WALTER WEIMER

The Psychology of Inference and Expectation: Some Preliminary Remarks

The problems of inference are manifold: the first task facing the theorist concerned with them is to attempt an ordering of priorities, to decide which of them to tackle. One must assay not only the relative importance of the various problems but also the possibilities of achieving their successful solution; only then may the individual problems profitably be addressed. Historically, the majority of philosophers have concluded that *the* problem of inference is that of induction, and more specifically the problem of the *justification* of induction. According to the point of view from which they operate, the prime task of the philosophy of science is to show how inductive inferences can result in genuine knowledge, or, what is to say the same thing, to show that science is a rational endeavor. Inductive inference is taken by almost everyone to be the method by which science operates: and *if* it is the method of science, *then* it must be shown to be a rational source of knowledge. Hence the philosophical or 'logical' justification of induction and the demonstration that one or another 'confirmation theory' or 'inductive logic' is actually a *rational* source of knowledge loom large in the literature — as the contributions to this volume amply indicate.

Plausible though they may seem, I reject not only the identification of the problems of inference with the justification of induction, but also the conceptual framework from which that identification stems. Neither that point of view nor the problems it suggests will increase our understanding of nondemonstrative reasoning. Of the many legitimate and worthwhile problems surrounding inference, I wish to concentrate upon two that I believe to be among the most important. They may be formulated in two straightforward questions: "*What* do

AUTHOR'S NOTE: This essay was supported in part by a visiting appointment at the Minnesota Center for Philosophy of Science during the summer of 1971. I am grateful to the editors for facilities and constructive criticism provided.

we learn from facts (and experience)?" and "*How* do we learn from them?" Instead of induction I wish to discuss *learning*, or better, the nature of knowledge and its acquisition: my thesis will be that psychology, rather than philosophy, is the domain of inquiry that will ultimately provide answers to the problems of inference. These problems cannot be successfully addressed in a framework that denies their inherently psychological nature; nor can the nature of man, the psychological organism, be divorced from the nature of our knowledge. Epistemology, as the theory of knowledge *and its acquisition*, is properly one of the psychological sciences. Put another way, the problems of inference and expectation are the problems of learning and knowledge; learning and knowledge (and thus epistemology) are problems for the psychological sciences. Hume was correct both in asserting that the senses in combination with logic cannot justify nondemonstrative inference and in relegating such reasoning to "mere animal belief."

The Dilemmas of Justificationism

The burden of this part of the essay is to lead us to an alternative point of view from which to conceive of science and philosophy *that does not lead to despair when the truth of both Humean claims is acknowledged*. To do this we must first explore the point of view from which the majority of philosophers operate and then expose several of the structural fusions and confusions that it engenders. For only if we abandon the justificationist's quest can the problems of inference be cast in such a form that insoluble dilemmas do not result at every turn. Unless it is clear that what I have elsewhere[1] called the justificationist metatheory of science *must* be abandoned, it will make little sense to conceive of the problems of nondemonstrative inference as problems for psychology.

Justificationism as a metatheory of science. Justificationism as a point of view or metatheory of scientific rationality is a "rule system" or "grammar" that specifies how our concepts can be formed. It enshrines

[1] See my *Psychology and the Conceptual Foundations of Science* (Englewood Cilffs, N.J.: Prentice-Hall, in press). The term 'justificational philosophies of criticism' was introduced by W. W. Bartley III in his brilliant and neglected book *The Retreat to Commitment* (New York: Knopf, 1962). By referring to justificationism as a *metatheory* of the scientific endeavor I have, I believe, only developed consistently ideas already implicit in Bartley's presentation.

a number of definitional fusions and confusions that are characteristic of its outlook, and these fusions and confusions, which *are* its rules of conceptual classification, are the defining traits of this point of view. By fusing certain concepts together it defines *in advance* what can be taken as appropriate answers to the fundamental questions of the nature of science, the nature of our knowledge and its acquisition, and even the nature of rationality (and rational inquiry). By listing *some* of these major structural fusions and confusions we can see how a justificationist *must* approach the questions and issues concerning our knowledge and its acquisition that are our concern in this essay.[2]

The justificationist conception of knowledge requires the identification of knowledge with both *proof* and *authority*. A putative knowledge claim cannot be accepted as genuine knowledge unless it can be proven, and it cannot be proven except by submission to the appropriate epistemological authority. For the empiricist such as Locke, *the* epistemological authority is the deliverance of sense: all genuine knowledge must be founded upon the authority of sense experience. For the intellectualist such as Descartes, the supreme epistemological authority needed to certify a knowledge claim as genuine is rational intuition. Regardless of the particular philosophy endorsed, the justificationist philosopher will not accept as genuine knowledge any claim that cannot be validated by whatever ultimate epistemological authority he accepts. To repeat: justificationism fuses and confuses knowledge with proof and with authority.

A third fusion follows from the identification of knowledge with proven assertion: the conception of 'eternally valid' knowledge gradually accumulating (through inductive inference) into the body of scientific knowledge. If knowledge is proven, then once certified it remains so forever. Thus scientific progress *must* be the accumulation of more and more eternally valid truths. Justificationist historiography, enshrining this 'cumulative record' position, rewrites history to guarantee the continuity (and therefore the rationality) of scientific progress.

The goal of science for classic justificationism was to establish that all the propositions of science have the same probability of being true:

[2] The few points we can mention, although crucial to an understanding of justificationism as a metatheory of science, are certainly not exhaustive of the position. A far more comprehensive account is found in my *Psychology and the Conceptual Foundations of Science*.

the probability value 1 or certainty. For every proposition h of science, given the relevant evidence e, the probability p of that proposition must, according to the demands of justificationist rationality and intellectual integrity, be shown to be expressed as $P(h,e) = 1$. One reads this formula as the probability of hypothesis h being true on the basis of evidence e is 1 (or certainty). To establish this, justificationism had to develop a logic of scientific assessment: an *inductive* logic.

By defining its concepts in this manner the justificationist approach *demands* that certain things obtain in science and its methodology. Half a dozen of the most important requirements are the following: First, there must be an 'empirical basis' of facts that are *known for certain*. This is the *foundation of empirical knowledge* (the basis for inductive inference). Second, 'theories' are second-class citizens, being 'derivative' from facts and accumulated generalizations (i.e., they are inductions based upon inductions). Third, science must be cumulative and gradual: facts must be piled upon facts to construct the edifice of scientific knowledge. Fourth, factual *meaning* must be fixed for once and for all *independently* of theory, and must remain invariant. Fifth, explanation consists in showing that a 'proven' factual proposition follows deductively (i.e., logically) from a theoretical proposition. Sixth, science evaluates *one* theory at a time, never two or more. Justificationism enshrines a *monotheoretical* model of assessment.

We could elaborate numerous other fusions and confusions at the heart of justificationism, but there is only one that is essential to the problems of inference: we must see how *certainty* (*proven* assertion) in the definition of knowledge becomes '*probability*'. Classic justificationism gave way to *neojustificationism* when *it* was gradually realized that *no* scientific proposition is provable, or can be "known for certain." (Fries showed this in 1828, as a special case of the logical thesis that the logical relations, such as provability, consistency, etc., can refer only to propositions. And propositions can be derived *only* from other propositions, never from 'facts'.) Thus *all theories are equally unprovable*, and $P(h,e) = 0$, always. Thus failed the method of justification for knowledge claims. The basis of knowledge failed also when Pierre Duhem noted that science is fact-correcting in its nature rather than fact-preserving. New theories often refute the facts of older ones. There is no eternally valid 'factual basis.' Popper and his students[3] have been calling attention to the failure of both the method and the basis of justificationism for several decades. But their criticism has in the main been ignored, because it has been assumed that with the relaxation of the definition of knowledge to allow probable assertions as genuine knowledge claims, inductive 'logic' could be rescued, i.e., *justified*, as a legitimate *source* of knowledge and simultaneously shown to be the *method* of its acquisition.

So in order to save the rationality of science, justificationists watered down 'proof' to 'probable' in the definition of knowledge. The development of *probabilistic inductive logic* (from the pioneering work of W. E. Johnson at Cambridge) has been the result. Intellectual honesty now demands far less of the 'formula' for the assessment of scientific warrant: today the goal is a "confirmation theory" that will assign a probability ranging from $0 < 1$ in the formula $P(h,e) = c$ (where c is interpreted as *degree of confirmation*). Practically nothing else has changed in the switch from the classic to the 'neo' form of justificationism. The chief problem is still to justify knowledge, but now it is the justification of probable rather than certain truth claims. Induction, as the means by which we acquire and certify scientific knowledge, must be rescued from the skeptical doubts of Hume. The epistemological authority must be *relocated* in the probabilistic inductive procedure: otherwise there is no authority, hence no proof, and therefore no rationality to science. At best, it will be where Hume left it — and animal belief, for the justificationist, is merely psychological and *therefore* irrational. Skepticism, for the justificationist, means that no *genuine* knowledge is possible at all and that our illusion of knowledge is a purely psychological phenomenon.

