus sampling theory. The mere use of mathematics and especially of mathematical symbols in a theory is no guarantee that any thing of mathematical or scientific interest is being done. In Hull's theory the feeling is too much that each experimental result, or almost each remark about an experiment, is being given a direct translation into mathematical symbols. In contrast, no powerful and simple set of theoretical postulates from which specific results can be derived, once initial and boundary conditions are described, is even remotely approached. The translation from ordinary mathematical statements into the still more formal apparatus of mathematical logic as exemplified in the 1940 work of Hull and others cited above is still more mistaken if the only objective is a translation. One of the great lessons of logic in the 20th century is that formal systems themselves as deductive instruments are not of as much conceptual importance as the mathematical study of properties of such systems, but it is precisely the mathematical analysis of psychological theory that is not even touched upon in Hull's work. Hull, on the other hand, is in good company in this mistaken move. It has taken so much time for the situation to become

clarified as a result of the large body of important work in mathematical logic and metamathematics in this century. In Volume III of Whitehead and Russell's *Principia Mathematica* (1913) a detailed theory of measurement is developed; but from a formal standpoint their results are elementary, the notation is so forbidding and the welter of symbols so profuse that very little use has subsequently been made of the material.

In order not to seem too dogmatic about this point, it is worth noting that the interest in the use of formal systems has returned in new guise in the form of programming languages, but here the orientation is very different from that to be found, for example, in Hull's *Mathematico-Deductive Theory of Rote Learning*.

In contrast to Hull's theory, stimulus sampling theory works in a way very analogous to that of physical theories. The initial probabilities of response correspond to initial conditions in a physical problem, and reinforcement schedules correspond closely to boundary conditions. The fruits of extensive collaboration with Estes in 1955-1956 did not appear until later. In fact, we have published only two articles together (1959e, 1974q). The first article appeared as a technical report in 1957 and the second one first appeared as a technical report in 1959. The collaboration with Estes in writing the two long technical reports, later condensed into shorter papers, has been one of my most satisfactory research efforts in psychology. In a genuine sense these two reports combined a concern for axiomatic foundations with a focus on new scientific results.

While Estes and I were together at the Center, he introduced me to his former graduate student, Richard C. Atkinson, and we arranged for Atkinson to spend the following academic year at Stanford as a research associate. He and I undertook an extensive series of investigations into the application of stimulus sampling theory to two-person interactions. The initial fruit of this collaboration was my first experimental article in psychology (1958a). (Unlike most psychologists, I published an experimental book (1957b) before publishing an experimental article.) Atkinson and I expanded this first effort into an extensive series of studies, which were published in our book, *Markov Learning Models for Multi-person interactions* (1960b). It was a great pleasure to me to work in this area of application of learning theory. It combined my fundamental interest in learning theory with my earlier interest in game theory, and I found that I had a natural taste for elaborate analysis of experimental data. (The book with Atkinson is much richer in data analysis and the testing of models than is the earlier book with Davidson and Siegel.) I think I work best with someone like Atkinson, who is extremely well organized and very good at designing and running experiments. I like to get into the action when the analysis of the data and the testing of theory are to be the focus. Working with Atkinson has the additional advantage that he is also an able theorist and has plenty of ideas of his own.

This was the period in which my theoretical interests in learning theory were flourishing. I also worked at this time on mathematical aspects of learning processes, particularly a study of their asymptotic properties in collaboration with John Lamperti, who was then a member of the mathematics faculty at Stanford. This work appeared in two publications (1959h, 1960i).

I spent a fair amount of time on the generalization of learning theory to a continuum of responses. The first step was to generalize the linear model (1959b), and the second step was to generalize stimulus sampling theory (1960h). In collaboration with Raymond Frankmann, Joseph Zinnes, and later Henri Rouanet and Michael Levine, extensive tests of learning theory for a continuum of responses were made in publications between 1961 and 1964. With Jean Donio, who like Henri Rouanet was at that time a young French scientist working with me, I worked out the generalization in another direction to a continuous-time formulation of the theory. I am still pleased with the character of this work. It took a certain amount of mathematical effort to get things straight, and some of the detailed empirical predictions were quantitatively accurate and surprising. It is especially in predicting something like the continuous distribution of responses that untutored intuition fails in providing anything like an accurate idea of what the results will be. The need of theory to make non-trivial predictions becomes especially evident. On the other hand, I do not think that this work has had very much impact in psychology. The developments have not been followed up in the directions that could have led to more powerful results, but I do not think we were walking down a blind alley in this research effort. The kind of approach developed will almost surely turn out to be of use in several directions when a larger number of psychologists with strong mathematical training and quantitative empirical interests come onto the scene to study the theory of motor skills and a variety of perceptual phenomena in detail.

Mathematical Concept Formation in Children

In 1956 my oldest child, Patricia, entered kindergarten and my interests in applications were once again stimulated, in this case to thinking about the initial learning of mathematical concepts by children. In collaboration with Newton Hawley, who was and still is a member of the mathematics faculty at Stanford, we began the following year, when our daughters were both in the first grade, the informal introduction of constructive geometry. At that time very little geometry was taught in the primary grades. A brief description of this first effort is to be found in Hawley and Suppes (1959g), but, more importantly, we went on to write two textbooks for primary-grade students in geometry, which have since been translated into French and Spanish (1960c, 1960d).

This practical interest in the mathematics curriculum in the schools almost inevitably led to trying to understand better how children learn mathematical concepts. Once again because of my continued

collaboration with Estes, I was fortunate to get Rose Ginsberg, who was just completing a PhD at Indiana with Estes and C. J. Burke, to join me at Stanford as a research associate with the express purpose of studying concept formation in children. This work resulted in a number of publications with Ginsberg (1962d, 1962e, 1963c) and also a number of publications of my own, of which I mention especially my monograph *On the Behavioral Foundations of Mathematical Concepts*, appearing in the Monograph Series of the Society for Research in Child Development (1965e), and also my article on the same topic in the next year in the *American Psychologist* (1966g).

In this work in mathematical concept formation in children, Ginsberg and I were concerned to apply, as directly and as naturally as possible, stimulus sampling theory to the learning of such concepts. We met with more success than I initially expected. The ability of relatively simple models of stimulus sampling theory to account for the learning of simple mathematical concepts in great detail is, I think, a surprising fact and an important one.

In 1960, I was finishing my textbook on axiomatic set theory and the question naturally arose of the relation between the set-theoretical foundations of mathematics provided by Zermelo-Fraenkel set theory and the learning of mathematics by children. If I had not been finishing that text-book at the time I might well not have embarked upon a number of the experiments that Ginsberg and I undertook to test the formation of elementary concepts about sets—for example, identity and equipollence—as well as elementary geometrical concepts. A practical fruit of these investigations was the undertaking of a new series of elementary mathematics textbooks entitled *Sets and Numbers*, which was published over several years in the sixties. In recent years the interest in mathematical concept formation has melded into my work on computer-assisted instruction, which I discuss in a later section.

Psycholinguistics

Another natural area of application of stimulus sampling theory is language learning. Again I was fortunate to get another former student of Estes and Burke from Indiana, Edward Crothers, to join me as a research associate at Stanford. We undertook a systematic series of experiments in second-language learning; this effort led to several publications but especially to a book in 1967, *Experiments in Second-language Learning*, which reported a large number of investigations on various elementary aspects of learning Russian.

At the same time, given my philosophical inclinations and training, it was natural for me to become interested in the broader range of controversies in psycholinguistics. To some extent the first chapter of the book with Crothers summarizes the kind of attack on the problems of language learning to be expected from a behavioral standpoint. I also attempted in the writing of this chapter to provide a

partial answer to the many criticisms that were being made of behavioral theories by psycholinguists.

Two years later I published a more thoroughly worked out answer in my article 'Stimulus-Response Theory of Finite Automata' (1969g). In this article I showed that from simple principles of conditioning one could obtain the kind of language behavior of which a finite automaton is capable. I made no attempt to relate the theoretical developments to the detailed and subtle learning that takes place in a child, but rather argued that the presumed theoretical limitations of stimulus-response theory were misunderstood by a number of linguistically oriented critics. I have continued to be involved in this controversy and some of my most recent articles are concerned with it (1975b, 1977d).

I have emphasized in my writings on this subject that the challenge to psychological theory made by linguists to provide an adequate theory of language learning may well be regarded as the most significant intellectual challenge to theoretical psychology in this century. At the present time numerous difficult problems of providing a completely adequate scientific theory of language learning and language performance are enough to make even the most optimistic theorist uneasy. In very developed areas of science or mathematics, it is familiar to find the statement made that certain kinds of problems are simply far beyond the resources currently available but that certain more restricted problems are amenable to serious attack and likely solution. For example, in quantum chemistry there is, with present intellectual and computing resources, no hope of making a direct attack on the behavior of complex molecules by beginning with the first principles of quantum theory. A problem as easy to formulate as that of deriving from first principles the boiling point of water under normal atmospheric pressure is simply beyond solution at the present time and is recognized as such. Within mathematics there are classical open problems in elementary number theory, group theory, differential geometry, and in fact almost any developed branch of mathematics. Psycholinguistics will be a far happier and more productive subject when the same state of developed theory has been reached. A frontal attack on the problem of giving a complete analysis of the speech of even a three-year-old child is certainly outside the range of our conceptual tools at the present time. What seems essential is to recognize this fact and to determine the appropriate range of significant yet possibly solvable problems that should be studied. One approach that has already been fruitful and will be significant in the future is the attempt to write computer programs that can understand and learn a natural language. Such enterprises must at present be restricted to a small fragment of a natural language, but the thorough investigation of such small fragments seems to me a promising arena for making progress in the way that is characteristic of other domains of science. I do not mean to suggest by this that study of natural language should be restricted to computers—

certainly not. There will continue to be a significant and important accumulation of fact and theory about the language learning of children. Our understanding of these matters will deepen each year, but what is not yet clear is the direction theory will take so as to be adequate to the limited domains of understanding we master. I recognize the inadequacies from an empirical standpoint of what can presently be said about language learning within a stimulus-response framework or a more sophisticated version of S-R theory in terms of procedures and internal data structure, but I also believe in emphasizing the theoretical thinness of any of the proposals about learning that come from the linguistic side of psycholinguistics. In fact, practically none of the ideas originating from linguistics about language learning has been sufficiently developed in a systematic fashion to permit any sort of theorem, asymptotic or otherwise, to be proved.

Some may say that it is scarcely required of an empirical theory that it be precise enough to permit the proving of asymptotic theorems, but in the present context it seems to me an important consideration. The actual empirical phenomena are too complicated for any one coming at them from any theoretical viewpoint to provide a detailed account. One test of a theoretical proposal is whether its structure is rich enough to permit in principle the learning of language as the amount of exposure or experience goes to infinity. A good recent effort in this direction that does permit a theorem to be proved, even if the concept of meaning underlying the theory is not at all realistic, is to be found in Hamburger and Wexler (1975)—and I am pleased to claim Wexler as a former doctoral student of mine.

Toward the end of this period in the sixties I also got involved in the detailed empirical study of children's first language. The initial work in this area was in the construction of probabilistic grammars (1970b) followed by the construction of model-theoretic semantics for context-free fragments of natural language (1973e). I have been especially skeptical of the semantical concepts used by linguists. The long and deep tradition of semantical analysis in logic and philosophy provides, in my judgment, a much sounder basis for the analysis of the semantics of natural language. In my address as recipient of the American Psychological Association Distinguished Scientific Award in 1972, I tried to lay out the virtues of the model-theoretic approach to the semantics of children's speech (1974m). This is an issue that is still before us, and it would be too easy for me to enter into the substantive debate in this essay. I cannot refrain, however, from a few remarks.

The concept of meaning has a much longer history in philosophy and logic than it does in psychology or in linguistics. It is possible to begin the philosophical story with Aristotle, but Frege and Tarski will do as modern points of departure. The important thrust of this work has been to describe how the meaning of a sentence is to be built up from the meaning of its parts. In Tarski's case this takes the

form of giving an explicit recursive definition of truth. One of the reasons for the large disparity between this developed logical literature and what goes under the heading of meaning in psycholinguistics is, I believe, the concern in psycholinguistics primarily for the meaning of individual words. In Roger Brown's recent book (1973) he talks a good deal about the meanings of individual words and even about the meaning of imperatives, but what he does not really face at any point is the Fregean task of trying to understand how the meaning of a complex utterance is built up from the meaning of its parts. Without this there can be no serious theory of meaning, and until this is thoroughly recognized by psycholinguists I am skeptical that a satisfactory psychological theory of meaning for use in studying the first language of children can be developed.

I have undertaken additional large-scale empirical work on children's language in collaboration with my former students, Dr. Robert Smith and Dr. Elizabeth Macken, as well as with Madeleine L  veill   of the Laboratory of Experimental Psychology in Paris. The first fruits of this collaboration are to be found in the reports by L  veill  , Smith, and me (1973g; 1974t), and in two recent articles, one with Macken (1978f) and one by me (in press—b).

Of all my work in psychology, that concerning the syntax and semantics of children's first language has had the closest relation to broad issues that are current in philosophy, and for this reason alone I expect my interest to continue unabated. I have more to say on these matters in the section on philosophy of language.

Behaviorism

In spite of my recent interest in psycholinguistics I have certainly been identified in psychology with behaviorism, and I have written several philosophical pieces in defense of behaviorism. It should be clear from many other things I have to say in this essay that I do not believe in some Skinnerian form of behavioristic reductionism. In fact, the kind of methodological behaviorism, or what I have sometimes labeled neobehaviorism, I advocate is antireductionist in spirit, and wholly compatible with mentalistic concepts. The central idea of methodological behaviorism is that psychology as a science must primarily depend on behavioristic evidence. Such evidence is characterized in terms of stimuli and responses described in terms of psychological concepts. It is certainly possible to ask for physiological and, indeed, physical or chemical characterizations of both stimuli and responses. In some kinds of work, characterization of stimuli at a physical level is highly desirable, as for example in some detailed studies of visual perception. On the other hand, I strongly believe that a reduction of psychology to the biological or physical sciences will not occur and is not intellectually feasible. I am not happy with leaving the statement of my views at this level of generality, and I consider it an intellectual responsibility of methodological behaviorists like myself to reach for a deeper and

more formal statement of this antireductionist position. What are needed are theorems based on currently reasonable assumptions showing that such a reduction cannot be made. I think of such theorems as being formulated in the spirit in which theorems are stated in quantum mechanics about the impossibility of deterministic hidden variable theories.

Given my earlier work, as reflected for example in my paper on the stimulus-response theory of finite automata (1969g), it may seem paradoxical for me to be arguing for such impossibility theorems, but the thrust to develop a psychological theory of computable processes is to be understood as an effort to bring the theory of language and other complex psychological phenomena within the framework of methodological behaviorism.