The major fusion of concepts introduced by neojustificationism is the equation of *induction* with *probability*. Having dropped the impossible

[3] The locus classicus of Popper's argument is *The Logic of Scientific Discovery* (New York: Harper & Row, 1959), especially sec. 1 and ch. 10. See also *Conjectures and Refutations* (New York: Harper & Row, 1963), especially ch. 1. The most telling single presentation thus far is Lakatos' essay "Changes in the Problem of Inductive Logic," in Lakatos, ed., The Problem of Inductive Logic (Amsterdam: North-Holland, 1968), pp. 315–417. See also Lakatos's "Falsification and the Methodology of Scientific Research Programmes," in Lakatos and Musgrave, eds., *Criticism and the Growth of Knowledge* (Cambridge: At the University Press, 1970); J. Agassi, *Towards an Historiography of Science*, Beiheft 2 of *History and Theory* (Middletown, Conn.: Wesleyan University Press, 1963); and J. Agassi, "Sensationalism," *Mind*, n.s. 75 (1966), 1–24.

quest for certainty, the sophisticated justificationist now aims at the next best thing: *near*-certainty. And the 'near certain' has been identified, with no reflection at all, with the 'probable.' Probabilistic inductive logic is the result: induction as the means of both knowledge acquisition and assessment has become fused to the calculus of probability. This new logic of scientific method is to be an algorithmic assessment procedure for the acquisition of probable knowledge rather than certain knowledge. That is, this 'logic' is supposed to *prove* putative scientific statements to be probable (rather than prove them true).[4]

The failure of justificationism. In a sense, the real failure of justificationism is that it creates for its adherents insoluble dilemmas. By fusing together concepts that do not belong together, it creates a metatheoretical structure that must inevitably lead even its most brilliant practitioners to an insurmountable wall of dilemmas, infinite regresses, and problems that admit no solutions. The justificationist quest, by its very nature, is impossibly difficult: the goal cannot be reached, no matter how ingenious its pursuer may be. The superiority of nonjustificational approaches to science and its growth is that they show, from alternative points of view, either how these dilemmas do not arise or how the problems they pose are soluble in a consistent manner.

Considering only the problem of inference (or knowledge acquisi-

[4] Logical positivism and empiricism, as the dominant 'received view' philosophies of science, are easily shown to be justificationist at heart. All the "theories" of cognitive significance (principles of verifiability, etc.) are explicit statements of the equation of knowledge with proof (or probability). *Meaning* is likewise assimilated to proof. The appeal is to a rational authority to say "What do you mean?" and "How do you know?" and science, which is the meaningful and the sensible, is demarcated from nonscience by the fact that the latter is unverifiable, i.e., meaningless, and also non-sense. The entire unity of science movement stems from the notion of a common, intersubjective, fixed-for-certain-for-all-time basis of "facts." All science must be Kuhnian normal science "fact accumulation." There can be no revolutionary science since science never overturns prior certified facts. There is no real problem in explaining the growth of scientific knowledge or of our acquisition of knowledge. There remains only the task of justifying induction as the means by which knowledge is achieved (i.e., the task of *proving* that induction yields *valid* knowledge). The construction of a theory of the scientific merit of a proposition is the prime task for philosophy, and that theory is universally assumed to be inductive logic or 'confirmation theory.' The growth of knowledge remains the gradual (inductive) accumulation of (probable) truths (because a theory of scientific merit is *automatically* one of growth). And the guidance of future research, i.e., the problems of pragmatic action in actual scientific practice, provides no problem at all: since induction is the rational means of knowledge acquisition, all who are rational will employ it.

tion),[5] it is obvious that probabilistic inductive logic and the neojustificationist conception of probable knowledge are no better off than the classic positions they are assumed to be such a significant improvement upon. If one understands Fries's argument that propositions cannot be proven by facts, then it is obvious that *they cannot be 'probabilified' either*. It will not do to substitute near-certainty, in the guise of the calculus of probability, for certainty. Fries's argument shows why neither can be obtained: inductive logic cannot assess the 'merit' of propositions where 'merit' is taken to be either truth or probability value. Nor, for that matter, can they be proven false. Within the justificationist framework the skeptic, who holds that no informative knowledge is possible, always triumphs over the 'positive' justificationist. The most that the positive justificationist can do is delude himself into thinking that it is a worthwhile task to render the skeptic's acknowledged victory as bloodless as possible.[6] No wonder that contemporary empiricists are the best representatives of existentialist despair and dread — by the consistent application of their own criteria of rationality, they have shown that their trust in empiricism was not justified — and have retreated to a *faith* in empiricism rather than a defense of it.[7] This is, as Bartley

[5] The futility of the other facets of the metatheory is becoming painfully obvious from both internal and external criticism. For example, the 'foundations of knowledge' (which justificationism requires as the basis of inference) conception has been demolished by K. R. Popper in *The Logic of Scientific Discovery*; by W. Sellars in *Science, Perception and Reality* (London: Routledge & Kegan Paul, 1963), especially chs. 3, 4, and 5); and by B. Aune in *Knowledge, Mind and Nature* (New York: Random House, 1967), especially ch. 3. Justificationist 'cumulative record' historiography has been severely criticized by T. S. Kuhn in *The Structure of Scientific Revolutions*, 2nd ed. (Chicago: University of Chicago Press, 1970); by Agassi, *Towards an Historiography of Science*; and by Lakatos, "Changes in the Problem of Inductive Logic," and "Falsification and the Methodology of Scientific Research Programs." The 'meaning invariance' thesis and the conception of explanation as logical deduction have been shattered by P. K. Feyerabend's criticism, especially in "Explanation, Reduction, and Empiricism," in H. Feigl and G. Maxwell, eds., *Minnesota Studies in the Philosophy of Science*, vol. 3 (Minneapolis: University of Minnesota Press, 1962), pp. 28–97. Feyerabend's philosophy of proliferation (see "Reply to Criticism" in R. S. Cohen and M. W. Wartofsky, eds., *Boston Studies in the Philosophy of Science*, vol. 2 (New York: Humanities Press, 1965), pp. 223–61) also demolishes the monotheoretical assessment model proposed by justificationism. Bartley, *The Retreat to Commitment*, pointed out the flaws in the justificationist conception of rationality and showed how the justificationist fuses the concept of criticism to that of proof.

[6] A representative example is A. J. Ayer in *The Problem of Knowledge* (London: Penguin Books, 1956), especially pp. 68–75.

[7] For example, Bertrand Russell in the last paragraph of *Human Knowledge: Its Scope and Limits* (New York: Simon and Schuster, 1948): "Empiricism as a theory

(1962) aptly remarked, a *retreat to commitment*, and it is no different, and no less *irrational*, when practiced by the 'scientific' empiricist than by the Protestant irrationalist. Whatever rationality there is to science must forever lie *outside* the confines of justificationism.

The promise of nonjustificational metatheories. For the justificationist, informative theoretical knowledge is an illusion — a mere 'animal' belief. But what is 'animal belief' if knowledge is *not* identified with proven assertion? If knowledge is instead identified with *warranted* assertion — such that a claim constitutes *genuine* knowledge *if* it is warrantedly assertible — then mere 'animal belief' *can* constitute knowledge. An animal belief could constitute a knowledge claim if there were 'good reasons' that could be adduced in its defense. Those 'good reasons' would never *prove* a claim to be warrantedly assertible: knowledge is always conjectural, always subject to revision and reformulation. What could constitute 'good reasons' for the defense of a conjectural knowledge claim? Why do we believe that science provides 'good reasons' for its claims?

One of the reasons we believe in science is that it advances *theories* (tentative, fallible, conjectural ones) that (attempt to) *explain* the phenomena with which they deal. Human knowledge, be it scientific or commonsensical, proceeds by what Russell called the method of analysis: given data, we attempt to construct theories that most satisfactorily explain that data. The 'good reasons' with which we defend our knowledge claims are *always* theoretical reasons: that is, we defend our knowledge claims by showing that they are theoretically motivated, that they follow from a theory that is sufficiently well corroborated that we deem it worthy of continued exploration and scrutiny. The hypothetico-inferential method of explanation is the characteristic pattern of support for all knowledge claims, scientific or commonsensical.

of knowledge has proved inadequate, though less so than any other previous theory of knowledge. Indeed, such inadequacies as we have seemed to find in empiricism have been discovered by strict adherence to a doctrine by which empiricist philosophy has been inspired: that all human knowledge is uncertain, inexact, and partial. To this doctrine we have not found any limitation whatever" (p. 507).

Compare with M. B. Turner, *Philosophy and the Science of Behavior* (New York: Appleton-Century-Crofts, 1967): "It is always a bit ironical when a house which professes to virtue topples under censure by its own precept. . . . Empiricism itself is culpable, yet we have found no reliable substitute for a knowledge supported by the fact of its public communicability. For the empiricist, the alternative to absolute skepticism is the wistful embrace of a principle of convergence" (p. 7).

But how do we arrive at tentative knowledge claims in the first place? How do our 'animal beliefs' arise? Why have we inferred that reality has the characteristics that common sense and science have attributed to it, and how and why do these accounts differ in their portrayal of reality? And exactly what sort of 'knowledge' is provided by our scientific theories? When and why are our inferences rational? These and other questions loom large in the discussion that follows, because they are central to an understanding of inference and expectation. The promise of addressing these issues from a nonjustificational point of view lies in the fact that answers — albeit tentative, conjectural ones in the form of theories that will explain the phenomena involved — *can* be adduced which are not doomed to incorporate insurmountable difficulties. That is, the psychology of inference and expectation can provide principled explanations of the nature of our knowledge and its acquisition that do not bog down in the impossible philosophical quest for a justification of induction. Hume's conclusions lead to skepticism *only* when the justificationist conceptions of knowledge and rationality are taken to be the only ones available. By abandoning the justificationist quest one can at least formulate problems of inference for which contingent theories can provide answers. And one can even reformulate old questions, such as the one about what guides scientific life, in a manner in which the incorrect answer "induction" does not arise. Consider how this can be so.

Inductive inference and the guidance of life. Because of the manner in which he fuses his concepts, inductive methods may be considered to be indispensable to the justificationist for two main problems. The first role for 'inductive logic' is as a theory of the instant *assessment* of the scientific warrant of a putative hypothesis. This leads to the quest for a computational formula, schematized as $P(h,e) = c$, where c or degree of confirmation stands for the rationality or warrantability of the hypothesis in question. That is, the probability of the hypothesis is equated with its scientific merit (both of which are assumed to be equivalent to its truth-value). I wish to ignore the assessment problem as beyond the scope of our present inquiry. Suffice it to say that I find the Popperian arguments against 'inductivism' as a theory of assessment quite convincing: so-called 'inductive method' *plays no role in the assessment* *of the scientific merit of theoretical statements.*[8] But there is another role for induction that the justificationist has up his sleeve, and it seems to be even more fundamental than the role of assessment. That is, there are justificationists who will concede that Popper's criticism of inductivism as a means of assessment is devastating, yet still hold that inductive methodology is indispensable to the guidance of scientific life. Thus the second major problem for which inductive method is deemed essential concerns knowledge acquisition: it is deemed indispensable in the quest for a guide to future scientific research. From the justificationist point of view, the problem of acceptance merges indistinguishably into that of providing a guide to scientific life. Inductive logic is not only a logic of assessment but it is supposed to be a logic of scientific discovery: the inductive judge (the confirmation formula) is also to guide scientific life. That is, it is assumed that information on the scientific warrant of propositions is somehow capable of guiding future research: that 'degree of confirmation' statements provided by the inductive judge can tell the researcher where to go in subsequent research.[9]

[8] See Popper, *The Logic of Scientific Discovery*, sec. 1 and ch. 10, and *Conjectures and Refutations*, chs. 1 and 11; and also Lakatos, "Changes in the Problem of Inductive Logic," and "Falsification and the Methodology of Scientific Research Programs"; also the contributions of Popper, Watkins, and Carnap in I. Lakatos, ed., *The Problem of Inductive Logic*.

[9] A representative argument in favor of the necessity of inductive method for guidance to future research is due to Wesley Salmon. See especially his book *The Foundations of Scientific Inference* (Pittsburgh: University of Pittsburgh Press, 1967) and such essays as "The Justification of Inductive Rules of Inference," in Lakatos, ed., *The Problem of Inductive Logic*, pp. 94–97.

Salmon's contention is that as long as any nondemonstrative inference (such as Salmon sees Popper's concept of corroboration to be) is deemed necessary for science, then the problem of induction is being smuggled in, and since *the* problem is that of justification, one must now justify this form of inference. But why must science necessarily contain some form of 'nondemonstrative' inference? Salmon feels that this is obvious because a science based upon observation statements (a factual basis) and deductive logic *alone* would be impotent: it could tell us nothing new, that is, could not tell us how knowledge is acquired and therefore could not guide future research. Since science "obviously" does tell us new and informative things, deductivism plus corroboration must be an ampliative inference procedure. Salmon writes: "Popper furnishes a method for selecting hypotheses whose content exceeds that of the relevant available observation statements. Demonstrative inference alone cannot accomplish this task, for valid deductions are nonampliative — i.e., their conclusions cannot exceed their premises in content. Furthermore, Popper's theory does not claim that basic statements plus deduction can give us scientific theory; instead, corroboration is introduced. Corroboration is, I think, a nondemonstrative kind of inference. It is a way of providing for the acceptance of hypotheses even though the content of these hypotheses goes far beyond that of the basic statements. *Modus tollens* without corroboration is empty; *modus tollens* with corroboration is induction" (1967, p. 28).

It is, of course, easy to see how a justificationist could be led to think that induction is a guide to scientific life. He would think that information gained by inductive methods could *enrich* the premises from which one could reason one's way to future conduct. For example, if I know that one horse running in a given race has won four out of his last five starts, while no other horse in the field has won more than one out of his last five starts, it would be "natural" (so the argument goes) to bet on that horse rather than one of the others. By similar reasoning, if a certain theory is highly 'inductively' established (i.e., has 'won' a large number of its experimental 'races' in the past) it would be "natural" not only to conclude that it is probably true (the assessment problem) but also to continue research based upon it. Information gathered by 'inductive inference' in the past is supposed to be relevant to the determination of future behavior.

But again, the appraisal of performance in the past is being confused with guidance in the future. It is obvious that the mere confirmation of an expectancy provides no basis for inference at all. Indeed, that *is* the classical problem of induction! The classic problem is that information about one event or series of events cannot provide any information about another event or series of events unless one can justify an inductive inference from the one to the other. (The temporal factor is irrelevant: it is the *inference* from past to present that is in question, not the change in time.) So the justificationist who thinks that induction is a guide to life must also think he has solved the classic problem of induction. Since no one *has* solved the justificationist problem of induction, anyone who believes induction is a guide to life is

He then proceeds to rehash the "necessity" for justification of nondemonstrative inference: "With the same force and logic with which Hume raised questions about the justification of induction, we may and must raise problems about the justification of any kind of demonstrative inference" (p. 28).

Notice two things about this argument. First, it is actually directed to corroboration as a problem of nondemonstrative 'inference' in the appraisal of theories. It ignores the repeated Popperian reply that corroboration is analytic and hence not nondemonstrative (see Watkins, "Non-inductive Corroboration," in Lakatos, ed., *The Problem of Inductive Logic*, especially p. 63), and therefore not an instrument of prediction or guidance. Second, since it is cast within the justificationist framework, it concentrates upon justification as the only significant problem. That is, it never gets to the problem of the guidance of scientific life at all — it merely states that if science is not to be totally empty there must be such a guide, and it assumes that that guide must be 'inductive inference'. There is no positive argument in favor of induction as a guide to scientific life at all, only the truism that science is not sterile.

simply being intellectually dishonest. To repeat: within this framework, we never get to guidance at all — since *all* the uses of inductive logic hinge upon its justification, *nothing*[10] gets off the ground unless a successful justification can be found.

Like the confirmation theorist, I take very seriously the problem of guidance in scientific research: indeed these problems of *expectation* are intimately related to what and how we learn from scientific "data." But I deny that the problem of guidance is in any simple way connected to the problem of acceptance: the acceptance of either a given proposition or an entire body of scientific knowledge claims says absolutely nothing about the course that future research should take. One could think that these two orthogonal problems were aspects of the same problem only if the justificationist point of view, which identifies both problems with 'induction,' is presupposed. To examine the problems of inference we must hold the problems of acceptance, which are the chief problems of methodology in scientific research, in abeyance. Instead of asking "What guides scientific life?" and providing yet another twist to the incorrect answer "inductive inference," we should ask "What is the nature of scientific concept formation?" and the further question "How is concept formation related to data?"