As I have argued more than once in the past, stimulus-response theory of behavior stands in relation to the whole of psychology in much the same way that set theory stands to the whole of mathematics. In principle it offers an appealingly simple and rigorous approach to a unified foundation of psychology. It is worth examining the extent to which a stimulus-response 'reduction' of the rest of psychology is feasible. The first difficulty is that most of the other parts of psychology are not formulated in a sufficiently general and mathematical form, contrary to the case of mathematics where what was to be defined in set-theoretical terms already had a relatively precise characterization. Because I am not persuaded that the reductionistic approach is of much interest at the present time, I turn to stimulus-response theory itself and its internal difficulties.

Although I shall not enter into the technical details here it is not difficult to show that given an arbitrary Turing machine a stimulus-response model of behavior with only simple principles of stimulus sampling, conditioning and reinforcement operating can be constructed that is asymptotically (in time) isomorphic to the Turing machine. The tape of the machine is represented by some potentially infinite sequence of responses, for example, responses to the numerals as stimuli. From this asymptotic representation for any Turing machine, we can construct a universal stimulus-response model corresponding to a universal Turing machine that will compute any partial recursive function. Thus in principle we can claim the adequacy of stimulus-response theory to give an account of the learning of any computable process, and presumably any human cognitive or affective behavior falls within this context.

From a sufficiently general philosophical viewpoint this stimulus-response representation of any behavior no matter how complex is of some interest. It shows, just as does the representation of all of classical mathematics within set theory, how simple the primitive concepts of a powerful theory can be when there is no severe limitation on the means of construction. In particular the representation provides an abstract reduction of all concepts of behavior to the simple set required for the formulation of stimulus-

response theory, but the word abstract needs emphasis because we certainly have no idea how to carry out the actual reduction of most interesting behavior.

The basic representation of a universal Turing machine by a stimulus-response model brought to isomorphism at asymptote requires learning procedures that consist only of conditioning and change of conditioning of responses to given stimuli. But there is a severe weakness of these asymptotic results. Nothing is said about the learning rate. To match human performance or to be of real conceptual interest, the learning of appropriately simple concepts must not be too slow. Take the case of first-language learning in the child, for instance. A rather extravagant upper bound on the number of utterances a child has heard by the age of five is ten million. A learning rate that requires three orders of magnitude of exposure beyond this is not acceptable from a theoretical standpoint, but the highly simplified inadequate representation of the genetically endowed structures and functions the child brings to the task is evident. A richer theory is required to deal with them, but almost certainly there will be no fully satisfactory theory developed at any time in the foreseeable future.

Philosophy of Language

I have already said something about my interest in language in the section on psycholinguistics. Much but not all of my formal work on the theory of language has been related to psycholinguistics. Some overlap of the earlier section will be inevitable, but I concentrate here on the work that is more or less independent of psycholinguistics. My paper on model-theoretic semantics for context-free fragments of natural language (1973e) was partly generated by thinking about children's language but also by the formal problem of putting together the kinds of grammar that have become current in linguistics and computer science with the kind of model theory familiar in logic. The general approach has been anticipated by Knuth (1968) but he had not brought to the surface the model theory, and I did not become aware of his paper until I had worked out the essentials. My paper was first circulated as a technical report in 1971, and the ideas were applied with great thoroughness and extended by Robert Smith in his dissertation (1972), written under my direction. The basic philosophical point is that we can provide a standard model theory for context-free fragments of English directly without any recourse to the model theory of first-order logic. Because of my conviction that this can be done easily and naturally, I have continued to argue for the inappropriateness of first-order (or second-order) logic as an analytical tool for the study of the semantics of natural language. I summarize some of the further developments in later papers. Before doing so I want to note that the restriction to context-free grammars is not essential. One can work out a corresponding model-theoretic semantics for transformations that map trees into trees in the standard linguistic fashion, but because there is a great deal of work to be

done just within a context-free framework I have not worked out many details from a wider perspective.

In my address (1973b) as outgoing president of the Pacific Division of the American Philosophical Association, I applied these semantical ideas to develop a general notion of congruence of meaning. I characterized this view as a geometrical theory of meaning because I developed the viewpoint that different weak and strong notions of congruence are appropriate to catch different senses of 'identity' of meaning. Moreover, it seemed to me then and it still seems to me that there is much to be learned from geometry about the concept of congruence and the related concept of invariance that is applicable to the theory of meaning. We have long ago abandoned the idea of one true theory of geometry; we should do the same for meaning.

The most important idea in the paper is to tie closely together the notion of congruence of meaning and the particular syntactical structure of expressions. Without such an explicit use of syntax I am deeply skeptical that any satisfactory general theory of meaning can be found. In particular, the features of similarity of meaning between two expressions seem to me lost in any translation into first-order logic, and just for this reason I am doubtful of the appropriateness of the standard notions of logical form. In fact, I suppose my view is that there is no serious notion of logical form separate from the syntactic form of an expression itself.

On the other hand, I accept as of great importance the crude notion of congruence characterized by logical equivalence of expressions. This extensional concept of congruence is the robust and stable one needed for much scientific and mathematical work. Of course, for most systematic contexts a weaker notion of equivalence is used, one in which the notion of consequence is broader than that of logical consequence, because of the assumption of various background theories—all of classical mathematics in the case of physics, for example.

There is a point concerning identity of meaning that I did not develop explicitly enough in that address. In geometry we have a clear notion of identity for geometrical figures but it is not a notion that receives any real use compared to the importance of the notion of congruence for the particular geometry under study. It seems to me that this is very much the case with meaning. I am not unhappy with a very psychological approach to meaning that takes the meaning of a term to be unique at a given time and place to a given individual. Thus in crude terms the meaning of a proper name in this sense might well be taken to be the set of internal programs or procedures by which the individual that uses or recognizes the proper name attaches properties or relations to the object denoted by the proper name. These procedures or programs internal to a particular language user are private and in detailed respects idiosyncratic to him. The appropriate

notion for a public theory of meaning is a notion of congruence that is considerably weaker than this very strong sense of identity. If the viewpoint I am expressing is near the truth, the search for any hard and fast sense of identity of meaning is a mistake. It rests hidden away in the internal programming of each individual. What we are after are congruences that can collapse these private features across language users to provide a public and stable notion of meaning.

In fact, this way of looking at the matter is in a more general philosophical way very satisfying to me. I have come to be skeptical of the long philosophical tradition of looking for various kinds of bedrocks of certainty, whether in epistemology, logic, or physics. Just as the natural notion of a person is not grounded in any hard and definite realization, and certainly not a physical one because of the continual fluctuation of the molecules that compose the body of the person, so it is with the meaning of expressions. In terms of what I have just said about an ultimately psychological theory of meaning at the deepest level, I also disagree with Frege's attempt to separate in a sharp and absolute fashion logic from psychology.

From a formal standpoint, my work on the semantics of natural language has recently taken a more radical turn. I now believe that the semantics of a significant fragment of ordinary language is most naturally worked out in a framework that is an extension of relational algebras as developed earlier by Tarski, McKinsey, and others. Moreover, the notation for the semantics is variable free, using only constants and operations on constants, for example, taking the converse of a relation, the image of a set under a relation, etc. In a recent paper (1976c) I work out the details of such an approach to the standard quantifier words, whether in subject or object position. Such a view runs against the tide of looking upon quantification theory in first-order logic as one of the prime logical features of natural language. But as has been known implicitly since the time of Aristotle, much natural language can be expressed within Boolean algebra and it is not a large step from Boolean algebras to relational algebras of various sorts. One of the points of my 1976 paper is to prove that if we use the standard linguistic parsing that makes quantifiers part of noun phrases—so that we treat *all men* in the sentence *All men are mortal* as being a simple noun phrase and have a tree structure that reflects this—then it is not possible under obvious and natural conditions to have a Boolean semantics of the sort that has been familiar for a hundred years for such utterances. The previous history of Boolean semantics did not emphasize that the syntax and semantics had to go together. The proof of the theorem depends upon this intimate marriage of model-theoretic semantics and context-free grammars. The line of extension to quantifiers in object position brings in relational algebras in an obvious way, but the elimination of any quantifier notation in the underlying semantical notation is based on the same concept as in the case of quantifiers in subject position.

This same framework of ideas is developed in considerable detail in a paper on attributive adjectives, possessives, and intensifying adverbs, by Macken and me (1978f). The full set of subtleties to be found in the ordinary use of adjectives, possessives, and adverbs is beyond the competence of any theory to handle completely at the present time. I do think we make a reasonable case for the kind of model-theoretic semantics without variables that I have described as providing a more detailed and intuitive semantical analysis than any of the theoretical approaches previously published.

In a paper I am just now finishing I go on to consider logical inference in English. To my surprise I have been able to find practically no papers dealing with such inferences in a direct fashion. The reason for the absence, I suppose, is the slavish adherence to first-order logic in too much of the tradition of semantical analysis of natural language. A semantics that fits more hand in glove with the syntax of the language is required to generate the proper feeling for rules of inference, even though, as a technical tour de force, translation back and forth into a formal language is possible.

My own program of research in the philosophy of language is firmly laid out in broad outline, if not in all details. I want to understand in the same way I believe I now understand quantifier words the many other common function words in English. I have already spent some time on the definite article and the possessive preposition *of* and I would like to do the same for the other high-frequency prepositions like *to*, *in*, *for*, *with*, *on*, *at*, *by*, and *from*—I have listed these prepositions in their order of frequency of occurrence in a large corpus collected by Kucera and Francis (1967). The frequency of these simple prepositions is among the highest of any words in English but their semantical theory is as yet in very unsatisfactory state. Detailed analysis of their semantics is undoubtedly too empirical a problem for many philosophers deeply interested in language. I do not know whose work it is supposed to be—perhaps the empirical flavor of it seems antithetical to what philosophy should be like—but my own empirical bent in philosophy is nowhere more clearly reflected than in my attitude toward the philosophy of language. I do not think it is the only task, but for me it is a primary task, to provide a formal analysis that is faithful to at least the main function words in English when used in a serious and systematic way—I would of course like to have a theory for all possible uses but that seems out of reach at the present time. Frege himself had little to say about such matters and seemed rather suspicious of natural language as a vehicle for communicating exact thoughts. This same Fregean attitude has continued to hold an important place in the philosophy of language, but it should be apparent that I consider this aspect of Fregean philosophy a clear and definite mistake. I can think of no more appropriate task for philosophers of language than to reach for an exact and complete understanding of the prepositions I just mentioned, quantifier words, the tenses of simple verbs, and the like. As long as there is one definite intuitive usage that remains

semantically unanalyzed, we have not completed the main task of any semantical theory of language. I do want to emphasize how complex and subtle I consider this task to be. I am sure it will continue to be an active topic in philosophy a hundred years from now.

Education and Computers

In the section on mathematical concept formation in children I mentioned the beginning of my interests in education in 1956 when my oldest child, Patricia, entered kindergarten. I cited there the work in primary-school geometry. An effort, also noted but briefly, that was much more sustained on my part was work in the basic elementary-school mathematics curriculum. This occupied a fair portion of my time between about 1956 and the middle of the sixties and led to publication of a basic elementary-school mathematics textbook series, *Sets and Numbers*, which was one of the more radical of the 'new math' efforts. Unlike many of my colleagues in mathematics and science who became interested in school curriculum after Sputnik, I had a genuine interest in the psychological and empirical aspects of learning and a traditional interest in knowing what had been done before.

When I began working on the foundations of physics after graduate school, I was shocked at the absence of what I would call traditional scholarship in the papers of philosophers like Reichenbach that I read, or even more of physicists who turned to philosophical matters such as Bridgman and Campbell. There was little or no effort to know anything about the previous serious work in the field. I found this same attitude to be true of my colleagues from the sciences who became interested in education. They had no desire to know anything about prior scholarship in education.

I found I had a real taste for the concrete kinds of questions that arise in organizing a large-scale curriculum activity. I shall not attempt to list all the aspects of this work here, but since, beginning in the mid-fifties, I have written a large number of research papers concerned with how students learn elementary mathematics and I have had a fairly large number of students from education or psychology write dissertations in this area. Most of the work in the last decade or so has been within the context of computer-assisted instruction, to which I now turn.

Computer-assisted Instruction

In the fall of 1962, on the basis of conversations with Lloyd Morrisett, Richard Atkinson and I submitted a proposal to the Carnegie Corporation of New York for the construction of a computer-based laboratory dedicated to the investigation of learning and teaching. The proposal was funded in January 1963 and the laboratory began operation in the latter part of that year as computing equipment that was ordered earlier in the year arrived and was installed. The laboratory was initially under the direction of an

executive committee consisting of Atkinson, Estes, and me. In addition, John McCarthy of the Department of Computer Science at Stanford played an important role in the design and activation of the laboratory. In fact, the first computer facilities were shared with McCarthy and his group.

From a research standpoint, one of my own strong motivations for becoming involved in computer-assisted instruction was the opportunity it presented of studying subject-matter learning in the schools under conditions approximating those that we ordinarily expect in a psychological laboratory. The history of the first five years of this effort, through 1968, has been described in great detail—probably too much detail for most readers—in two books (1968a, 1972a) and in a large number of articles. I shall restrict myself here to a few general comments.

To some extent those initial hopes have been realized of obtaining school-learning data of the sort one expects to get in the laboratory. Massive analyses of data on elementary-school mathematics have been presented in my own publications, including the two books listed above, and a comparable body of publications has issued from the work of Atkinson and his colleagues on initial reading. My own experience has been that even a subject as relatively simple as elementary-school mathematics is of unbounded complexity in terms of understanding the underlying psychological theory of learning and performance. Over the past several years I have found myself moving away from the kind of framework that is provided by stimulus sampling theory and that has been so attractive to me for so many years. The new ideas are more cognitive in character and organized around the concept of procedures or programs as exemplified, for instance, in a simple register machine, that is, a simple idealized computer with a certain number of registers and a small, fixed number of instructions (1973c). I think that the ideas of stimulus sampling theory still have importance in terms of learning, even in the context of such procedures or programs, but certainly there is a shift in conceptual interest characteristic not only of my own work but also of that of a great many psychologists originally devoted to learning.

One of my initial interests in computer-assisted instruction was the teaching of logic at the elementary-school level and subsequently at the college level. Once complexity of this level is reached, psychological theory is in a more difficult spot in terms of providing appropriate conceptual tools for the analysis of student behavior. Currently my work in computer-assisted instruction is almost entirely devoted to university-level courses, and we are struggling to understand how to analyze data from the sorts of proofs or logical derivations students give in the first logic course or in the course in axiomatic set theory that follows it.

Although there are many questions about the psychology of learning and performance in elementary-school mathematics that I do not understand, still I feel that I have a relatively deep conceptual grasp of what is going on and how to think about what students do in acquiring elementary mathematical skills. This is not at all the case for skills of logical inference or mathematical inference, as exemplified in the two college-level courses I have mentioned. We are still floundering about for the right psychological framework in which to investigate the complete behavior of students in these computer-based courses.

There are other psychological and educational aspects of the work in computer-assisted instruction that have attracted a good deal of my attention and that I think are worth mentioning. Perhaps the most important is the extent to which I have been drawn into the problems of evaluation of student performance. I have ended up, in association with my colleagues, in trying to conceive and test a number of different models of evaluation, especially for the evaluation of performance in the basic skills of mathematics and reading in the elementary school. Again I will not try to survey the various papers we have published except to mention the work that I think is probably intellectually the most interesting and which is at the present time best reported in Suppes, Fletcher, and Zanotti(1976f), in which we introduce the concept of a student trajectory. The first point of the model is to derive from qualitative assumptions a differential equation for the motion of students through the course, initially the drill-and-practice supplementary work in elementary mathematics given at computer terminals. The constants of integration of the differential equation are individual constants of integration, varying for individual students. On the basis of the estimation of the constants of integration we have been able to get remarkably good fits to individual trajectories through the curriculum. (A trajectory is a function of time, and the value of the function is grade placement in the course at a given time.) The development of these ideas has taken me back to ways of thinking about evaluation that are close to my earlier work in the foundations of physics.

Research on computer-assisted instruction has also provided the framework within which the large-scale empirical work on first-language learning in children has taken place. Without the sophisticated computer facilities available to me at Stanford it would not have been possible to pursue these matters in such detail and on such a scale. Even more essentially, the presence of a sophisticated computer system in the Institute for Mathematical Studies in the Social Sciences has led to the computer-based approach to the problems of language learning and performance mentioned earlier. One of our objectives for the future is to have a much more natural interaction between student and computer program in the computer-based courses we are concerned with. Out of these efforts I believe we shall also come to a deeper understanding of not only how

computer programs can best handle language but also how we do, in fact, handle it. (Part of this search for naturalness has led to intensive study of prosodic features of spoken speech and how to reproduce them in computer hardware and software.)

I have not yet conveyed in any vivid sense the variety of conceptual and technical problems of computer-assisted instruction that I have tried to deal with in collaboration with my colleagues since 1963. This is not the place to undertake a systematic review of these problems, most of which have been dealt with extensively in other publications. I do, however, want to convey the view that the best work is yet to be done and will require solution of formidable intellectual problems. The central task is one well described by Socrates long ago in Plato's dialogue *Phaedrus*. Toward the end of this dialogue, Socrates emphasizes that the written word is but a pale image of the spoken; the highest form of intellectual discourse is to be found neither in written works or prepared speeches but in the give and take of spoken arguments that are based on knowledge of the truth. Until we have been able to reach the standard set by Socrates, we will not have solved the deepest problems in the instructional use of computers. How far we shall be able to go in having computer programs and accompanying hardware that permit free and easy spoken interaction between the learner and the instructional program is not possible to forecast with any reasonable confidence, for we are too far from yet having solved simple problems of language recognition and understanding.

At the present time we are only able to teach well skills of mathematics and language, but much can be done, and it is my conviction that unless we tackle the problems we can currently handle we will not move on to deeper solutions in the future. Because I am able to teach all my own undergraduate courses in a thoroughly computer-based environment, I now have, at the time of writing this essay, the largest teaching load, in terms of number of courses, of any faculty member at Stanford. During each term I offer ordinarily two undergraduate courses, one in logic and one in axiomatic set theory, both of which are wholly taught at computer terminals. In addition, I offer either one or two graduate seminars. As I have argued elsewhere on several occasions, I foresee that computer technology will be one of the few means by which we can continue to offer highly technical and specialized courses that ordinarily draw low enrollment, because of the budgetary pressures that exist at all American universities and that will continue unremittingly throughout the remainder of this century. Before I am done I hope to add other computer-based courses in relatively specialized areas, such as the foundations of probability and the foundations of measurement. The enrollment in one of these courses will ordinarily consist of no more than five students. I shall be able to offer them only because I can offer them simultaneously. My vision for the teaching of philosophy is that we should use the new technology of computers to return to the standard of dialogue and

intimate discourse that has such a long and honored tradition in philosophy. Using the technology appropriately for prior preparation, students should come to seminars ready to talk and argue. Lectures should become as passé as the recitation methods of earlier times already have.

In 1967, when computer-assisted instruction was still a very new educational technology, I organized with Richard Atkinson and others a small company, Computer Curriculum Corporation, to produce courses in the basic skills that are the main focus of elementary-school teaching. In retrospect it is now quite clear that we were ahead of our times and were quite lucky to survive the first five or six years. Since about 1973 the company has prospered, and I have enjoyed very much my part in that development. I find that the kind of carefully thought out and tough decisions required to keep a small business going suits my temperament well.

I have not worked in education as a philosopher. I have published only one paper in the philosophy of education and read a second one, as yet unpublished, on the aims of education, at a bicentennial symposium. Until recently I do not think I have had any interesting ideas about the philosophy of education but I am beginning to think about these matters more intensely and expect to have more to say in the future.

Philosophy and Science

From the standpoint of research I think of myself primarily as a philosopher of science, but to a degree that I think is unusual among professional philosophers I have had over the period of my career strong scientific interests. Much of this scientific activity could not in fact be justified as being of any direct philosophical interest. But I think the influence of this scientific work on my philosophy has been of immeasurable value. I sometimes like to describe this influence in a self-praising way by claiming that I am the only genuinely empirical philosopher I know. It is surprising how little concern for scientific details is to be found in the great empirical tradition in philosophy. It has become a point with me to cite scientific data and not just scientific theories whenever it seems pertinent. I recently made an effort to find any instances in which John Dewey cited particular scientific data or theories in his voluminous writings. About the only place that I found anything of even a partially detailed character was in the early psychology textbook written in the 19th century. When it comes to data, almost the same can be said of Bertrand Russell. It is especially the case that data from the social and behavioral sciences are seldom used in any form by philosophers. In my monograph on causality (1970a) I deliberately introduced detailed data from psychological experiments to illustrate some subtle points about causality. In a recent paper on distributive justice (1977a) I went so far as to calculate Gini coefficients for the distribution of salaries at the various professorial ranks at Stanford and several other universities.

Set-theoretical Methods.

One of the positions in the philosophy of science for which I am known is my attitude toward formalization. In various papers I have baptized this attitude with the slogan "to axiomatize a scientific theory is to define a set-theoretical predicate." A large number of my papers have used such methods, and I continue to consider them important. I should make clear that I am under no illusion that in any sense this method originated with me. My distinctive contribution has been to push for these methods in dealing with empirical scientific theories; the methods themselves have been widely used and developed in pure mathematics in this century. To a large extent, my arguments for set-theoretical methods are meant to be a constructive criticism of the philosophical tendency to restrict formal methods of analysis to what can be done conveniently within first-order logic.

I do not think of set-theoretical methods as providing any absolute kind of clarity or certainty of results independent of this particular point in the history of such matters. They constitute a powerful instrument that permits us to communicate in a reasonably objective way the structure of important and complicated theories. In a broad spirit they represent nothing really new; the axiomatic viewpoint that underlies them was developed to a sophisticated degree in Hellenistic times. Explicit use of such methods provides a satisfactory analysis of many questions that were in the past left vaguer than they need to be. A good example would be their use in the theory of measurement to establish appropriate isomorphic relations between qualitative empirical structures and numerical structures.

The many recent results in the foundations of set theory showing the independence of the continuum hypothesis and related assertions are reminiscent of what happened in geometry with the proof of the independence of the parallel postulate. But, as Kreisel has repeatedly urged, the parallel postulate is independent in second-order formulations of geometry having a strong continuity axiom, whereas the continuum hypothesis is not independent in second-order formulations of set theory. The great variety of recent results in the foundations of set theory have not really affected the usefulness of set-theoretical methods in the analysis of problems in the philosophy of science, and I am certain such methods will continue to be valuable for many years to come.

At one time I might have been upset by the prospect of moving away from set-theoretical methods to other approaches, for example, the kind of deeply computational viewpoint characteristic of contemporary computer science, but now I see such developments as inevitable and indeed as healthy signs of change. It seems likely that the theory of computation will be much more fundamental to psychology, for example, than any development of set-theoretical methods.

Schematic Character of Knowledge

In 1974, I gave the Hågerström lectures in Uppsala, Sweden, entitled *Probabilistic Metaphysics* (1974a). In those lectures I took as my starting point Kant's criticism of the old theology; my purpose was to criticize various basic tenets of what I termed the new theology. The five most important are these:

1. The future is uniquely determined by the past.
2. Every event has its sufficient determinate cause.
3. Knowledge must be grounded in certainty.
4. Scientific knowledge can in principle be made complete.
5. The grounds of rational belief and action can be made complete.

It is not appropriate here to develop in detail my arguments against determinism, certainty, and completeness, but it is my conviction that an important function of contemporary philosophy is to understand and to formulate as a coherent world view the highly schematic character of modern science and the highly tentative character of the knowledge that is its aim. The tension created by a pluralistic attitude toward knowledge and skepticism about achieving certainty is not, in my judgment, removable. Explicit recognition of this tension is one aspect of recent historically oriented work in the philosophy of science that I like.

It seems evident (to me) that philosophy has no special methods separate from those of science and ordinary life and has no special approaches to problems of inquiry. What makes a problem philosophical is not some peculiar intrinsic feature of the problem but its place as a fundamental problem in a given discipline or in some cases the paradoxical qualities it focuses on and brings to the surface. I am sometimes thought of as a primarily formalist philosopher of science, but I want to stress that at least as much of my scientific activity has been spent on detailed data analysis as it has on the construction of formal theories. My attitudes toward induction and the foundations of statistics, for example, have been conditioned by the extensive work in applied statistics I have done as part of other research efforts in psychology and in education.

I pointed out earlier that I thought my work in the foundations of physics was not as significant as the work in psychology because of the absence of an original scientific component. It is one of my regrets that I have not been able to do more in physics, especially in terms of empirical data. I have together with my students pursued certain questions of data in physics with some persistence and great pleasure. If I had the time and energy to write my own ideal book on the philosophical foundations of quantum mechanics, it would present a rigorous and detailed analysis of the relevant data as well as of the theory.

I was especially pleased to receive in 1972 the Distinguished Scientific Award of the American Psychological Association in recognition of my activities as a psychologist. This dual role of philosopher and scientist would not suit everyone's taste but in my own case it has been a happy one for the vitality of my intellectual life.

III. Personal Reflections

My entire academic career has been spent at Stanford, so I have divided this section into three periods: the first five years at Stanford, the next ten, and the last twelve. What I intend is to make some more general and more personal remarks in this part of the essay and especially to comment, beyond the remarks made earlier, on some of the people who have had an influence on me.

I came to Stanford in 1950 immediately upon receiving my PhD from Columbia and I have remained here without interruption. I have had the usual sabbaticals and I have traveled a great deal, perhaps more than most of my academic colleagues, but still I have remained relatively fixed at Stanford and undoubtedly will do so throughout the remainder of my career.

1950-1955

I have already commented on the influence that McKinsey and Tarski had on me during my first years at Stanford. From another direction, as I have already mentioned, I was strongly influenced by working with David Blackwell and M. A. Girshick on their book, *Theory of Games and Statistical Decisions* (1954). Working with them I learned a lot about both mathematical statistics and decision theory that was very useful to me later. I also began at this time working with Herman Rubin in the Department of Statistics, and I learned a great deal from Rubin, who has perhaps the quickest mathematical mind I have ever had the pleasure to interact with in any extended way. His error rate is reasonably high by ordinary standards, but the speed and depth of his reactions to a problem posed are incomparably good. During this period I also learned a great deal from my colleague in philosophy, Donald Davidson, and began with him the joint work in decision theory I mentioned earlier. In many ways we nicely complemented each other, for he comes at philosophical problems from a different direction and from a different background than I do, but our common agreement on matters of importance was more than sufficient to give us a good basis for collaboration.

I remember these early years at Stanford with great pleasure. I was working hard and intensely and absorbing a great deal about a great many different things. On the other hand, it is useful to say something about how slow one can be in taking in new ideas. I learned much from Blackwell and Girshick and also from Rubin about the foundations of statistics, but as I look back on the splendid

opportunity that was available to me in working with the three of them it does not seem to me that I got the kind of grip on the subject that I feel I have acquired since that time. I cannot help but feel that an opportunity was wasted and that a delay of years was imposed by my failure to reach a deeper understanding at that time. I think one of the difficulties with my earlier work is that I did not sufficiently appreciate the necessity of getting a strong intuitive or conceptual feeling for a subject. I probably tended to operate in too formal a manner, at least so it seems to me now in retrospect. All the same, some of my best papers were written in this period, and I do not want to sound overly negative about all that I had the opportunity to learn and do in those early years.

1956-1965

I have already mentioned the important influence of Estes during the year 1955-1956 at the Center for Advanced Study in the Behavioral Sciences. The continuation of the work in learning with applications to multiperson interactions and to mathematical concept formation in children was intellectually a major part of my life during the ten years ending in 1965. The work with Estes continued; we spent many summers together. We planned a monograph as the outgrowth of our work but for various reasons did not complete it. We did write the two long technical reports that were eventually published in shortened form as papers. But the extent of Estes's influence on my thinking during this period is underestimated by referring simply to the publication of two papers.

In the summer of 1957 there was a Social Science Research Council workshop, or rather collection of workshops, at Stanford. An outgrowth of the workshop on learning was the volume *Studies in Mathematical Learning Theory* (1959), edited by R. R. Bush and W. K. Estes, in which I published several papers, including the first paper with Estes. Perhaps the most important intellectual event for me that summer was the encounter with Duncan Luce and the famous 'red cover' report that later was published by Luce as his classical book *Individual Choice Behavior* (1959). He and I had great arguments about the exact interpretation of his axioms. I initially thought he had wrongly formulated his central choice axiom but he succeeded in persuading me otherwise, and out of those first encounters has grown a strong personal friendship and a large amount of collaborative work.

Although during this time I published two logic textbooks, *Introduction to Logic* (1957a) and *Axiomatic Set Theory* (1960a), in this ten-year period more than any other time in my career most of my effort was devoted to psychological research rather than to work in philosophy. I have already mentioned a good many of the individual psychologists I had the pleasure of working with in these years.

Another important influence on me was interaction with Dana Scott over several years, beginning with the period when he was an undergraduate at Berkeley and carrying through intermittently until he joined the faculty at Stanford some years later. Among mathematical logicians I have known, Scott is unusual in possessing a natural taste for philosophical problems and great interest in their analysis. We wrote only one paper together (1958b), but our conversations about a range of intellectual matters have extended over many years. Scott has the kind of clarity typical of logicians put at an early enough age in the Tarski mold. In some ways I am definitely less compulsive about clarity than I was in the days when I was working with McKinsey and later with Scott. Whatever one may say about the psychological healthiness of reducing compulsiveness of this kind, I am not at all sure it has been a good thing intellectually. It is perhaps an inevitable aspect of my widening intellectual interests since the late fifties.