One way of restating the so-called logical problem of induction within the psychological framework is to note that the difficulty is that there are an infinitude of inferences (or theories, or models) that are equally *not supported* (i.e., neither determined to be true nor shown

[10] Assuming that the *only* information available was an appraisal of past success, *would* it actually be rational to guide conduct (to "bet") on that basis? Certainly not. This is because there is *always* an infinitude of "theories" that are consonant with any "data" whatsoever, and unless one has other, *independent* reasons for believing one of them, it is not rational to support (i.e., to make a commitment, such as placing a bet) *any* of them. The race track analogy disguises this because everyone brings their "intuitive" knowledge of the racing situation to bear upon the information at hand. For example, the "theory" that the horse which has won in the past will win the "present" race is certainly compatible with the "data" as given, i.e., his past performance. But equally compatible a priori are the "theories" that he will be sick or break a leg, have a bad jockey, be frightened by a colorful spectator, etc., etc., and *lose*. The "good reasons" that can be adduced for betting on a horse that has won in the past, which render such behavior rational, always involve knowledge of reality that transcends empirical and logical considerations. And this nondemonstrative knowledge will in its turn have to be justified before it is rational to accept one out of an infinitude of potential theories of future performance. Once again: the justificationist cannot put inductive inference to *any* use until it has been justified as a rational source of knowledge. And it is precisely this that he has never succeeded in doing.

to be trustworthy guides) by any given data. The transition from the justificationist's problem of sanctioning inference to the nonjustificational one of learning and concept formation is made by asking what *constrains* our patterns of inference in such a manner that we sometimes hit upon a 'good' concept or theory. We *do* learn informative things with the aid of (some) data in the process of scientific concept formation. It will not suffice to explain how and why this is so by vague reference to "insight" or "creative imagination." And the problem cannot be avoided by sharply distinguishing between the contexts of discovery and justification. Nor will the neo-Popperian ploy of distinguishing pure science from applied technology avoid the issue.[11] Sooner or later we must face the question of how we learn to form scientific concepts and how scientific concept formation is related to the guidance of scientific practice. The problems of inference and expectation require a psychological solution. Hume was correct in dropping these problems into the (according to the philosopher, obviously disreputable) hands of the psychologist. But before we despair of this fate, let us see what the psychologist must account for, and then ask what theories he has available to explain the phenomena in question. Like Russell, I think that such explanations are all that we, as prudent men, can hope for from a contingent theory.

The Nature of Scientific Knowledge

As far as logic is concerned, *all* scientific and 'common-sense' knowledge is on a par: it is all equally nondemonstrative in character and,

[11] Popperians freely admit that the logic of science is deductive only and that the corroboration of a theory, as an analytic appraisal of past success, says nothing about the future success that can be expected for the theory. Indeed, they often claim that science has no guidance at all (other than creative imagination, which they do not even try to explain) and that only (mere) technology requires *reliable* theories that can guide future practice (i.e., provide safe bets). Both J. Agassi, "The Confusion between Science and Technology in Standard Philosophies of Science," *Technology and Culture*, 7 (1966), 348–66, and "Positive Evidence in Science and Technology," *Philosophy of Science*, 37 (1970), 261–70, and T. W. Settle, "Induction and Probability Unfused," in P. A. Schilpp, ed., *The Philosophy of Karl Popper* (La Salle, Ill.: Open Court, 1974), pp. 697–749, have taken this approach. In the first place, the problem of guidance is not *just* the problem of finding reliable or technologically safe theories: everyone (including especially the Popperians) wants to *increase our knowledge* by exploring more adequate theories. Second, the problem of creativity, which is at the heart of the matter, is not addressed at all by this approach. Scientific concept formation is beyond the pale of methodology for the Popperian.

hence, equally incapable of justification. The *assumption* that human beings are in possession of any knowledge *whatsoever* about the world can only be defended by adducing two classes of further assumptions in its support. The first class of assumptions that is necessary to defend the claim that we have knowledge (*any* knowledge) concerns the actual structure of reality — the world we live in as opposed to the possible ones. The second class of assumptions concerns the nature of the organism that *has* knowledge — we must be structured in such a manner that we can come to know the environment in which, as a matter of (contingent) fact, we happen to exist. Indeed, we could not have *knowledge* in (or from) the 'trivialities' of immediate perception (as the empiricist claims) unless our theories of how the knowing organism is structured are correct. The so-called 'direct' knowledge disclosed by *perception is as inferential* in character as our most speculative conjectures. (That is, the assumption that we have even 'trivial' knowledge by perceptual experience depends upon our theories of cognizing organisms. It is the further assumption that we have other, *non*trivial contingent knowledge that is dependent upon the structure of reality.) Although we cannot justify the belief that we actually possess knowledge, we most certainly can adduce good reasons for thinking that we do have (and can gain more) knowledge. Those 'good reasons' are our contingent theories about the nature of our knowledge and the nature of knowledge acquisition. Although we cannot *prove* that we have knowledge (by this line of defense), there is nothing whatsoever irrational about seeking it and continuing to subject every theory that we deem worthy of entertaining to critical scrutiny. Constructing contingent theories of the nature of ourselves and our knowledge would be irrational *only* within the justificationist framework, where rationality remains equated with proof and certainty and justification.

The justificationist is literally a prisoner within the confines of his own framework: because of the structural fusions and confusions inherent in it, that framework prevents many crucial problems from ever being satisfactorily resolved. The nature of our knowledge and the means of its acquisition are problems that it has done a particularly fine job of botching up, and it will take the remainder of this essay to begin to locate the directions in which genuine solutions to these problems should be sought. To begin this task of "relocating" the problems posed by inference and expectation we shall spend the rest of this part of the essay exploring the character of so-called 'scientific' knowledge, to see what the end product is that our theory of inference must produce. In the subsequent part the nature of the organism that makes inferences will be located within a biological, evolutionary framework. Only after these preliminaries are completed may we address the psychology of inference and expectation.

Theories as structural representations of reality. The classic conception of scientific theories as an amalgam of postulates, syntactic calculus, correspondence rules, and a model, while perhaps not too incorrect as a description, is totally incapable of explaining what theories *do,* i.e., how they function in embodying and generating scientific knowledge. In order to understand the role of inference and expectation in the acquisition of knowledge we must know both the nature of our knowledge and how it is embodied. Only after we are clear on what knowledge is and how it is embodied may we ask how it is transmitted and how we acquire new knowledge. Thus we must begin by asking how theories, as embodiments of scientific knowledge, function in the scientist's head. That is, how does the conceptual structure of a theory make 'understanding' possible?

The thesis I wish to defend is that theories are nothing more, nothing less, than conceptual *points of view* that organize 'observations' or 'factual data' into meaningful patterns. *Theories are a way of seeing reality.* They render their subject matter intelligible by exhibiting the *structure* of the phenomena with which they deal. Theories are structured representations that constitute a way of 'seeing' or perceiving reality. Theories function as instruments of understanding *because* they are structural representations of a domain.

The conception of theories *as a way of seeing* is very old: *theoria* for the ancient Greeks meant a 'vision,' and metaphysicians have long been known as 'visionaries.' This way of 'picturing' theories has the advantage of bringing to the fore their inseparable relation to perception: we perceive reality by means of the structural representations which are our theories. Theories let us see the structure of reality. And the key feature of theories as 'points of view' or 'ways of perceiving' is that they are *generative* or *productive* (in the transformational grammarian's sense): they enable us to *infer* to indefinitely extended instances of unobserved cases. Theories allow us to represent an indefinitely large number of data points. States of affairs are represented by theories, and that representation very literally *is* our understanding of them.

The thesis that theories are structural representations of reality was admirably defended by N. R. Hanson.[12] Consider this comment by Hanson:

Scientifically understanding phenomena x, y, and z consists in perceiving what kinds of phenomena they are — how they relate one to the other within some larger epistemic context; how they are dependent upon, or interfere with, one another. Insights into such relations "out there" are generable within our perceptions of the structures of theories; these theoretical structures function vis-a-vis our linguistic references to x, y, and z in a way analogous to how the scene stands to the tree and hill "out there" and also to the painted patches on canvas. Thus, I suggest that in contrast to the delineation of theories as "ideal languages" or "Euclidean hypothetico-deductive structures," the important function of scientific theory is to provide such structural representations of phenomena that to understand how the elements in the theoretical representation "hang together" is to discover a way in which the facts of the world "hang together." In short, scientific theories do not always argue us into the truth; they do not always demonstrate deductively and forcefully what is the case. Often they show what could be the case with perplexing phenomena, by relating representations of those phenomena in ways which are themselves possible representations of relationships obtaining "out there." Theories provide patterns for ordering phenomena. They do this, just as much as they provide inference channels through which to argue toward descriptions of phenomena (1970, p. 240).

Scientific theories represent the phenomena of their domains by being structural patterns that are isomorphic to the structural relations that obtain in the phenomena themselves.