In the last several years of this period, one of the strongest influences on my own work was the succession of able graduate students who wrote doctoral dissertations under my guidance and with whom I often collaborated. I mention (in chronological order) especially Helena Kraemer, Jean Donio, Barry Arnold, M. Frank Norman, and Paul Holland, all of whom took degrees with me in the Department of Statistics at Stanford and all of whom were concerned with mathematical or statistical problems in the foundations of learning theory. (Since 1960 I have had a joint appointment in the Departments of Philosophy and Statistics.) During the same period Michael Levine worked on a dissertation in psychology, which he completed in a formal sense a year or two later. Both Michael Levine and Frank Norman were as graduate students great sticklers for mathematical precision and correctness of formulation of theorems and proofs. I remember well the pleasure they took in correcting any mistakes I made in my graduate course on mathematical learning theory.

By the end of this period my attention was moving to the kinds of psychological questions, many of them applied, that arose in connection with computer-assisted instruction, and, on the other hand, I began to return to a more intense consideration of purely methodological problems in the philosophy of science. This does not mean that my interest in psychological research ended but rather that the 'learning theory' period running from 1955 to 1963 was reaching a natural end.

At the end of this period Duncan Luce and I undertook to write a long chapter on preference, utility, and subjective probability for Volume III of the *Handbook of Mathematical Psychology*. Writing this long article introduced me to Luce's awesome habits of work. Very few people I know are able to meet deadlines for completing a piece of work on time; practically no one is able to complete an agreed-to assignment in advance of the deadline. Luce is one of the

few that can; but it is not simply the meeting of the deadline that is impressive, it is his clear and relentless pursuit of the details of a particular theory or argument. I learned a great deal from him in writing this long survey article and have continued to do so. As the impact of his book *Individual Choice Behavior* has shown, he has a superb gift for simple formulation of quite general concepts and laws of behavior.

One important event for my own work and life that took place during this period was the founding, together with Kenneth Arrow, of the Institute for Mathematical Studies in the Social Sciences at Stanford. Stanford was then administered in a sufficiently informal way that it is not easy to peg the exact date on which the Institute was formed. It was a natural outgrowth of the Applied Mathematics and Statistics Laboratory, which had been put together in the late forties by Albert H. Bowker, now Chancellor at the University of California at Berkeley, to provide an organizational framework and an intellectual home for a wide range of work in applied mathematics and statistics. My own research began there in the summers, starting with the apprenticeship to Blackwell and Girschick mentioned earlier. The forming of the Institute followed in the late fifties. I have continued to direct the institute since 1959, and it has been a pleasant and constructive home for most of my research efforts since that date.

It was also during this period that I had my one serious flirtation with university administration. I was a half-time Associate Dean of the School of Humanities and Sciences at Stanford for three years and during the last term, the fall of 1961, Acting Dean. During this period I had several opportunities to assume full-time administrative positions at Stanford and some offers to do so elsewhere. I enjoyed the administrative work during this period and think that I have some flair for it, but certainly one of the wisest decisions I have personally ever made was to move away from administration and back into a regular position of teaching and research.

During this period I would probably have left Stanford except for the continued support of Bowker, first when he was Chairman of the Department of Statistics and Director of the Applied Mathematics and Statistics Laboratory, and later when he was Graduate Dean. He more than anybody else was responsible for creating for me an intellectual atmosphere at Stanford and a context for constructive research that have been so attractive that for many years I have not thought seriously about leaving.

1966-1978

To continue this theme of administration as I move into the final period of these personal reflections, I have found that the large-scale computer activities on which we embarked in 1963 have turned out to be a sizable administrative problem in their own right. At the peak of our activities in 1967-1968 we had almost 200 persons, including staff, research associates and graduate students, involved in the

Institute. The activity is much smaller now and I am thankful for that, but it continues to be a relatively complex affair and demands, as it has demanded since 1963, a fair share of my time. I do not regret the time spent, for my temperament is such that I would be restless in a purely sedentary life of paper-and-pencil research. The complex problems of running a computer operation on the frontiers of the technology currently available have provided just the right kind of stimulation to keep me contented and not inclined to seek administrative outlets of a more extensive nature, except those already mentioned at Computer Curriculum Corporation.

The efforts in computer-assisted instruction during these last ten years have been in association with a very large number of people, too many to mention here. My numerous joint publications in this area provide partial evidence of the extent of this collaboration, but I should also mention the many extremely able young programmers and engineers with whom I have worked and who have contributed so much to our efforts. At first I did not rightly appreciate the important role that able technical people can play and indeed must play in any successful research effort involving complex hardware and software. I had in the past heard such stories from my physics friends, but this was my first opportunity to learn the lesson firsthand. It is humbling to direct a complex activity in which you know that you yourself are not competent in most of the aspects of the operation. Large-scale research and development work has this character; I am not entirely happy with it. Increasingly in parts of my own research I have enjoyed working alone, but this is certainly not true of the work in computer-assisted instruction as it is not true of my experimental work in psychology.

Also, there are certain kinds of extended efforts that I would simply not be capable of carrying through on my own. I have in mind especially the effort I have engaged in jointly with David Krantz, Duncan Luce, and Amos Tversky in the writing of our two-volume treatise, *Foundations of Measurement*. This has been my longest collaborative effort, and I am pleased to say that we are still all speaking to each other.

Friendships that extend over many years have not been a topic I have emphasized in this essay, but a number have been important to me. One sad fact is that after staying at Stanford so many years I find that all of the persons with whom I formed relatively close personal ties in the 1950s have now departed. This includes Albert Bowker, now at Berkeley, Donald Davidson, now at the University of Chicago, William Estes, now at Rockefeller University, and Richard Atkinson, currently Director of the National Science Foundation. There are, of course, a number of individuals on the campus, especially colleagues in Philosophy and in the Institute for Mathematical Studies in the Social Sciences, that I work with and enjoy interacting with, but most of them are a good deal younger and are not friends of many years standing. Jaakko Hintikka has been my part-time colleague at

Stanford for many years. We have edited several books together and given a number of joint seminars that have been rewarding and pleasurable. On several occasions, Julius Moravcsik also participated in these seminars and I have benefited from my discussions with him about the philosophy of language. A third colleague in philosophy who has been at Stanford for many years but increasingly on a full-time basis is Georg Kreisel, and in the last few years we have developed the habit of talking extensively with each other about a wide range of topics. I value the regular round of talks with Kreisel and look forward to their continuing in the years ahead. Although Kreisel primarily works in the foundations of mathematics, he has a long-standing interest in the foundations of physics. He is especially good at giving broad general criticisms of the first drafts of manuscripts on almost any subject about which I write, and I try to do the same for him.

I close with some brief remarks about my personal life. I was married to my first wife, Joanne Farmer, in 1946, and together we had three children, John, Deborah, and Patricia, who at the beginning of 1978 are 17, 20, and 26 years of age; Joanne and I were divorced in 1969. In 1970, I married Joan Sieber and we were divorced in 1973. Since 1956 I have lived on the Stanford campus, and since 1961 in the house that I now occupy. It is a large, comfortable house, built in 1924, and is located in the old residential section, San Juan Hill, on a spacious lot. I feel completely anchored to Stanford and the Bay Area. It is unlikely that I will ever move anywhere else. All in all, I feel fortunate to have had the kind of life I have had for the past twenty-seven years at Stanford. Another twenty of the same sort is more than is reasonable to ask for, but I plan to enjoy as many as I can.

Stanford,
March, 1978.

Note

1. There is a certain amount of overlap in the content of this self-profile and an autobiography I wrote earlier (1978a), focused on my interest in psychology. I am indebted to Georg Kreisel for a number of useful criticisms. I also wish to thank Lofti Zodeh for the photograph which forms the frontispiece of this volume.

References

- Blackwell, D., and Girshick, M. A. (1954) *Theory of Games and Statistical Decisions*, Wiley, New York.
- Brown, R. (1973) *A First Language*, Harvard University Press, Cambridge, Mass.

- Bush, R. R., and Estes W. K. (Eds.) (1959) *Studies in Mathematical Learning Theory*, Stanford University Press, Stanford, Calif.
- Estes, W. K. (1950) 'Toward a Statistical Theory of Learning', *Psychological Review* 57, 94-107.
- Hamburger, H., and Wexler, K. (1975) 'A Mathematical Theory of Learning Transformational Grammars', *Journal of Mathematical Psychology* 12, 137-177.
- Hull, C. L. (1943) *Principles of Behavior*, Appleton-Century-Crofts, New York.
- Hull, C. L., Hovland, C.I., Ross, R.T., Hall, M., Perkins, D.T., and Fitch, F.B. (1940) *Mathematico-deductive Theory of Rote Learning*, Yale University Press, New Haven, Conn.
- Kahneman, D. and Tversky, A. (1972) 'Subjective Probability: A Judgment of Representativeness', *Cognitive Psychology* 3, 430-454.
- Knuth, D.E. (1968) 'Semantics of Context-free Languages', *Mathematical Systems Theory* 2, 127-131.
- Koopman, B.O. (1940a) 'The Axioms and Algebra of Intuitive Probability', *Annals of Mathematics* 41, 269-292.
- Koopman, B.O. (1940b) 'The Bases of Probability', *Bulletin of the American Mathematical Society* 46, 763-774. Reprinted in *Subjective Probability* (ed. by H.E. Kyburg and H.E. Smokler). Wiley, New York, 1964.
- Kucera, H., and Francis, W. N. (1967) *Computational Analysis of Present-day American English*, Brown University Press, Providence, R. I.
- Luce, R.D. (1959) *Individual Choice Behavior: A Theoretical Analysis*, Wiley, New York.
- Mackey, G. W. (1963) *Mathematical Foundations of Quantum Mechanics*, Benjamin, New York.
- Noll, W. (1964) 'Euclidean Geometry and Minkowskian Chronometry', *American Mathematical Monthly* 71, 129-144.
- Robb, A. A. (1936) *Geometry of Space and Time*, Cambridge University Press, Cambridge.
- Savage, L. J. (1954) *The Foundations of Statistics*, Wiley, New York.

Sen. A. K. (1970) *Collective Choice and Social Welfare*, Holden Day, San Francisco.

Smith, R. L. , Jr. (1972) 'The Syntax and Semantics of ERICA', Technical Report 185, Institute for Mathematical Studies in the Social Sciences, Psychology and Education Series, Stanford University.

Tarski, A. (1935) 'Der Wahrheitsbegriff in den formalisierten Sprachen', *Studia Philosophica* **1** , 261-405.

Tversky, A., and Kahneman, D. (1971a) 'Availability; A Heuristic for Judging Frequency and Probability', *ORI Research Bulletin* **11** (6).

Tversky, A., and Kahneman, D. (1971b) 'Belief in the Law of Small Numbers', *Psychological Bulletin* **76** (2) 105-110.

Whitehead, A. N., and Russell, B. (1913) *Principia Mathematica* (Vol. 3), Cambridge University Press, Cambridge.

Woods, J., and Walton, D. (1977) 'Post hoc, ergo propter hoc', *Review of Metaphysics* **30** , 569-593.

Zeeman, E. C. (1964) 'Causality Implies the Lorentz Group', *Journal of Mathematical Physics* **5** , 490-493.

Intellectual Autobiography, 1979–2006

I ended Part I of my Intellectual Autobiography, written in 1979, with the year 1978. So I now take up the narrative from that point onward. I begin, as earlier, with a long section on research, followed by a section on personal reflections. Although my interests have changed somewhat, I follow the same main seven divisions I used earlier: foundations of physics; theory of measurement; decision theory, foundations of probability, and causality; foundations of psychology; philosophy of language; education and computers; and philosophy and science.

IV. Research 1979–2006

Foundations of Physics

Most of my work in this period was, first, in collaboration with Mario Zanotti, as a continuation of our earlier research begun in 1974; and second with Acacio de Barros, a Brazilian physicist.

Zanotti and I wrote, between 1980 and 1991, five papers on quantum mechanics, all focused, from one angle or another, on the quantum phenomena of entanglement, first brought to prominence by the physicist John Bell. We considered the Bell inequalities in all of these articles. I won't give a technical account here of those inequalities, but they are used to show that there is no proper classical theory of the experiments to which they are applied. A sketch of the argument goes like this. In classical physics and, in fact, almost all branches of science except quantum mechanics, when we find simultaneously correlated events, such as fever and headache, we search for a common cause. Einstein believed to the end that such common causes, even if not necessarily observable in any direct way, should be found for quantum phenomena as well. The search for such causes is the search for hidden variables in the language of this enterprise, when pursued in the quantum domain. The Bell inequalities give necessary and sufficient conditions for such observable correlated phenomena as the spin of particles, often represented in entanglement experiments by the polarization of photons, to have hidden variables as common causes of the correlations. The point of the inequalities is to give a condition, just in terms of the observable correlations, for there to be such hidden causes. It turns out, not initially clear that this is the case, that the inequalities are equivalent to requiring that the random variables or probabilistic events having the pairwise correlations must have a joint probability distribution of all four together. So Bell's requirement was just equivalent to a well-known probabilistic problem for a great variety of special cases: when do random variables having certain marginal distributions have a joint distribution? The subject is messy and full of tantalizing partial

results. Some of the nice earlier results, with no connection with quantum mechanics, were given by the distinguished French mathematician Maurice Frechet, whom I met in the 1960s in Paris, but, as is usually for me and other academics, who talk to a lot of different people, I cannot remember anything about our conversations. (John Bell I did meet once and we had a very good talk about the foundations of physics at the Stanford Faculty Club; Bell, like a lot of other scientists who are well-known for conceptual innovations, had a strong natural interest in philosophy.)

I won't enter into the details of what Zanotti and I did in these five papers, except to mention one general result that has applications elsewhere. This is our proof that the existence of common causes for phenomena of any kind is exactly equivalent to their having a joint probability distribution (Suppes and Zanotti, 1981). We proved it for finite random variables, and it was easily extended to continuous random variables by Paul Holland and P. R. Rosenbaum (1986). (I am proud to say that I was Paul's Ph.D. advisor in the Department of Statistics many years before.) One immediate application is that the search for latent variables, as common causes are often called in the social sciences, always exist for any social phenomena that have a joint probability distribution, which in such research is implicitly always assumed. This means that further structural conditions must be imposed to get a result on latent variables of any scientific interest.

Years later, in fact in 2000, Acacio de Barros and I derived a set of similar inequalities for what are called Greenberger, Holt and Zeilinger (GHZ) configurations of three entangled particles. The surprising result is that we found inequalities that looked just like the Bell ones when correlations between pairs of particles are replaced by moments for three particles, i.e., expectations of the products of the spins of three particles. It was a particular pleasure for me, now 78 in 2000, to have this article published in *Physical Review Letters*, the leading physics journal in the world in the opinion of many people.

As I have emphasized in several publications, entanglement is the great mystery of nonrelativistic quantum mechanics. As far as the current evidence shows, particles can be entangled, even though separated by large distances. This seems to be a revival of action at a distance, as occult as was gravitation in the seventeenth century. Remember, Newton was quite clear that he did not understand the cause of gravity, and it violated all the precepts of the Cartesian physics that was still being taught, even at Cambridge, in the second half of the seventeenth century.

I now examine carefully some possible consequences of these ideas. The primary criterion of adequacy of a probabilistic causal analysis is that the causal variable should render the simultaneous phenomenological data conditionally independent. The intuition back

of this idea is that the common cause of the phenomena should factor out the observed correlations. So we label this principle the *common-cause criterion*. Although satisfaction of it is not the end of the search for causes or probabilistic explanations, it does represent a significant and important milestone in any particular investigation

Much of the earlier discussion of hidden variables in quantum mechanics centered around the search for deterministic underlying processes, but for some time now the literature has also been concerned with the existence of probabilistic hidden variables. It is a striking and important fact that even probabilistic hidden variables do not exist when the observable random variables do not have a joint distribution.

The next systematic concept is that of *locality*. I mean by this what John Bell (1966) meant:

It is the requirement of locality, or more precisely that the result of a measurement on one system be unaffected by operations on a distant system with which it has interacted in the past, that creates the essential difficulty. ...The vital assumption is that the result B for particle 2 does not depend on the setting \mathbf{a} , of the magnet for particle 1, nor \mathbf{A} on \mathbf{b} .

Quantum entanglement, as is well known, violates both the principle of a common cause and the principle of locality. There are other consequences, even more directly challenging, to the edifice of much modern physics. My objective here is to describe these additional difficulties, without the necessary technical developments for a really clear and transparent exposition.

The more recent entanglement experiments demonstrate two critical features:

- (i) Entangled particles can be experimentally separated by distances of the order of 100 kilometers (Zeilinger, 2005).
- (ii) Changing the angle of the measurement apparatus for one particle can be done just before the other particle reaches its measurement apparatus. So the causal effect of this change is transmitted, in recent experiments, at least five times faster than the speed of light.

These experimental results seem to present us with the following stark alternatives:

- A1. Reject the special theory of relativity and return to the Galilean invariance of Newtonian mechanics, where even instantaneous action at a distance is possible, and any finite velocity creates no problem.

A2. Hold fast to the special theory of relativity and accept macroscopic backward causation.

Backward causation follows from (i), (ii) and the special theory of relativity. With speeds faster than that of light, particles separated by a space-like interval can causally interact. In that case there is an inertial frame F_2 of reference, related by a Lorentz transformation to the initial frame F_1 , such that in this frame F_2 the causal effect of changing the measurement angle of one particle is an event A in the future of the event B of changing the other particle. Thus, backward causation. Event A occurs, relative to frame F_2 after event B , but A is the cause of B . I find this hard to accept. But perhaps it also helps explain why relativistic quantum mechanics seems still to lack a satisfactory foundation.

In 1991 Zanotti and I published three articles on entanglement. This was our last work on foundations of physics. In that same year I met Acacio, who joined us from Brazil to teach physics in the Education Program for Gifted Youth at Stanford. (More on EPGY later.) He and I also began talking to each other about a variety of issues in foundations of physics. Our first paper, published in 1994, was on a random walk approach to optical interference. This was the beginning of several papers on a semiclassical theory of optics that postulated photons, but with definite trajectories, and the associated electromagnetic field was now “reduced to” a probability distribution of photons. We published four more papers broadly in this area and entanglement, before the 2000 article on GHZ quantum entanglement of triples of particles, I will not go into the details, but remark on two aspects of this research that has been significant for me. It represents the first time I did any sustained research with a working physicist. Back in the early 1960s, when I wrote my first article on probability in quantum mechanics, I got some help from Sidney Drell, who became one of Stanford’s distinguished theoretical physicists. But it was different with Acacio. We worked day after day in a sustained effort to understand problems we found puzzling. Earlier coworkers, like Zanotti, had a background closer to mine from philosophy, foundations of mathematics and mathematical statistics. Acacio is a true-blue physicist and we often differed sharply in our natural approach to the problems we tackled, but this meant that the union of what we did know covered much more ground.

I end this part with reference to only two other relevant research efforts. In 1998 I published a paper on pragmatism in physics, to which I return later. The second effort was begun in 2000 with Acacio to apply the theory of oscillators to the representation of brain activities. This effort is continuing, but I will say something about this, beginning in 2000, under the general heading of Foundations of Psychology.

Theory of Measurement

The big event was the completion of the treatise *Foundations of Measurement*. Volume I was published in 1971. Eighteen years later, Volume II was published in 1989, and Volume III one year later in 1990. Dave Krantz had been the lead author in finishing Volume I; I took that responsibility for Volume II; and Duncan Luce did so for Volume III. (The four authors of the volumes were David Krantz, Duncan Luce, Amos Tversky and me. Unfortunately, Amos died early in 1996, just about at the apex of his brilliant scientific career.) Writing a treatise like the one on measurement is a mixture of many intellectual activities. Perhaps the most important and the most time consuming is organizing sets of results in a coherent fashion that brings out concepts and their relations in a way that was not as evident in their original development, or in the direction of their most recent application. This sounds very general and abstract. I will give a couple of examples from Volume II. Chapter 13, which I mainly wrote, presents in the first 38 pages a variety of classical geometric structures, such as absolute, affine and projective spaces, mainly according to their axiomatic foundations. To write something detailed, but not too lengthy, took much more effort than I expected, probably because I misjudged the depth of my knowledge at the beginning.

Another problem, more serious in a way, is that I initially naively intended to do something much more original. I wanted to emphasize geometric constructions, such as the familiar affine one of bisecting a line segment, to generate only finite geometries that have exact representations. By 'exact' I mean representations unique up to the relevant classical group, for instance, the projective or affine group of transformations. Well, this is a reasonable project, but one I had not thought through, and when I began thinking about hyperbolic and elliptic spaces I needed to do more work than I had planned on. My attachment to these finite geometries marks a philosophical attitude of mine toward the foundations of mathematics. Geometry has been neglected, following the arithmetization of analysis in the latter part of the nineteenth century. Moreover, the search for a foundation has been dominated by the search for a monolithic one that would work for all of mathematics. That ideal is still in place, and, of course, I fully recognize the value of it. But as a thoroughgoing pragmatist about science and philosophy, I have with experience (for which age is a surrogate) become ever more skeptical of the need for such results. Indeed, focus on them can be intellectually mistaken. A good, but different example was the long search in philosophy until recently for an absolutely certain grounding for epistemology. We can ring the changes on the great philosophers who sought such a foundation from Kant to Carnap. In applications of mathematics, including geometry, a pluralistic view is easy to defend. I mean by this that limited use of mathematics, in fact, mainly elementary use, is made in almost all applications in physics, biology, economics or elsewhere. Such limited use ordinarily can be given a respectable finitistic foundation if one is needed.

Let me be more concrete to illustrate this last point. In order to prove the standard representations in analytic geometry of the various kinds of structures I mentioned—an enterprise which has been famous since Descartes, it is necessary to add some kind of axiom of continuity, which is usually formulated in a strong form equivalent to a completeness axiom for the real numbers, such as: any nonempty set of real numbers that has an upper bound has a least upper bound. Such an axiom, not known in this form, or its geometric equivalents, in ancient Euclidean geometry, says something that is important to be clear about in pure mathematics. But it most surely expresses a principle that will not be tested directly in any physical or other scientific experiments. Having such strong smoothness conditions was no doubt important in nineteenth-century mathematical physics, when the aim was to find symbolic methods of computation to generate closed forms, i.e., finite expressions as solutions to problems, whenever possible. Already before the end of the twentieth century it was evident that the scene had changed. Most problems that arise naturally in physics now do not have closed-form solutions. The beautiful classroom examples of classical mechanics, such as Newton's solution to the isolated two-body problem, are rare in nature or the modern physics laboratory. Numerical work more or less dominates all the way down, and what is symbolic foundation can ordinarily be given a rather straightforward constructive foundation. These comments reflect my biases about measurement as well. I personally have always, and now ever more so, stressed the need to keep the subject closer to the laboratory and the practice of actual measurement, in the lab and outside as well.

So I will only briefly consider a more radical second example. The second part of Chapter 13 is devoted to classical space-time and that of special relativity. Standard strong assumptions about the smoothness of space-time are used. But it is evident enough they are in no sense subject to a detailed positive confirmation. It needs little argument to show that no experiments as we know them will refute the assumption of a sufficiently fine discrete space-time. In fact, high-energy particle physics, when pushed, would favor the latter. Experiments that try to work with ever finer volumes of space-time do not find smoothness, but explosive reactions—this sounds too vague, but it can be technically expanded on in a standard way. For details on quantum fluctuations in a vacuum, see (Milonni, 1993).

This finitistic laboratory orientation of my recent thinking about measurement is brought out in a paper I recently wrote (Suppes, 2006) that is about indistinguishability of observations or what is being observed, the subject of Chapter 16 of Volume II, which is entitled "Representations with Thresholds", also a subject dear to my heart. But in this recent paper, rather than taking the classical psychological stance, well documented experimentally, that indistinguishability within a threshold is a nontransitive relation, I considered measurements with finite scales. So, for example, if I have weights to measure accurately to a gram, the relation of

indistinguishability of the weight of objects falling between 2 and 3 grams is transitive, but this is all that can be said. We can make up an exact number for computational purposes, but the result is indeterminate, and should, strictly be represented by a pair of upper and lower measures, the upper measure being 3 grams, and the lower one, 2 grams. Do not, when you read this, make the mistake of saying, "Well, anybody can improve such a coarse measurement as this." The point is that, unlike the nineteenth-century fantasy that classical physical measurements could always be indefinitely improved, we now recognize that this is not the case. I can always find, at a given level of technological development of measuring apparatus, a lower bound beyond which improvements cannot go. And there are clear theoretical arguments I will not enter into here, as to why improvements will not converge in time to perfection. Making up an exact result is nice for computations, but should not be part of thinking about what the world is really like. Another point to emphasize. This indeterminate approximation precedes, rather than follows, statistical calculations, but this point also I shall not go into now in detail.

These remarks about indeterminate approximate measurements fit in well with the pragmatism I endorsed earlier. It is part of such a program to insist on plurality and partiality, as well as, indeterminacy of results. What I said, for instance, about weight, holds even more for subjective probability, or, if you wish, for propensity interpretations of probability as well.

Decision Theory, Foundations of Probability, and Causality

During the first ten years, 1979–1988, I published two books and a lot of articles under this heading. The list is too diverse and varied to summarize in a synoptic conceptual fashion. But certain themes are constant. In 1981 *Logique du Probable* was published in Paris. The text was a revision of a series of four lectures given at the College de France in November 1979. I enjoyed the extensive discussions that followed each lecture, and so, I think, did Jules Vuillemin, who, as Professor of Philosophy, at the College, sponsored my appearance. But I don't think the book was much read in France. With its emphasis on the philosophical relevance of many concepts and results from decision theory and the foundations of probability, it strayed quite far from the main concerns of French philosophy. I ended the book with a summary of eleven principal theses I had defended. Looking back now, my favorite two are these:

Intuitive judgments that are not able to be expressed verbally are an essential and uneliminable feature of our processes of decision making.

Determinism is dead; its short existence, from Laplace to Heisenberg, Born and Dirac, has ended.

No doubt too brief to be taken seriously, but meant to summarize a longer argument in the text. Jules was a good friend for many years. He was the most learned person I have known in the history of philosophy and science. He often corrected my historical errors on many points. I continue to miss him since his death in 2001.

Three years after the publication of *Logique du Probable*, in 1984 I published *Probabilistic Metaphysics*, based on the Hågerström Lectures, which I gave in Uppsala, Sweden in 1974, on the invitation of Stig Kanger and other members of the Department of Philosophy at Uppsala. Many themes of this longer book overlapped with those of the earlier one. The headings of the five chapters after the introductory one give a quick overview of the focus: randomness in nature, causality and randomness, uncertainty, incompleteness, and the plurality of science. The seventh chapter applies many of these ideas to the analysis of language, with an appendix on probabilistic grammars. The eighth and last chapter is on rationality and the orientation is close to that of the earlier book in French.

On one major point I have since changed my mind. In *Probabilistic Metaphysics* I argued strongly that deterministic theories of most natural phenomena must be false, and that determinism as a general view of the natural world is a false one. I have been persuaded by striking invariance results in ergodic theory established by my colleague Don Ornstein and others (see especially Ornstein and Weiss, 1991). They prove the impossibility of distinguishing between the correctness of a deterministic theory and a stochastic one of the classic billiard-ball example of motion, once only a finite accuracy of measurement, bounded away from zero, is assumed. In this case, two theories that are mathematically inconsistent cannot even in principle be distinguished by an unbounded but finite number of observations of the behavior of a billiard ball on a table with a convex obstacle in the center. To play off the themes of Kant's antinomies in the *Critique of Pure Reason*, especially the third one on causality, the revised view is that determinism is not false, but transcendental. I elaborated on this idea in several articles in the 1990s. The view I now advocate is that either universal determinism or universal indeterminism are transcendental, i.e., beyond experience. So a choice between them is a transcendental metaphysical one. Moreover, the invariance I mentioned is undoubtedly present in almost every variety of complex natural system from chaotic billiard balls to quivering *Aplysia* or your favorite bacteria. That transcendental determinism will not disappear as an attractive doctrine is supported by its presence even in theories of quantum mechanics, notably the Bohmian approach.

My interest in the foundations of probability continued during this period, but I will not enter into the technical details of various papers. I summarized my thoughts, and answered some questions I now think I understand, relatively recently in the long Chapter 5 of *Suppes* (2002).

Foundations of Psychology

I continued work on most of the psychological topics discussed in the earlier period prior to 1979, but I will only mention here two new areas of psychological research I had not contributed to or even really thought about before. One is the study of eye movements and the other is the study of electrical and magnetic activity in the brain.