What is it that allows one to say that an artist has captured "the same" scene as the actual landscape "out there"? Subject matters and their representations must share something in common: *structure*. Consider a song being sung, its recording as a record, and its musical score on note paper: it is knowledge of their common structure that enables us to say that these disparate manifestations all represent the same (abstract) entity. Likewise, scientific theories represent external reality. But representation is not just (iconic) picturing: theories are not pho-

[12] See his *Patterns of Discovery* (Cambridge: At the University Press, 1958), and "A Picture Theory of Theory Meaning," in R. Colodny, ed., *The Nature and Function of Scientific Theories* (Pittsburgh: University of Pittsburgh Press, 1970), pp. 233–74.

tographs. Scientific theories are *abstract* representations — they are like graphs or flow charts rather than pictures. The correlation between theoretical terms and external objects is conventional, as is the correlation between the symbols of a map and the terrain it represents. But a map is informative because it shares relational structures with the terrain in question. Charts, graphs, etc., enable one to *see* the dynamical structural characteristics of the phenomena they represent. "They provide a pattern through which the multiform and chaotic manifestations of the original appear as correlated parameters. These patterns provide conceptual gestalts which allow inferences from one parameter to another parameter throughout a charted system of data lines" (*ibid.*, p. 250).

How far is it from graphs and flow charts to our sophisticated mathematico-physical theories? Not far at all: mathematics is simply a tool for transforming the *perceptual* structural representation provided by a graph or flow chart into an *equation* that *shares that same structure*. "Scientific theories enable us to understand perplexing phenomena precisely because they enable us to see on the page some of the same structures which are there in the phenomena themselves. The theory allows us to comprehend what makes things "go" — and to work our ways into the phenomena, along the dynamical structures (as it were) by way of inferences through the algebra which itself has the same structure as the phenomena, or at least a structure compatible with the phenomena" (*ibid.*, p. 273). All our 'perceptual' knowledge of the structure of reality is (in principle) capable of transformation by mathematics into equations that equally represent the structure of reality. Our understanding of reality is nothing more, nothing less, than our ability to perceive such structures (i.e., to infer it from our phenomenal experience). The problem for a psychology of inference and expectation is to characterize the organism that does this 'perceiving,' in order to understand how our knowledge is acquired.

Structural realism and the nature of scientific knowledge. What scientific theories do is to represent, in concise and manageable form, the structure of their domains: they are structural representations of reality. But what is 'structure,' and why is it that when we represent reality we represent only the structure that is definitive of it? Why is our knowledge of the external world structural? The answer, simply put, is that the only nonstructural or intrinsic properties of existents of which we have any comprehension whatsoever are the properties of our own phenomenal experience. Our knowledge by acquaintance, our phenomenal experience, is qualitatively different from our knowledge of the nonmental realm. What we know of 'external' (external to the mental realm of 'direct awareness') reality is knowledge *by description*: we are not acquainted with the intrinsic properties of external objects at all.

The genesis of scientific knowledge of reality, insofar as it is knowledge by description only, may be compared to the unfolding plot of a detective story. Assuming that a crime has been committed, the detective must track down the culprit. Now if no one saw the commission of the crime, if it was not an 'ingredient' (to use Russell's term) in the perceptual knowledge by acquaintance of anyone, then the detective is in exactly the same situation as the scientist. Although the detective is in a sense at a loss because no one saw the culprit, he is not entirely hamstrung: he can gain sufficient knowledge to enable the guilty person to be "brought to justice." That is, the detective can come to *know* the culprit, by a (usually quite involved) process of inferential reasoning, without ever becoming directly acquainted with him. All the "clues" regarding the nature and whereabouts of the culprit, from those present at the scene of the crime through to the inevitable chase scene, can provide the detective with knowledge by description, knowledge of structural properties, of the guilty party. He can come to learn, for instance, that he is male, between 6′ and 6′2″ tall, that he weighs approximately 175 lbs., has a gold ring on a certain finger, walks with a limp, etc., etc. In short, the detective can come to identify positively an individual without ever having 'seen' or become 'acquainted' with him. The guilty party can be known in exactly the same way that an abstract and unobservable scientific entity, such as a proton or a phoneme, can be 'known.'

But there is one major difference between the fictional detective story and the scientific detective story: the analogy is incorrect in one crucial respect. At the end of the novel, the culprit is identified — somebody shouts, "I know him, that's Joe Smith!" or some such. What they mean by identifying the culprit by his proper name is that they are "acquainted" with him, and hence their knowledge of Joe Smith is not only knowledge by description but *also* knowledge by acquaintance. But proper names do not occur in scientific investigation: the scientist never concludes by saying, "And that's how I met Joe Neutrino." This is because the scientist is never directly acquainted with the intrinsic properties of *any* 'external object': *all* the properties which we normally attribute to nonmental objects exist *wholly* within the mind. As Berkeley pointed out, Locke's distinction of primary and secondary qualities is invalid: *all* the properties of objects, including what Locke thought were primary or intrinsic properties, exist only within the mental realm of the perceiver. All such properties and qualities are "in" the mental world rather than "in" the objects themselves. Can science conceive of *any* properties of the entities of the nonmental world? According to the view I am proposing, which is really Bertrand Russell's view,[13] what we can know about external or nonmental objects (including even our physical bodies in the latter class) is only their structural properties and relations, never their intrinsic properties.

Commenting on the question of what science knows of the nonmental realm, Maxwell[14] has said:

> [T]he only aspects of the nonmental world of which we can have any knowledge or any conception are purely structural (or, in other words, purely formal). The *details* of the answer consist of an explication of "structural" or "formal" adequate to the task at hand. . . . The notion of *form* or *structure* needed here may accurately be said to be logical (and/or mathematical) and, in a sense, abstract; *characterizations* of instances of it will be in terms of logic alone, i.e., the logical connectives, quantifiers, and variables — they will contain no descriptive terms. . . . *Structure*, in the sense we require, *must be factual and contingent and*, at least in its exemplifications, *concrete*. Furthermore, it cannot be emphasized too strongly — what should already be obvious — that structure is *not* linguistic nor, even, conceptual in character; it is an objective feature of the real world (p. 153).

But *structure* should not be identified *solely* with the notion of form: rather "it is form plus causal connections with experience" (*ibid.*, p. 154).

In overview, structural realism is *structural* in the sense that our knowledge of the entire nonmental world, from our bodies to physical objects, is of the structural characteristics of that world rather than

[13] As first presented in *The Analysis of Matter* (London: Allen and Unwin, 1927), and developed consistently in *Human Knowledge: Its Scope and Limits.*

[14] G. Maxwell, "Scientific Methodology and the Causal Theory of Perception," in Lakatos and Musgrave, eds., *Problems in the Philosophy of Science* (Amsterdam: North-Holland, 1968), pp. 148–77.

of the *intrinsic properties* of those objects composing it. Structural realism is a *realism* in that physical objects *do* causally influence our perceptions: properly stated, the causal theory of perception is true.

Structural realism may be stated in terms reflecting (and related to) Kant's distinction of *phenomena* from *noumena*:

On the one hand there is the realm of phenomena. These are wholly *in the mind* (in our sense). Of the phenomena and only of the phenomena do we have *direct knowledge*. On the other hand, there are the things in themselves, and here our divergence from the views of Kant is great: although we have no *direct* knowledge of the latter, the bulk of our common sense knowledge and our scientific knowledge *is* of them. Among them are not only electrons, protons, forces, and fields but also tables, chairs, and human bodies. All of our knowledge of these is, of course, indirect and may be generally characterized as hypothetico-deductive (or, better, hypothetico-inferential). . . . The implications of this are manifold and profound. In fact, it requires considerable time and thought — for most of us at least — to realize what a drastic revision of our usual conceptions, including our scientific conceptions is required (*ibid.*, pp. 154–55).

One of the concepts most in need of revision (actually reinterpretation) is that of *observation*. Both common sense and unreflective scientific thought speak commonly of "observing physical objects." But, strictly speaking, there can be no observation of public objects at all:

Observation, as usually conceived, is a naive realist concept through and through. Therefore, if structural realism is true, then, in any usual sense of "observation," we observe neither public objects nor entities in our minds: we never observe anything at all. For example, if I have a dream, no matter how vivid, of seeing a white dog, I cannot be said to have observed a dog or, indeed, anything, since there existed nothing corresponding in any straightforward way to the ostensibly observed object. On the other hand, if I do what ordinarily could be called "actually seeing a white dog," then . . . I do not actually observe anything (or even *see* in the usual sense), for there *is* nothing external to me which is white in the usual, qualitative sense, or etc., etc. (*ibid.*, p. 167).

Talk of the observation of objects, as is common in science and common sense, must be taken as a shorthand formulation, couched in terms of naive realism, for discourse about sense impressions and their structural relation to the external objects which are their causes.