Eye Movements

I cannot now remember how I first got interested in eye movements. But it was about 1980. Somewhat rashly I agreed to host an eye-movement laboratory run by Jim Anliker, which was supported by one of the government agencies. Jim was an experienced scientist, with a doctorate in psychology, who had devoted at least the previous ten years to eye-movement experiments. But he had published very few papers, and did not have much interest himself in building stochastic models of the obviously probabilistic behavior of eye movements. So it was a desirable collaboration, with each of us contributing, but also depending on the other. We did some experiments I still like in which we recorded the eye movements of students doing arithmetic exercises. The actual behavior was much more complex than the hallowed explanations of how students work such exercises and use the algorithms they were taught. Two detailed papers on this work were published in 1982 and 1983. I was then invited by Eileen Kowler, another specialist in eye movements, to write a review paper on eye-movement studies of both arithmetic and reading. This time no new experiments were conducted, but I assiduously studied the large literature on eye-movements in reading, which I previously knew only in a superficial way. I organized my ideas to formulate some stochastic models of eye movements in reading in this article (1990) and in a second one (1994) that grew out of my participating in a conference on autism in Sweden. A couple of years later, Julie Epelboim received a postdoctoral fellowship to study mathematical models in psychology with me. Besides Jim Anliker, she was the only experimental psychologist, well-trained in the specialized lore of eye-movement research, with whom I have worked. We published in 2001 a paper I still like a lot on eye movements in solving geometric problems with diagrams. The eye movements revealed directly the subtle interplay in the shifting back and forth between the texts of the problems and the accompanying diagrams. The model we developed was used to estimate the size of working memory in solving such problems.

Brain Experiments

I was invited to give a paper and participate in a conference at the University of Minnesota on high performance computer applications in the behavioral sciences in May of 1996. I don't really remember the details of my own paper, but I was fascinated by the lecture of Professor Sam Williamson. This decisive event led to research that

has occupied me since 1996 more than any other area of research and seems likely to continue to do so in the future.

Here is a record of what I wrote about that lecture and what I thought about in the days following it (May and June, 1996).

Saturday, 11 May 1996

A lecture on current work on Magnetic Imaging (MEG) by Sam Williamson, a physicist at NYU who has been one of the early workers in developing the use of SQUIDS for recording the magnetic field generated by electrical activity in the brain. I asked Sam various questions in a skeptical mode about what could and could not be done. I also continued to chat with Sam and got a better sense of the nature of the magnetic recordings. This is the first time that I had heard about any of the details of magnetic recording with SQUIDS (superconducting quantum interference devices).

22 May 1996

On the way to Europe, I stopped with Christine in New York and during the afternoon of May 22, I met Sam Williamson at NYU's Center for Neural Science.

4 June 1996

The following memo of record is directly transcribed from the writing of this note on 4 June 1996 on United Airlines Flight 955 from London to San Francisco.

Memo of Record

...The purpose of this memo is to record my meeting with Sam Williamson at NYU – Sam is a physicist at the NYU Center for Neural Science. We met from 2 PM to 3 PM on Wed., May 22 – Christine and I were on our way to Europe and we stopped overnight in New York. I called Sam in hope of meeting him after hearing a talk by him on MEG (magnetoencephalography) – the technique using SQUIDS for measuring magnetic fields arising from neuronal activity. The fields can be measured externally but close to the scalp.

In the course of our intense one-hour conversation, I suggested to Sam that it might be possible to recognize internal speech, initially just individual words such as “Yes” and “No”, by analyzing the magnetic wave produced by the neuronal activity by an approach similar to that used for standard speech

recognition. My conjecture is that the wave pattern of internal speech must, in a precise sense to be determined, be approximately isomorphic to the actual speech wave.

A main basis for the conjecture is that the time course and sense of prosodic emphasis in the conscious awareness of the internal speech seems, as would be expected, very similar to that of the corresponding actual speech.

Although the physical details are bound to be subtle and complex, there is much to support this conjecture. We produce too much speech of which we are consciously aware, and it is too easy to produce corresponding internal speech for the relation between the two to be impossibly complex. The subjective sense of similarity is of the greatest importance.

Let us suppose that the area of the cerebral cortex where internal speech takes place can be localized, because it is in fact a relatively local phenomenon. If, on the other hand, production of such speech involves a distributed network of neurons, then internal speech recognition will be much more difficult.

Assuming, then, localization of activity, the subject must, during training trials, tell us what was said silently on each trial. Put better, confirm that what was said was the word or phrase requested by the experimenter. A variety of machine-learning techniques can be applied to classify correctly what was silently said on the basis of the magnetic wave form recorded.

In principle, if the technique works for a very small number of words, it should be relatively straightforward to extend it to a large vocabulary, as is the case with current research on actual speech – in fact, pretty good commercial products exist. As in the case of actual speech, continuous internal speech will be much more difficult.

Of course, a different kind of machine learning may be required to learn the “meaning” of the magnetic waves. But given the abundance of current techniques, the difficulties should be not too great, granted it is technically possible to record meaningful magnetic waves.

If internal speech can be recognized, many further research questions immediately arise. Perhaps the first of great interest would be this. When we do an associative word search of memory, do the magnetic waves of the neuronal activity reflect a meaningful

use of internal speech. The conscious subjective impression is certainly affirmative, but is it so?

Here are some more questions briefly stated. Could the execution of simple arithmetic algorithms be similarly interpreted in a meaningful way from observation of the magnetic field when done purely mentally? For example, if I silently say "Nineteen plus fifteen is thirty-four"?

A large bundle of questions of a different sort concern internal conscious recognition of mental visual images and their manipulation or more detailed construction or enhancement. Internal speech recognition seems much the easier problem, but the massive work on visual pattern recognition should apply rather directly, as far as general technique of recognition is concerned.

If, like the work I did in the 70's on speech recognition over the telephone, we can succeed with "Yes", "No" and the ten digits, we will have taken the crucial first step. It remains to be seen if the current level of technology, and our understanding of how to use it, is up to the task.

Much has happened since 1996, and my understanding of the many complex aspects of signaling in the brain has vastly improved. The first important finding was that in the initial experiments at Scripps Institute in La Jolla with magnetoencephalography (MEG), we were fortunate that the standard electroencephalography (EEG) was also used by the placing of such sensors in the standard EEG 10-20 configuration. What we found to our surprise in our first analyses, as reported in the initial paper (Suppes, Han, and Lu 1997), was that we were able to recognize brain waves better by using the EEG results, rather than the 140 some MEG sensors. From the standpoint of continued pursuit of the research, this was a lucky break, because it is relatively easy and cheap to run EEG experiments. The technical details of MEG experiments are advanced lessons in modern superconducting devices: for example, SQUIDS must be kept at a temperature close to 0° Kelvin. In contrast, EEG research requires equipment that has not been drastically modified for more than fifty years. Moreover, it is widely used, not only in research, but in thousands, indeed hundreds of thousands of medical offices and clinics throughout the world.

So after these initial experiments, which had some success, I and my collaborators rapidly published five articles in the Proceedings of the National Academy of Sciences, of which I am a member, having been so elected in 1978. These articles marked the first burst of activity. Since then we have settled down to a more elaborate program of experimentation, with the computational analysis and modeling requiring vastly more time and effort than the experiments themselves.

This is not the place to describe the work in technical detail, which is properly reported in a series of short articles that have been appearing over the last few years, and also in Chapter 8 of Suppes (2002), but I would like to give a sense of the kind of prospects we see for the research we have embarked on. One way of describing the effort is that we are like cryptographers trying to understand how the brain encodes language and may use it internally. In support of such internal use is the surprising finding that when very simple visual images are shown to experimental subjects there is rather good evidence that within a few hundred milliseconds they will activate in the brain an image of the word corresponding to the concept. For example, if we show subjects a circle or a triangle, then this surprising finding has relevance to an old philosophical controversy. It is a controversy about how the brain or the mind represents abstract ideas, such as the general concept of a particular color or shape. I quote from our summary of this reference to the eighteenth-century dispute in Suppes, Han, Epelboim, and Lu (1999b, p. 14663).

Early in the 18th century, Bishop Berkeley (1710) famously criticized John Locke's Theory of abstract ideas (1690). David Hume (1739, p. 17) later summarized succinctly Berkeley's argument.

"A great philosopher [Berkeley] has disputed the received opinion in this particular, and has asserted, that all general ideas are nothing but particular ones, annexed to a certain term, which gives them a more extensive signification, and makes them recall upon occasion other individuals, which are similar to them."

Berkeley's views are well supported by our results. After visual display of a patch of red or of a circle, the image is represented in the cortex by the brain wave of the word *red* or *circle* within a few hundred ms of the display and somewhat quicker than is the representation in the cortex of the spoken word *red* or *circle*. To the skeptical response that we do not really know it is the word *red* or *circle* that is being represented in the cortex, as opposed to the particular visual image, we respond that everything we have learned thus far about the one-dimensional temporal representation of words, presented either auditorily or visually, supports our inference, the spatial unidimensionality of the temporal representation used for recognition, above all. Perhaps just as important, the filtered brain waves representing the spoken color or shape words conform closely to the brain waves of the many

other words whose brain waves we have identified in our earlier work.

A second slightly earlier finding, also surprising, was that we could use brain waves of one subset of subjects to recognize the brain waves of a disjoint subset of persons. This provides evidence of an important kind of invariance in the brain-wave representations of words between individuals (Suppes, Han, Epelboim, and Lu, 1999a). Having such an invariance is certainly not something that can be derived apriori, but it greatly facilitates the efficiency of the complex system of communication of passing messages or other linguistic expressions from one person to another. If my physical procedures for processing a given word or sentence are quite different from another person's, one can anticipate it will lead to more difficulties in fast decoding of what was said by each of us to the other. It is easy to expand on this theme. It is not my purpose here to do so.

More recently (2006), in the search for very high rates of recognition, that is, rates on the order of ninety percent for at least two or three hundred individual trials, with the null hypothesis being that the recognition rates would by pure chance be only fifty percent, we have moved from words to more general auditory and visual images. We have made a big effort to find stimuli that could easily be discriminated, that is, easily discriminated in terms of the brain waves generated by subjects listening or looking at the auditory or visual stimuli. Most of this work has not yet been published. It had taken us in a direction I had not really anticipated at the beginning. In fact, our best results, the very best results, not surprisingly really, are for the binary contrast between a highly selected auditory stimulus and a highly selected visual stimulus. The outstanding instance of this so far is the auditory word *Go* juxtaposed in trials to the visual stimulus of a bright red stop-sign. One of the impressive things about these experiments is that the rate of trials can be quite fast. We randomly display one of the two stimuli every second so that in twenty minutes we already have 1200 trials. For several subjects we have tested, we have recognition rates above ninety-five percent on test trials for which no parameters of any kind have been estimated. These are very high recognition rates under classical tests of the null hypothesis. The probability of the null hypothesis being true, i.e., the p -value of the null hypothesis, is an extremely small probability and the results are highly significant. For example, it is not uncommon to have $p < 10^{-50}$ and sometimes even $p < 10^{-100}$. These are, by ordinary psychological standards, extraordinary statistical results, which are highly desirable, in order to have a firm basis for further analysis.

I end this discussion with two additional examples. The first one is about comparing stimuli of the sort referred to. On one trial the subjects sees a bright blue circle, and on the next, with the same luminance, a bright red triangle: other trials use the other two combinations, blue triangle and red circle. In analysis of these data, we are able to discriminate, if we hold shape constant, the color, and

if we hold color constant, the shape. These findings support rather strongly the psychological experiments on attention of some years ago (Treisman and Gelade, 1980). But at the same time they exhibit some contextual effects that cannot be ignored. So that, for example, which color is being held constant can make a difference. The important point is that the concept of an invariant, so nicely formulated for physical quantities like length, is more complex when the invariance is for less idealized properties of concrete physical stimuli. Some contextual effects seem inevitable, and were the source of many of the difficulties in earlier psychological research, especially in the theory of perception. What is heartening, on the other hand, is the robustness of what can still be achieved without having a simple exact invariant.

The second example is also related to many kinds of psychological and philosophical arguments of the past. Are imagined brain images, when compared to ones resulting from the presentation of physical stimuli, as in our experiments, of the same sort or different? So, for example, if I take the red stop-sign mentioned earlier and if I now ask a subject to imagine a red stop-sign, will the brain image generated be similar to the brain image generated from looking at the red stop-sign? The answer is positive. We have been able to classify the stimulus-driven brain image using the imagined brain image even though, as one would expect, the recognition rates are not as good, but still significant statistically. Second, and perhaps more important, is the earlier question, can we classify and recognize such images at all and the answer to that is again positive. We do best, of course, when comparing one imagined image, for example, of the auditory stimulus *Go* against the imagined red stop-sign, that is, imagined image compared to imagined image of another sort. The recognition rates, again, are not as good as for the strong, and as Hume would say, vivid impressions of a stimulus, but still quite significant. (In this discussion, the use of the word *image* is not meant to imply that brain images have direct perceptual qualities.)

Well, it is easy for me to occupy too many pages with the discussion of our current brain research. The last examples would seem to suggest that we have given up on language, but this is not so. We think that there is still much that we can do and have hopes of being able to accomplish. One of the examples is showing a structural isomorphism between sentences and their generated brain waves. The mapping here is in terms of words to brain images of words, so that, for example, the word *Paris* in a sentence maps to a similar brain image when the word *Paris* occurs in different sentences and in different grammatical roles. It is easy to take this intuitive idea and give a formal definition of structural isomorphism for the correspondence between sentences and their brain images. Exactly the same problem arises for such a correspondence between words and their phonemic or syllabic structure, and the corresponding brain images. Here too we have some successes, but also problems to deepen our understanding. The work done so far is not yet published.

I should also mention one other approach, mainly theoretical, on which Acacio de Barros and I have spent considerable time, with as yet nothing published. In the recent work we have been joined by Gary Oas, head of physics instruction at EPGY. A great variety of theoretical and experimental research in biology and in physics shows that collections of neurons often behave as coupled electromagnetic oscillators that may be synchronized by an appropriate external stimulus, or even by attention to an internal computation or imagined brain image. We are now working on modeling these neural oscillators by weakly coupled phase oscillators satisfying Kuramoto's nonlinear equations (1984). Our initial efforts have focused on modeling with such oscillators the relatively simple but experimentally quite successful behavioral stochastic models of learning, initiated by a classic paper Estes (1950). Our current work specifically uses behavioral models and experiments from Estes (1959), Suppes and Atkinson (1960), and Suppes and Ginsburg (1963). We introduce sets of oscillators to represent sets of stimuli, other sets of oscillators to represent sets of responses, and similarly for reinforcements. Conceptually, we think of this modeling network of oscillators to be a special case of modeling more general associative networks. Some quite positive specific results are about to be published, but I shall not try at this point to describe what we have done in more detail.

It is easy to mention in closing many things we have not touched, nor has anyone else as yet successfully. Probably the most interesting example would be detailed recognition of the brain waves retrieving memories. For instance, retrieving words from long-term storage when listening to speech or reading, or even the more complicated case of producing speech, that is, of speaking.