Observation and the deductive unification of experience. The ordinary concept of observation, as structural realism makes clear, is a naive realist relic in scientific discourse. But we are supposed to observe facts, and facts are alleged to be the indispensable basis upon which the edifice of scientific knowledge is erected (by so-called inductive inference). Although there is increasing acceptance of the inherent theoretical nature (or contamination) of facts, i.e., acceptance of the doctrine of factual relativity,[15] even among staunch justificationist empiricists, there has been almost unanimous assent to the claim that the ultimate foundation of science is the perceptual experience of the scientist. There has been no general acknowledgment that science is *not based upon perceptual experience at all*, nor has the import of the idealized and abstract nature of scientific entities been acknowledged. But the total *disconnection* of theory from experience, as well as the abstract nature of scientific entities, can easily be demonstrated. When this is done, the nature and role of inference in scientific thought are seen to be drastically different from what the received view doctrines have proposed: that is, the basis of inference must be radically relocated from generalization from facts to the functioning of the nervous system. Let us follow up our presentation of structural realism with a look at the abstract and nonperceptual nature of scientific knowledge in order to relocate the real problems that they pose for inference and expectation.

[15] Factual relativity, or the conceptual nature of facts, is the doctrine that factual propositions exist only within a given conceptual framework rather than independently of theories and conceptual schemes. The data of sensory experience do not come with tags proclaiming their factual status. Observation is more than merely becoming aware of sensory input: it is the assimilation of input into a classificatory scheme that is logically prior to that input. Karl Popper is largely responsible for forcing the admission that the 'basic statements' of science are theoretically determined upon the members of the Vienna Circle; see his *Logic of Scientific Discovery*, especially section 30. Thomas Kuhn has continually emphasized factual relativity in scientific revolutions in *The Structure of Scientific Revolutions*. The pragmatist C. I. Lewis also ably defended the doctrine under the guise of the pragmatic conception of the a priori; see especially chs. 5 and 8 of *Mind and the World Order* (New York: Scribner's, 1929). N. R. Hanson also drove home the relativity of facts, as in this passage from his *Patterns of Discovery*: "If in the brilliant disc of which he is visually aware Tycho sees only the sun, then he cannot but see that it is a body which will behave in characteristically "Tychonic" ways. These serve as the foundation for Tycho's general geocentric-geostatic theories about the sun. . . . Tycho sees the sun beginning its journey from horizon to horizon. He sees that from some celestial vantage point the sun (carrying with it the moon and planets) could be watched circling our fixed earth. Watching the sun at dawn through Tychonic spectacles would be to see it in something like this way. Kepler's visual field, however, has a different conceptual organization. Yet a drawing of what he sees at dawn could be a drawing of exactly what Tycho saw, and could be recognized as such by Tycho. But Kepler will see the horizon dipping, or turning away, from our fixed local star" (1958, p. 23).

We must support two claims: first, there are no "perceptual experiences" whatsoever at the "basis" of science; second, the fundamental propositions of scientific explanatory discourse deal with abstract and ideal entities. Let us consider the latter contention first. It may be restated as the claim that science can never deal directly with particulars or individual things *at all*, but only with "abstract entities" or thing-*kinds*. Most contemporary "philosophy of science," insofar as it deals with substantive theory construction in science, is concerned with the purely formal structure of the subject matter theories it examines (in particular, those that are or can be axiomatized). Thus every textbook has its chapter on "theories in general" which treats the nature of formal systems and their utilization in theory construction and scientific explanation. But one factor that has been quite neglected is what Stephen Körner has examined:[16] the restriction which the logical structure of scientific axiomatic systems imposes upon the subject matter of such theoretical systems.

The empiricist "foundations" view of knowledge holds that that which is given in perceptual experience is to be taken as particular. But one may pose an embarrassing question at this point: how does the nominalist-empiricist *know* (or come to recognize, be acquainted with, or recollect, etc.) his given particulars as singular instances or *particulars*? That is, if a "thing" is proffered to me as "an X," where X is any description (i.e., classification) of it *whatsoever*, how can I claim to know that it *is* an X without first knowing what it is to be an instance of a thing-*kind*, namely, of kind X? In short, in order to recognize a thing (or a fact, etc.) mustn't one presuppose knowledge of, or operation within, a framework (theory) of thing-kinds? Reflection on the problem indicates that one must acknowledge that *classification is fundamentally a process of abstraction*: that is, the process of abstraction is a sine qua non for determination of concreta. Let us develop this point in another way by considering what is involved in hypothetico-deductive "scientific" explanation.

The question we must ask is this: to what extent does employment of the hypothetico-deductive method *idealize* the domain to which it is applied? In other words, what does the process of deduction require

[16] See pt. II of his *Experience and Theory* (London: Routledge & Kegan Paul, 1966).

of the empirical predicates of a science before it may legitimately be applied?

Körner considers two classes of such constraints. The first constraint requires the elimination of inexactness and indefiniteness of all the predicates. The logic of the H-D framework is an unmodified classical two-valued logic. Thus, strictly speaking, the H-D framework admits no inexact predicates whatsoever. So the H-D framework must, strictly speaking, be the logic of the finished science report, i.e., admit no inexact or 'open concepts' as Pap[17] termed the concepts of a growing science. The point is simply this: the H-D system is *not* connected to experience directly — the empirical predicates with which it deals *must* be idealizations or abstractions. Raw perceptual experience is rendered into "concrete" categories by the process of abstraction. The abstraction distinguishes between the relevant and irrelevant determinable characteristics and discards the latter. It creates new abstract determinables that *replace* the original perceptual ones at the "basis" of scientific explanatory systems. "This type of abstraction which, in order to distinguish it from other kinds, I shall call 'deductive abstraction,' and which replaces perceptual by abstract determinables, reinforces the general effect of the restrictions which the logico-mathematical framework of every theory imposes upon perceptual characteristics" (Körner, 1966, pp. 166–67).

But even this is not yet sufficient to render a so-called 'empirical' predicate fit for a position in a deductively unified system. Further idealization effectively removes the predicate from perceptual experience *entirely*:

> The disconnection of the theory from its perceptual subject-matter can now be also expressed by saying that no perceptual proposition and no perceptual predicate occurs in any deductive sequence. That this must be so is clear. All inexact perceptual predicates are precluded from occurring in any sequence, and the exact — though internally inexact — determinables have been replaced by non-perceptual predicates through abstraction — to say nothing of the further replacements due to the conditions of measurement, general and special.
>
> No perceptual proposition will be a last term in a deductive sequence. The theory will be linked to perception not by deduction but by identification (*ibid.*, pp. 168–69).

[17] See ch. 11, "Reduction and Open Concepts," in his *Semantics and Necessary Truth* (New Haven, Conn.: Yale University Press, 1958), pp. 302–60.

This means that the reference of empirical concepts cannot be given, but is rather idealized, which is to say, *constructed*, by the active abstraction of our conceptual frameworks. The concept of an ideal straight line within an axiomatic system is a far cry from *any* empirical line, such as a light ray, the surface of a straight edge, etc. But our theories of reality deal with the abstract, idealized elements such as "straight lines" and not with the lines drawn with rulers or physical objects. With that admission, *our theories become nonempirical* in the sense that their reference is to ideal, nonperceptual entities rather than concrete physical ones. The empirical becomes identified with the ideal.[18]

This effectively complicates the traditional usage and interpretation of correspondence rules or bridge laws between 'theory' and 'observation.' It is commonly assumed that correspondence rules somehow link or "translate" empirical predicates into theoretical ones. But clearly there is a "linkage" that precedes the typical notion of correspondence rule, which is required to identify an empirical predicate in the first instance. As Körner notes, most philosophers who talk of correspondence rules "are not aware that one of the fundamental transactions consists in the transposition of internally inexact, empirical predicates from a modified into an unmodified two-value logic and in their consequent replacement by internally exact ones" (*ibid.*, p. 90).

Körner's claim, then, is that the traditional usage of correspondence rules (to link empirical predicates to theoretical propositions) covers only half the problem of reference for theoretical terms: there must be additional "correspondence" rules linking the phenomena of experience, sense data if you like, with empirical predicates. And this linkage is not unambiguous or in any sense "direct": constructing an empirical predicate out of the deliverances of sense requires a theory of idealization and abstraction. That "theory of abstraction" means that cognition is not a passive register of "objective" sensations, as Locke claimed, but rather an active, indeed constructive, process is another way of stating the abstract, conceptual nature of facts. But the "theory" determining the facts in *this* instance is actually the genetic, physiological, and psychological structure of man. We must now turn to a study

[18] It is the *identification* of the empirical with the ideal that strikes at the heart of the matter: "A hypothetico-deductive system is . . . not directly connected with experience. In linking it to experience by 'identifying' some of its predicates and propositions with internally inexact empirical ones one is not ascertaining that they are identical, only treating them as if they were" (Körner, 1966, pp. 89–90).

of the a priori constraints upon human knowledge that are imposed by our physiological and psychological nature. What constrains the input of the potential environmental flux into human consciousness is actually "the way the mind works." And the way the mind works is as an inference machine. As we shall now see, this inference machine not only makes inferences in accordance with the data it has available but also constructs the very data from which it infers.

The Biological Basis of Cognition: The Organism as a Theory of Its Environment

Hume relegated nondemonstrative inference to "animal belief." Justificationist philosophers, while not denying such base accompaniments to this pattern of reasoning, have been trying to salvage a respectable *logical* problem of 'induction' ever since. Typically their resultant 'respectable' problem utilizes two tacit assumptions (among others): first, that we have 'facts' or 'observational statements' as unproblematically given, and that induction leaps beyond the firm foundation of evidence to risky theoretical formulations; second, that the psychological nature of the organism who utilizes nondemonstrative inference to obtain knowledge is all but totally irrelevant to the problems of philosophy. In a sense, the burden of this essay is to argue that these two assumptions are grossly mistaken and that the very nature of the legitimate problems of inference cannot be discerned until these assumptions are exposed as totally indefensible and then their pernicious effects eliminated. Only by purging these inadequate conceptions of knowledge and human nature from our thinking can we understand the indispensable role of inference in cognition and the relation of inference to knowledge and its acquisition. Not surprisingly, these two assumptions of traditional philosophy are intimately related: they are based on an incredibly inadequate psychology of knowledge acquisition (to say nothing of a chimerical conception of contingent yet somehow certain knowledge).

To counter the justificationist approach to the problems of inference we must get clear on the nature of human knowledge and the nature of the organism that acquires that knowledge (as well as indicate *how* knowledge is acquired). We have begun to counter the justificationist conception of knowledge and its acquisition by examining the function of scientific theories and the structural nature of our knowledge of the nonmental realm. We have also indicated the abstract and conjectural nature of observational statements, and discussed the nonperceptual nature of scientific propositions. Now we must turn even more in the direction of theoretical psychology, to examine the nature of the 'inferring' organism. In so doing we will be examining the a priori constraints upon human knowledge and its acquisition that are imposed by our psychological nature.

The fundamental thesis I wish to defend, and ultimately to elaborate, is that the problems of inference can only be understood in the context of a psychology of learning and concept formation which takes as its central tenet the thesis that *organisms* (more properly, the central nervous systems of organisms) *are theories of their environment*. The problems of nondemonstrative inference are the problems of understanding how the nervous system,[19] from its preconscious determination of the orders of sensory experience to the higher mental processes such as human thought, attempts to model, with ever increasing degrees of adequacy, the environment that the organism confronts. There is no difference in kind between the scientist inferring the most esoteric theory of reality, on the one hand, and the simplest organism's inferring the presence of food or danger in its environment. In both cases the fundamental *activity* of the nervous system is *classification* (or abstraction) and the fundamental *function* the nervous system performs is *modeling* (of the environment). Nondemonstrative inference, the process by which we gain our contingent knowledge of reality, is not different from the psychological phenomena of concept formation and learning. And we shall see subsequently that the key to these latter phenomena is found in the psychology of perception and memory.

The proposal that organisms are theories of their environments is to be understood within an adaptational or evolutionary perspective. If an organism is to survive in an uncertain world (such as the one we find ourselves inhabiting), it must be able to operate as a biological mechanism within that world. The mere fact of survival implies that

[19] Although it is more correct to say that the entire organism is a manifestation of a theory of its environment, the nervous system is the key component in the acquisition and generation of knowledge. The central nervous system (CNS) is also a theory of the environment of the CNS, which *includes* the remainder of the organism. Thus we may talk of the CNS as "the organism" and as a theory of the organism's environment without oversimplification or distortion.

an organism has been effective in maintaining an appropriate commerce with its environment. The struggle for survival is the struggle to adapt to, or to learn to utilize effectively, the environment in which the organism is situated. The survival of a species *implies* that the species is efficient in dealing with the contingencies of its environment. It is the nervous system that is ultimately responsible for the organism's perception and knowledge of, and its commerce with, the environment. The question then arises, "How does the nervous system come to have knowledge of the environment and its contingencies?" that is, "How does the CNS function?" The answer is that the nervous system functions as a theory of the environment of the organism: *its job is to make inferences* (about environmental contingencies). To the extent that the theory that the nervous system instantiates is adequate to its task and the inferences are successful, to that extent the species survives. If the theory is inadequate, the species may be expected to die out. Evolution, from this point of view, is a mechanism which allows nervous systems to construct more and more adequate theories of their environment. Man, the conscious, thinking (and sometimes) rational animal, is an emergent phenomenon in that only his CNS has developed to the extent that it can consciously create and reflect upon, as well as operate according to, theories of the environment.

The question now arises *how* organisms (such as scientists) construct their theories of the environment. The answer to this question lies in a proper unpacking of the psychological phenomena of "learning" and "concept formation." That is, this question really asks, "How do organisms form their concepts of reality?" and it is at this point that the psychology of inference and expectation enters, as an indispensable ingredient, into epistemology. The questions that must be asked are "What are 'facts' or 'data'?" and "What and how do we learn from data?" These questions take on a very strange light once it is realized that the reality which not only preconscious sensory experience but scientific thought attempts to model is known *only* as a result of the classifications imposed by the central nervous system.[20] Once (at least rudimentary) answers to these questions of perception and learning are at hand, the problems that "induction" has been presumed to be

[20] This is not, of course, a claim that the nervous system is a *homonculus* whose job is that of a taxonomist; rather it is the case that the nervous system classifies information, a task which homonculi are often postulated to perform.

necessary to solve, such as the guidance of scientific life, can then be addressed.

But first we must explore the ramifications of considering organisms as theories. There would seem at first blush to be little reason to construe either the activity of preconscious nervous processes as instantiating a theory or the conscious activity of all but trained professionals as actually theory construction. After some time is spent developing the approach with the higher mental processes, the manner in which the thesis applies to the functioning of the nervous system as a whole will be considered. But the point of all this discussion remains: unless we can understand the nature and functioning of the organism that makes inferences, we shall never understand either what or how we learn from 'facts'.

Thought and symbolism. So far as I am aware, there has been only one serious hypothesis on the nature of thought offered within the context of an adaptational or evolutionary approach to the problem. This is an hypothesis due to Kenneth Craik.[21] Craik's hypothesis is simply "that thought models, or parallels, reality — that its essential feature is not 'the mind', 'the self', 'sense-data', nor propositions but symbolism, and that this symbolism is largely of the same kind as that which is familiar to us in mechanical devices which aid thought and calculation" (p. 57). The fundamental feature of our neural machinery is its power to parallel or model external events. In terms familiar from the discussion of Hanson's conception of theories as structural representations, thought literally *is* a representation of reality. Man, the conscious organism, as a theory of its environment, models reality by means of thought. The knowledge that resides "in the head" of the scientist resides in his internal model — thinking about the environment in which we find ourselves is ipso facto to model that environment. That is, modeling (as an activity of CNS's) is a richer concept, or a more generic one, than 'representation' (which in turn is richer or more inclusive than iconic 'picturing'). And modeling, as Russell and Maxwell have been at pains to emphasize, is a structural means of representation. The models which constitute scientific theories are (to greater or lesser degree) structurally isomorphic to the reality to which they pertain. To repeat, our scientific knowledge of reality, which is embodied in our theories, which

[21] In his neglected monograph *The Nature of Explanation* (Cambridge: At the University Press, 1943).

are in turn representations, which are in turn the product of thought, which is in turn modeling, is purely of the structural rather than the intrinsic properties of reality. Craik was very clear on the structural nature of models: "By a model we thus mean any physical or chemical system which has a similar relation-structure to that of the process it imitates. By 'relation-structure' I do not mean some obscure non-physical entity which attends the model, but the fact that it is a physical working model which works in the same way as the process it parallels, in the aspects under consideration at any moment" (*ibid.*, p. 51).

Thinking, if Craik's hypothesis is tenable, is always *symbolic*: the model which is the thought is only structurally related to the reality which it imitates. The activity of neural excitation per se is totally unlike, say, the patterns of stress in a bridge, yet the *patterns* of excitation which constitute the thinking about (or calculating) that stress are isomorphic (in a structural sense) to the stress itself. "It is likely then that the nervous system is in a fortunate position, as far as modelling physical processes is concerned, in that it has only to produce combinations of excited arcs, not physical objects; its 'answer' need only be a combination of consistent patterns of excitation — not a new object that is physically and chemically stable" (*ibid.*, p. 56).