I want to make a final philosophical remark of a general sort about our brain research. Neural experiments, as opposed to purely psychological ones on mental phenomena, are almost certain to encourage an empirical brand of nominalism. The reason is that the brain itself is not going to recognize, in any direct sense, an abstract idea or abstract concept. There must always be some physical representation. As the axiom goes in physics: all information and all computations are physical in nature. There is no nonphysical information and there is no nonphysical computation. So in order to plan anything, to think about anything, or to feel anything, concrete physical representations of brain activity are needed, just of the kind that nominalists such as Ockham, and in their own way, in the controversy mentioned earlier, Berkeley and Hume supported. The radical nature of this nominalism has not been adequately absorbed in modern philosophy of mind. Clever philosophers will find ways to work around it and to hold on to their traditional and cherished ideas. My own prediction is that as detailed work of the kind described, and ever better levels of precision and depth are achieved in the future, there will inevitably be a change. The seemingly permanent abstract

mentalese of philosophy and folk psychology will be challenged by a deep skepticism about the actual empirical validity of this traditional approach. It works only when not too many details are called for. This doesn't mean, on the other hand, that the language of folk psychology and of common sense will be abandoned. Just as we haven't abandoned the talk about ordinary objects and events like tables, chairs, and water running in the kitchen, because of all the many results of modern physics. We will retain folk psychology and ordinary mentalese, but that doesn't mean we believe they can express anything like the exact scientific account of what is going on, or that they are adequate for analyzing the structure and functions of the brain.

Philosophy of Language

During this period I have written a number of different papers in the philosophy of language, but I think I will restrict my remarks here to one sustained effort that I enjoyed a great deal and hope to continue, even though the last paper was written in 2001. This effort is a series of nine papers written on the learning of robotic natural language. Here the effort was not with learning on the part of humans, but how to teach robots to understand ordinary natural language, for example, such commands as *Put the bolt in the red box*, or *Get the screw*, *Put the nut on the screw*, *Get the large washer*, or more elaborately *Get the red screw behind the plate, which is left of the washer*.

The important feature of this work is that the robots' learning did not start absolutely from scratch. It was assumed they already had machine-language instructions for carrying out the robotic manipulations required. A demo was implemented on an "academic" but real robot by my former student and now successful entrepreneur in China, Dr. Lin Liang. What was important was learning how to interpret ordinary language in the several machine-language levels used in implementation. This is a nontrivial task. To show that we were not simply oriented toward English, and also because my two collaborators on most of this work were Li Liang who is Chinese and Michael Böttner who is German, we always did experiments in at least Chinese and German as well as English. Indeed, in one paper we studied the learning of interpretations of the machine language in ten different natural languages. What does learning an interpretation mean? Here is the answer. We had a set of formal axioms for learning that we implemented. As you might expect, they are associative in nature. Then we had axioms about changes of state in long-term memory using semantic categories and the semantically interpreted internal language already available to the robot. We also needed and had axioms for denotational learning and here we had denotational learning computations. Given the attention that has been devoted to denotation in the philosophy of language and the philosophy of mind, what is interesting here, and probably most distinctive, are the specific denotational learning axioms. Seldom, if

ever, in the ordinary philosophical enterprise have such denotational learning concepts been explicitly used and tested.

Finally we had various computations using the concept of congruence I introduced in my presidential address to the Pacific Division of the American Philosophical Association (1972), where I urged that a geometrical notion of meaning, using various strong and weak concepts of congruence, was the way to proceed without being obligated to adopt some single notion of synonymy. Such a notion of congruence is absolutely critical to semantic paraphrase and must play a role, in my judgment, in any extended development of robotic learning of natural language. All nine papers were written with the collaborators already mentioned. What I liked was the success we had in using classical concepts from psychology and from philosophy of language in developing the very specific theory embodied in the kind of axioms mentioned above.

Education and Computers

During the period running from 1979 to 1992, really all my work in this area was focused on the use of computers for instruction, what was called at that time *computer-assisted instruction* and now is described in various ways, for example, *computer-based instruction*. The important point is that there was much to be done, and I was able to round up resources to do a good deal of research on such teaching, not only in elementary mathematics where I had started back in the 1960s, but also in teaching of language, such as various levels of teaching of English from elementary-school to freshman courses in college, and especially the teaching of foreign languages. Much of the work over a considerable period was summarized in the book I edited in 1981 entitled *University-Level Computer-Assisted Instruction at Stanford: 1968–1980*. It provides an overview of what was done at the university level. In fact, this 930-page book is perhaps the most detailed analysis of such work published by anyone during the period when it began in the 1960s until the present time. Indeed, a book of this length will probably not be published in the future, because such extensive details will be reported on the web in electronic form.

Over the next few years after 1992, I published very little on the use of computers for educational purposes, but in 1992 a new effort began. This was the Education Program for Gifted Youth (EPGY) that I, with others, organized at Stanford. I have served as director since the beginning, and continue to do so today. The aim is to provide online courses for precollege students. We are not just focused on the last few years of secondary school but begin instruction in kindergarten in the case of mathematics. Building on my earlier work,—first a series of textbooks for elementary-school mathematics in the early 1960s, then corresponding computer-based courses of the same sort—, I revised the material extensively and quickly created a K-7 mathematics course for EPGY, which has

been, up until now, the course with by far the largest enrollment. I list among my publications the CDs produced for the various levels of that course until 2005. However, that CD-type of publication has now ceased, and the course is offered on the web through a browser. It is clear that future developments will be directly online on the internet.

I now have something that I particularly value that I did not have in the extensive earlier work. This is a centralized data base from all the sites on which the courses are being used, so we have, as a result, massive files that can be used to guide revisions of lectures and exercises, by analyzing student responses, and to test various psychological models of student learning and performance. The resulting publications are now principally being placed on the web site of EPGY. I look upon such electronic distribution as the primary method for this kind of publication in the future.

I have put a lot of effort into EPGY from 1992 to the present. We have extensive courses in mathematics, physics, and English for gifted students, many of whom begin as early as four-years old. But we also offer advanced courses in mathematics and physics for secondary-school students who are able first to complete a calculus course, either in their own local school or with EPGY, well before they graduate from high school, so that the last several years before college can be spent studying what are really online versions of undergraduate mathematics and physics courses at Stanford, and for which they get a Stanford University transcript.

During this period, up until 1990, I was also CEO of Computer Curriculum Corporation, a commercial company dedicated to offering computer-assisted instruction in schools. In 1990 the company was sold to Paramount Corporation, primarily known for movies but at that time also the owner of a large number of educational publishing houses. Returning actively in 1992 to the development of new computer-based courses within the framework of Stanford had academic aspects that I very much appreciate. Here I mean *academic* in the following sense. There is no real push to create something that must be profitable, and it is possible to spend much more time on research questions. Something that is perhaps even more important than anything aimed directly at published research is having the opportunity to explore uses of computers that would be considered impractical in a commercial setting. Let me just give two recent examples on which I have been working the last several years.

Beginning as long ago as the 1960s, I introduced elementary theorem proving on the computer to very bright elementary-school students in Palo Alto and other places. The proofs focused on are those for elementary theorems of arithmetic, which follow from the elementary axioms for the ordered field of rational numbers. In the beginning, only the very first parts of the natural axioms in this subject were used by the students. This was repeated in several

different forms over many years, each year increasing the number of students and often having the course used by older students. The course was then suspended for at least ten years, and I am pleased to have restarted it once again in a more elaborate form. So, we are now offering for elementary-school grades 4 to 7 a quite elaborate development of the theorems based on the rational operations of addition, multiplication, subtraction, and division, and including, for example, the definition of absolute value and the tricky theorems about absolute value and inequalities that are so useful later in proving the standard ϵ - δ theorems in analysis.

As part of this effort, we also offer what is I think one of the best features introduced earlier, but now done much more thoroughly. The students are given a great many exercises in which they are not told whether the exercise has a sentence that is to be proved by using theorems already proved, or is to be shown false by giving a counterexample. The point of the exercises is to make the students think conceptually and strategically whether a given formula can be proved or is a counterexample. If it is a counterexample, they must give specific numerical values to the variables to show that it is. Through such exercises, students learn some useful methods of problem analysis they might otherwise not. So as I write this, we are completing the new version of something begun more than forty years ago in its first version (1964). This will be the best and most complete version, and undoubtedly will have the most students using it. It will establish, in a better way than in the past, the ability of able students in elementary school to learn to give rigorous elementary mathematical proofs, with validity checked explicitly by appropriate computer programs.

The second example is connected with one of my most cherished ideas in the philosophy of language. I have already mentioned paraphrase. Both philosophers and psychologists have neglected this concept. If I hear a lecture, which I often do,—in fact one of the main pleasures and also problems of being in a place like Stanford is that there are too many lectures that I would like to hear each week, and now that I have a strong interest in neuroscience far too many—, so it is a real problem to choose what I want. In any case, when I go to one of these lectures and someone asks me afterward, “What was the lecture about?” I can, of course, not possibly repeat in serious literal detail what was said. I give a paraphrase that varies in the coarseness of its summary, and that does not use the exact words and phrases of the lecturer, but my own. This is the glory of paraphrase, one of the great syntactic and semantic features of human speech not nearly remarked upon enough.

What are the mechanisms of paraphrase, how do we do it? Well, there are many things to be said, but I will not try to go into the open research questions that I think should be of current interest. Rather, I will describe a direct application in our current EPGY teaching of elementary-school language arts, English as a second language, and

Chinese as a second language. For students taking language arts in American schools in English, or students learning English or Chinese as a second language, one problem is getting them to practice writing large numbers of sentences in English or Chinese, as the case may be, to enhance their command of the language. It is a fact of life, much regretted by many of us, that it is simply the case for a large number of reasons that students' writing of English in American schools is not evaluated very thoroughly and concretely by teachers. Certainly not as much as was often the case in the past. The same is true of foreign-language instruction. So the problem is to see how much we can do with sophisticated computer programs. The results, of course, are not going to be perfect. But the finer points of grammar are not well controlled by elementary-school teachers of language arts either. Our objective is to do at least as well and at the same time to give students extensive practice in writing. So, we have constructed a large sequence of limited environments, which we can manage rather well in terms of highly detailed analysis of what the student is doing. We present on the computer screen, in any of the kind of cases mentioned, a fixed vocabulary of between twenty to fifty words. Students can click and drag on these words to write sentences. Clicking and dragging is faster, particularly for young students, than typing. We have written very large and very specific context-free grammars for the given restricted vocabulary—which can be quite accurate because of the restriction.

We also do computations of paraphrase to check the semantic correctness of students' written answers to exercises. What students are ordinarily doing is answering questions or giving a brief description of some kind. We do not expect the student to give a unique answer. That would be inappropriate for any rich use of language. Moreover, in most cases the given vocabulary list requires students to have to write something very different from what they have read. So, for example, they must learn to use in proper anaphoric fashion the many kinds of pronominal and other anaphoric constructions important in English. The computation of a paraphrase being correct is the right semantic approach for this purpose. We have not, in any complete sense mastered the subject, but we are continually improving our computer programs and learning ourselves from the mistakes students make.

Because formal rules of paraphrase have not been widely discussed in the literature of mathematical and computational linguistics, in contrast to context-free grammars and the rules that can generate them, some brief comments at least are in order. The rules of paraphrase are, at bottom, semantic rather than grammatical in spirit, for instance, rules of deletion and generalization, which work this way.

John and Mary are jogging in the park.

Question: *Where are they?*

Paraphrased answer: *They are in the park.*

In this answer, the information about jogging was deleted as unnecessary in an adequate paraphrase.

Alice is eating a plum, and Bob is eating an apple.

Question: *What are they eating?*

Paraphrased answer: *They are eating fruit.*

These exercises are very simple, but it is easy to generate more subtle and difficult cases. Many of the main heuristic principles of paraphrase are easy to state. Here are two examples, much in the spirit of Paul Grice's maxims of conversation (1989).¹

- (1) Delete information not relevant to the question asked or point being focused on in a conversation.
- (2) Do not add information not present in the original text or speech being paraphrased.²

I'll end with what I consider is the biggest failure of technological development of educational use of computers, one that I forecast in the 1960s would be widely used by now. This is speech interaction between student and computer. Now, of course, it is easy enough to talk to the student with the kind of programs I have just been discussing and we have a great deal of audio instruction in the courses. But we do not, at present, have extensive use of speech recognition for responses of students. This, it seems to me, is a significant failure, one that we actually will be able to master once we make a sustained effort. It is a failing that is more general than just our relatively small effort at EPGY. One of the great surprises that by now (2006) operating systems for computers that include sophisticated recognition software for speech interaction have not been developed and more widely used than is actually the case. Certainly, in this century, one of the significant computer developments, from the standpoint of broad use, will be the reduction of the use of the keyboard and the natural use of the human voice to interact with hardware devices everywhere. As we do so, even the psychology of the way we think about our computers will change.

¹ The first two Grice maxims (1989, p. 26) are the following:

1. Make your contribution as informative as is required (for the current purposes of the exchange).
2. Do not make your contribution more informative than is required.

Surprisingly, Grice has only one short indexed reference to paraphrase in the 1989 volume.

² Systematic paraphrasing has a long and distinguished history that reaches back to at least late ancient times. One of the most influential examples has been Themistius' self-conscious paraphrase of Aristotle's *De Anima* and other works, written in the fourth century A.D. His paraphrastic commentary on the *De Anima* is now available in an English translation (1996).

Philosophy and Science

In what I said about the earlier part of my career under this heading, I mentioned that I had worked both in philosophy and science, though from the standpoint of research, I primarily thought of myself as a philosopher of science. I think this is probably not really the most accurate characterization. I continue to do and continue to have great interest in the philosophy of science, but it is certainly also true that, in many respects, more of my energy in the last quarter of a century has been devoted to scientific activities. A good record of these activities is to be found in my too lengthy 2002 book, *Representation and Invariance of Scientific Structures*. This large book summarizes much of the work I have done in many different areas. I think of it as written for technically minded philosophers of science, but the details go into individual sciences where I have done research at various times. The two main areas are psychology, above all, and physics.

I could take another line and say that a distinction between philosophy of science and science is in itself incorrect. In many ways I am sympathetic with such a summary of Quine's view, namely, that philosophy should mainly be philosophy of science and philosophy of science should mainly itself be science. This is a way of saying that philosophy is not privy to any special methods different from the methods used in the sciences. I certainly very much agree with these ideas, but also think that there are special aspects of problems that are of particular philosophical interest and often cultivated only by philosophers. Current philosophers of physics do not expect to develop special theories of space and time, but rather, make philosophical commentaries on the work done by physicists. It is one of the mistakes in the philosophy of mind not to have a similar attitude toward psychology.

The two big topics from psychology that I cover in the 2002 book and that reflect my perennial interests are visual perception, especially visual perception of space, and learning theory. Both of these have vast empirical and theoretical literatures, extensive already in the nineteenth century, much of the work is not really known by most psychologists and certainly by even less philosophers. Many current questions are of permanent conceptual interest, with a strong natural philosophical aspect. I will just mention some that I discuss in detail in the book. First, concerning perception, there is the complicated and subtle query of whether or not visual space is Euclidean in character. I will not even try to summarize the kind of answer I think is appropriate, but refer to Chapter 6. Another subject in psychology, treated even more extensively, is, the foundations of learning theory and, especially the question of whether or not it is possible to build from elementary concepts of learning, such as association and similarity, the complex cognitive concepts that we think of as higher order. In spite of the fact that we know very clearly that the answer in the case of mathematics to such questions is affirmative, many psychologists

refuse to believe that really at bottom its association, as I sometimes like to say, all the way down and all the way up. But consider the mathematical case. An arbitrarily complex recursive function can be computed by a universal Turing machine with a small number of states and a small number of output symbols. It is one of the mathematical triumphs of the twentieth century to show that such a reduction of computation to very simple devices is possible. That lesson has still not sunk in as far as it should, and as widely as it should, in either psychology or philosophy. These matters are discussed extensively in Chapter 8 of the 2002 book.

The two topics from physics that I know something about in detail and can therefore write about with confidence are special relativity and quantum mechanics. I include detailed discussions of the foundational literature on special relativity, but not extended to the more complicated case of general relativity. A natural foundational question is that of giving elementary axioms for special relativity and this I try to survey and give one example of. As in all such large subjects, there is extensive literature going back to the early years of the twentieth century. I am referring here not to Einstein's famous and important 1905 paper, but actually the attempts to give axiomatizations in the synthetic and qualitative spirit familiar in geometry of the same time, i.e., axioms formulated very much in the style of Hilbert's influential nineteenth-century book on foundations of geometry. The other topic is entanglement in quantum mechanics. I said enough about that already, so I need not say more here, except that I give a summary of earlier results in Chapter 7 of the 2002 book.

Although the book is long, there are many subjects, dear to my heart and about which I know something, that I had to leave out for various reasons. The most important omission was to not include a chapter on statistics following the long chapter on probability. I originally planned such a chapter, which would move from the foundations of probability (Chapter 5), and questions about probability theory in general, to the much more particular questions of statistics and analysis of data. I have spent a great deal of my academic life analyzing data and I certainly fully intended to have quite a bit to say about statistics in the 2002 book. But I kept putting off the writing of that chapter and when I finally needed it, not enough had been done to include in the actual publication. This is a shame, because one of the great developments of the twentieth century was the move from probability, thoroughly developed in many ways by nineteenth-century mathematicians starting with the magnificent work of Laplace at the end of the eighteenth century, to statistics. As a formal subject, statistics really only began in the twentieth century, but much was done during the hundred years of that century. It is still the case that not very many books in philosophy of science discuss with any thoroughness the foundations of mathematical statistics and how those foundations carry over and affect applications. Now, I certainly admit there is not some Chinese

wall between probability theory and statistics. The concepts run back and forth between the two subjects, and yet, there is still a clear and important distinction that I did not represent well.

The other big omission was not to have much more to say about statistical mechanics, which ironically is not really about statistics, in the sense just used, but an application of probability theory to one of the most important scientific cases. I said toward the end of the book that the subject was too difficult to present rigorously. I think now, in retrospect, I could have written a short chapter, certainly, not one that covered many topics, but, of course, I didn't try to do anything like that in quantum mechanics either. The only sustained argument close to the spirit of statistical mechanics was the analysis in Chapter 4 of entropy as a complete invariant of many isomorphic stochastic processes.

Maybe I don't even think that statistical mechanics is the second most important, maybe it was the omission of what I like to call the ergodic theory of free will. I have written some papers about the application of the sort of ergodic theory discussed at the end of Chapter 4 under questions of invariance, to questions of free will. I am firmly convinced that many of the tangles of compatibilism and incompatibilism can be resolved by application of beautiful and original ideas in modern ergodic theory showing, as I like to put it, the inability to distinguish between stochastic or deterministic models of many natural phenomena. It is recognized that in almost all of the important physical cases where continuous quantities are measured, there must be errors of measurement. In practice these errors of measurement are, essentially without exception, bounded away from zero no matter how good the experiments are. Given finite errors of measurement and accompanying kinds of statistics, there are beautiful proofs, especially by Ornstein and Weiss (1991), of the following sort of theorem. Consider the motion of a billiard ball with a convex object in the middle of the table to disturb its motion. For such a setting, two theories, the deterministic one of classical mechanics plus measurement errors and a stochastic Markov theory, will be mathematically inconsistent as theories of the billiard-ball motion. But we cannot distinguish between the empirical correctness of them, no matter how many observations we take. This makes as I like to say, both determinism and indeterminism in a universal sense transcendental, and creates the right kind of arena in which to discuss intentionality and free will. It is not possible to go into my analysis of such matters here, but just to express my regrets that I did not include it in the 2002 book.

Pragmatism

There is one point about philosophy and science that I have not really emphasized in any definite way, but that has become important to me. This is my much stronger explicit interest in pragmatism. In earlier years, I was put off by its superficially developed

philosophical doctrines. They seemed to be lacking in depth and, perhaps to make a joke, any serious model theory. But this was a mistake on my part. With the modern move away from foundations as an explicit aim of most philosophical work in the sciences or mathematics, I have come to see that pragmatism now fits in very well.

It is fair to say, as I have emphasized earlier in these pages, my thinking about the relationship between philosophy and science for a very long time was in terms of constructing explicit formal structures that gave a detailed sense of how a particular part of science would look when given the kind of explicit treatment characteristic of that given structures in modern mathematics. My 2002 book gave a good many examples of this. I am not against those examples now, but already as I was writing the final version, I found myself moving toward a more pragmatic view of science. I will give just two examples.

The first concerns how I ended up treating the variety of approaches to the foundations of probability in Chapter 5. This is the longest chapter in *Representation and Invariance of Scientific Structures*. I started out, when I was writing a semi-final draft of the chapter, say five years before publication, revising material from much earlier years, that I would be particularly sympathetic to a Bayesian approach. The more I got into it, the more I realized this was not really the way I now felt. One of the things that changed by mind was when I explicitly noticed that the qualitative axioms that I liked, in terms of thinking about the formal foundations of subjective probability, for instance, the qualitative axioms for a weak ordering, were not restricted to subjective ideas about probability. I could not imagine why I hadn't thought about it more clearly earlier, but, in any case, in the final version of Chapter 5, I made a number of remarks that such qualitative approaches were also very natural for objective propensity interpretations of probability. I gave in the chapter several examples. These examples were more in terms of qualitative axioms to construct a density, such as a discrete density for the geometric distribution, or the corresponding exponential distribution for continuous phenomena, as in the case of radioactive decay.

Then I found something very reinforcing. I have always liked Fred Mosteller's down-to-earth approach to statistics and the wisdom he conveys to those of us like me who are not as well educated as he is in all matters statistical. Well, I found in examining Fred's wonderful treatise with D. L. Wallace (1964/1984) on authorship of the *Federalist Papers* that it reflected, in a way that I felt extraordinarily sympathetic to, a pragmatic approach, which I summarized at the end of Chapter 5. Mosteller reports that Bayesian friends asserted that much of the analysis seemed really Bayesian, and objectivist friends said the same. So he scarcely knew how to classify the statistical approach he and Wallace used.

Reflecting on this example caused me to go back to something that I had looked at before, namely, what about the attitude of physicists to probability, especially, in that decisive case of modern physics—the probabilities that occur in quantum mechanics? So, I put in this same last section of Chapter 5 quotations from some of the most distinguished physicists who worked on quantum mechanics in the early days. Their wholly pragmatic attitude toward probability is evident. They didn't really see it as necessary, in any sense whatsoever, to make a commitment to a foundational view, but they understood very well that the computational aspects of probability were exactly what they needed for the new theoretical treatment of quantum phenomena. I'm not going to repeat here what I say there, but I am trying to give a sense of how pragmatism has more and more dominated my own thinking.

The second example, of quite a different kind, is a paper I wrote in 1999 on pragmatism in physics, where I was concerned with earlier historical episodes. I began with the history of ancient astronomy, a subject in which surely I am a rank amateur, but from which I could not resist drawing some general parallels. I felt particularly encouraged by reading Noel Swerdlow's attractive book on Babylonian astronomy (1998). I remember asking Noel, do you read the cuneiform tablets? His answer was, "Of course not." What I loved about that answer is that Noel is one of the distinguished historians of ancient and medieval astronomy. In our conversations, if I speak carelessly and make a mistake about Ptolemy, he tells me so at once. I love the fact that he had ventured into writing this excellent analysis of the Babylonian attitude to planetary motion, and yet he himself did not read the original texts. So, I tried to push a pragmatic theory, looking at the broad history from ancient astronomy to Kepler, to show how many concepts that were important to Babylonians for making omens and the like, and later, many aspects of Greek thought as well, were simply pushed out of the way and ignored. But the varied and detailed observations made by the Babylonian astronomers and used by Ptolemy more than five hundred years later, are even of some use today. Ptolemy's own central work was preserved in the tradition of a millennium and a half span leading up to Kepler and including, of course, the less important work of Copernicus. In this long period two important things were preserved: the observations reaching back to Babylonian times, and many of the Ptolemaic methods of computation, which Copernicus himself continued to use and were only changed by the new astronomy, as Kepler called what he introduced. The whole subject was then given a much greater state of perfection by Newton, with the introduction of gravitational dynamics. But much of what Kepler and Newton did rested on the shoulders of these observational and calculational giants of the distant past. It is this that is pragmatic—keeping the useful and letting go of the rest. Let me reformulate this last remark in a more purely pragmatic way. So what usually happens in the history of science is that which is true and

useful is kept, that which is false and useless is dropped. This leaves two other cases of course. What about false and useful? Well, those traditions can last a very long time, and it is easy to cite instances of something that is not literally true but is close enough as an approximation to be very useful. (Approximations are central to much important science.) And then there is the fourth case of true and useless. There are many banal claims that are true but not useful and they get lost in time, except for the arcane interests of a few benighted scholars. This is too optimistic for a good many political and social historians who do not believe in progress, but scarcely any serious historian of astronomy or physics can or does hold their pessimistic position. Even more universally accepted is the recognition of the massive progress made in astronomy and astrophysics since ancient times.

I emphasize in saying this, I am not endorsing a pragmatic account of truth as usefulness. For me they run on separate tracks that are often correlated, but one is not definable or reducible to the other.

In 1990 I received the U.S. National Medal of Science, which I was surprised and pleased to get. At that time very few of the medals of science had been awarded for work in the social and behavior sciences. I cite the brief statement attached to the award, written, I am sure, by some of my friends, because I am proud of the summary of four decades of my work, and it may provide a quick overview for some readers:

“For his broad efforts to deepen the theoretical and empirical understanding of four major areas: the measurement of subjective probability and utility in uncertain situations; the development and testing of general learning theory; the semantics and syntax of natural language; and the use of interactive computer programs for instruction.”

References

Bell, J.S. (1966). On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics* **38**: 447–52.

Berkeley, G. (1710). *Principles of human knowledge*. Dublin: Jeremy Pemyat.

Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review* **57**: 94–107.

Estes, W. K. (1959). Component and pattern models with Markovian interpretations. In R. R. Bush and W. K. Estes (Eds.) *Studies in mathematical learning theory*. Stanford: Stanford University Press, pp. 9–52.

Grice, P. (1989). *Studies in the way of words*. Cambridge, London: Harvard University Press.

Holland, P. W. and P. R. Rosenbaum (1986) Conditional association and unidimensionality in monotone latent variable models. *The Annals of Statistics* **14**: 4, 1523–1543.

Hume, D. A. (1739). *A treatise on human nature*. London: John Noon.

Krantz, D. H., R. D. Luce, A. Tversky, and P. Suppes (1971). *Foundations of measurement, vol. I: additive and polynomial representations*. New York: Academic Press.

Kuramoto, Y. (1984). *Chemical oscillations, waves, and turbulence*. Berlin; New York: Springer. Republished by Dover Publications, 2003.

Locke, J. (1690). *An essay concerning human understanding*. London: Thomas Basset.

Luce, R. D., D. H. Krantz, P. Suppes and A. Tversky (1990). *Foundations of measurement, vol. III: representation, axiomatization, and invariance*. New York: Academic Press.

Milonni, P. W. (1993). *The quantum vacuum: an introduction to quantum electrodynamics*. London: Academic Press, Inc.

Mosteller, F. and D. L. Wallace (1964/1984). *Applied Bayesian and classical inference: the case of the federalist papers*. Springer Series in Statistics. New York: Springer.

Ornstein, D. S. and B. Weiss (1991). Statistical properties of chaotic systems. *Bulletin of the American Mathematical Society (New Series)* **24**: 11–116.

Suppes, P. (2002). *Representation and invariance of scientific structures*. Stanford: CSLI Publications.

Suppes, P. (2006). Transitive indistinguishability and approximate measurement with standard finite ratio-scale representations. *Journal of Mathematical Psychology*, **50**: 329–336.

Suppes P. and R. C. Atkinson (1960). *Markov learning models for multiperson interactions*. Stanford: Stanford University Press.

Suppes P. and R. Ginsburg (1963). A fundamental property of all-or-none models, binomial distribution of responses prior to conditioning, with application to concept formation in children. *Psychological Review* **70**: 139–161.

- Suppes, P. and M. Zanotti (1981). When are probabilistic explanations possible? *Synthese* **48**: 191–199. Reprinted in *Foundations of probability with applications. Selected papers, 1974–1995*. Cambridge: Cambridge University Press, 1996.
- Suppes, P., D. H. Krantz, R. D. Luce, and A. Tversky (1989). *Foundations of measurement, vol. II: geometrical, threshold and probabilistic representations*. New York: Academic Press.
- Suppes, P., Z.-L. Lu, and B. Han (1997). Brain-wave recognition of words. *Proceedings National Academy of Sciences* **94**: 14965–14969.
- Suppes, P., B. Han, J. Epelboim, and Z.-L. Lu (1999a). Invariance between subjects of brain-wave representations of language. *Proceedings of the National Academy of Sciences USA* **96**: 12953–12958.
- Suppes, P., B. Han, J. Epelboim, and Z.-L. Lu (1999b). Invariance of brain-wave representations of simple visual images and their names. *Proceedings of the National Academy of Sciences USA* **96**: 14658–14663.
- Swerdlow, N. M. (1998). *The Babylonian theory of the planets*. Princeton: Princeton University Press.
- Themistius (350 A.D./1996). *On Aristotle's On the Soul*. Translated by Robert B. Todd. New York: Cornell University Press.
- Treisman, A. M. and G. Gelade (1990). A feature-integration theory of attention. *Cognitive Psychology* **12**: 97–136.
- Zeilinger, A., (2005). The message of the quantum. *Nature* **438**: 743.