If thinking *is* modeling, then 'the organism is a theory of its environment' follows automatically. If the organism carries around in its CNS an internal model of external reality (and, of course, a model of itself and its capabilities) then it will be able to adapt to, to survive within, that environment. It will be able to try out and assess various alternative actions, react to future situations before they arise, and in general *anticipate* the environment in which it finds itself. (Our nervous systems are instruments of adaptation to our environment *because* they permit trial of alternatives for future conduct in an economical manner: thought models *potential* realities.) Modeling has survival value — it enables an organism to anticipate the future course of events *and* to *act* in accordance with that information. As shoud be obvious by now, the function of thought, as modeling, is identical to that which is attributed to inductive or nondemonstrative inference: thought-as-modeling is the vehicle by which we gain our contingent knowledge of reality.

What is the relation of thought (as an obviously "higher" form of evolutionary adaptation) to the other activity of the CNS, specifically to the unconscious and unverbalized aspects of human (and animal) neural functioning? Is 'conscious' thought, as it occurs in words, the only symbolic activity of the nervous system which may be said to model reality, or is there other evidence for modeling? If symbolism is defined in a very general way, as the ability of processes to parallel or imitate each other, then it becomes clear that other aspects of neural activity are equally as symbolic as thought. Indeed, *modeling is the fundamental function of the central nervous system* (of all species possessing such nervous systems), and it is exemplified in every instance of assimilation and accommodation that occurs in the life of an organism. The non-human cases are of only academic interest (for our purposes) and need not be pursued. But the case of *unconscious* modeling is worthy of note, for it makes the point that virtually all neural activity is symbolic and that words and conscious awareness are inessential to this mode of functioning.

One of the simplest appearing neural activities is the phenomenon of habituation. When an organism is exposed to stimulation for an extended duration, its reaction (in terms of its neural responsivity) diminishes. At first it was natural to assume that this phenomenon of habituation to stimulation was the result of "inhibition" that simply made the nervous system less sensitive to input: i.e., stimulation was said to decrease sensitivity if it endured over time, as if the system had a refractory period to overcome. But the lack of reaction to continuing *constant* stimulation is only part of the story of habituation. If the stimulus *changes* slightly, all the alerting and orienting reactions that accompany the onset of *novel* stimulation occur. The work on physiological correlates to stimulation is primarily Russian, stemming from the continued interest in Pavlovian conditioning. The demonstration of the paradoxical nature of habituation, that it *cannot be due to decreased sensitivity*, is largely the work of E. N. Sokolov.[22] Sokolov's findings indicate that what is going on in the cortex during habituation is that the present stimulus input is being matched to a standard that represents prior stimulation. That is, the current input is compared with an internal standard, or model, that is the result of prior stimulation. If the match of sample to standard obtains, "habituation" results; if the match does not obtain, if the new input does not

[22] See R. Lynn, *Attention, Arousal and the Orientation Reaction* (New York: Pergamon Press, 1966); E. Sokolov, "Neuronal Models and the Orienting Reflex," in M. A. B. Brazier, ed., *The Central Nervous System and Behavior* (New York: Josiah Macy, Jr., Foundation, 1960), pp. 197–276; also the references in Karl Pribram, "The Brain," *Contemporary Psychology*, September 1971.

match the stored or *expected* representation, then orienting and investigatory reactions occur (unless the stimulation is extremely strong: then a defense reaction occurs). But the point to note is that habituation, one of the simplest and most ubiquitous phenomena of neural activity, is not just a passive phenomenon. Rather, as Karl Pribram[23] has said, "habituation thus does not indicate loss of sensitivity by the nervous system; it shows that the brain develops a neural model of the environment, a representation, an expectancy against which inputs are constantly matched" (p. 46). Modeling, then, is not just the fundamental function of thought — it appears to be the fundamental function of the nervous system in general.

Now this is not an overstatement: modeling *is* really *the* fundamental function of the CNS. What the CNS "does" is to pick up information and then act upon it. And it is easy to see that all "picking up" of information requires the system to make a choice, to effect a classification, to determine whether a present level of sensory stimulation is to count as input or not. This is because the living organism never is or inhabits a static system in which no "stimulation" occurs — all input occurs against a background of ongoing neural activity. The sensory apparatus of living organisms has neural activity occurring *all* the time. As long as this is so, then the habituation situation is always present: in order to code *anything* as an input, the nervous system must compare it to an internalized model of *what input is*. Surprisingly, the nervous system cannot know what in the flux of its milieu is an input unless it has some internal standard or criterion of what constitutes an input. Thus there are two types of classification inherent in each of our sensory impressions: first, the determination that an ongoing level of activity constitutes an input; second, simultaneous to that determination the input 'event' is classified as instantiating a particular *kind* of sensory event. Pribram's comment cannot be ignored: the brain develops a neural model of the environment, a representation, an expectancy against which inputs are constantly matched. But to this picture must be added the startling rider that the nervous system *creates* its own inputs.

The sensory order. The fact that it is the structure and functioning of the nervous system that create the richness of sensory experience, rather than the environment external to the nervous system, is not

[23] In his article "The Brain." See also his book *Languages of the Brain* (Englewood Cliffs, N.J.: Prentice-Hall, 1971).

often recognized. That this must be so was clearly stated in a fascinating and neglected monograph by F. A. Hayek.[24] Hayek's psychological theory is concerned with sensory *perception* and the physiological correlates of our psychological abilities. The fundamental thesis Hayek advances is that no sensory input is "perceived" (i.e., inputed through the active CNS) *at all* unless it is perceived as one of the kinds of input accepted by the (innate or learned) classes of sensory order. As noted in the discussion of habituation, sensory perception is always an act of classification; the input signal is "processed" (matched to a standard) by any member (to which it "keys") of the sensory "orders" which impart to the phenomenal event the intrinsic properties that we experience. No sensory input can be perceived unless it can be isomorphically accepted as a match by the classes of sensory order. That is, no objects or constructions of phenomenal existents are possible except in terms of the (prior) apparatus of classification inherent in the operation of the functional nervous system. In other words, unless an environmental "event" gives rise to a pattern of sensory input that fits the organism's preexisting (either innate or learned) system of natural-kind classification, it is not perceived *at all. Perception is thus never of the intrinsic properties or attributes of "objects" in the real world*: instead, they (objects) are the results, the *abstractions*, of the actual organization and *memory* of the central nervous system. According to Hayek,

> The point on which [this] theory of the determination of mental qualities . . . differs from the position taken by practically all current psychological theories is thus the contention that the sensory (or other mental) qualities are not in some manner originally attached to, or an original attribute of, the individual physiological impulses, but that the whole of these qualities is determined by the system of connections by which the impulses can be transmitted from neuron to neuron: that it is thus the position of the individual impulse or group of impulses in the whole system of such connexions which gives it its distinctive quality; that this system of connexions is acquired in the course of the development of the species and the individual by a kind of "experience" or "learning"; and that it reproduces therefore at every stage of its development certain relationships existing in the physical environment between the stimuli evoking the impulses. . . . This central conten-

[24] Hayek's psychological theorizing is found primarily in *The Sensory Order* (London: Routledge & Kegan Paul, 1952; reprint, 1963, Phoenix Science Series). A more recent statement is found in "The Primacy of the Abstract," in A. Koestler and J. R. Smythies, eds., *Beyond Reductionism* (London: Macmillan, 1969), pp. 309–33.

tion may also be expressed more briefly by saying that "we do not first have sensations which are then preserved by memory, but it is as a result of physiological memory that the physiological impulses are converted into sensations. The connexions between the physiological elements are thus the primary phenomenon which creates the mental phenomena" (1952, p. 53).

Thus Hayek's thesis is that an "event" (an external physical energy source) is not coded at all unless it is assimilated to a system of classification that already exists within the CNS. In a later paper Hayek makes clear the implications of this thesis for the organismic basis of our knowledge: "What this amounts to is that all the 'knowledge' of the external world which . . . an organism possesses consists in the action patterns which the stimuli tend to evoke, or, with special reference to the human mind, that what we call knowledge is primarily a system of rules of action assisted and modified by rules indicating equivalences or differences or various combinations of stimuli" (1969, p. 316).

Thus there is a circular interdetermination at work in the functioning of the CNS as it gains its knowledge of the environment. The primary mode of functioning of the nervous system is modeling: the input to the system is compared to an internalized standard, and the action patterns that result have their basis in this information. But not only does the nervous system compare its inputs to a standard, it also creates them. We have no direct commerce with external reality at all: every input that (causally) results from an external energy source reflects the intrinsic properties of the CNS rather than the intrinsic properties of the external source. This is, of course, the physiological correlate to the distinction of knowledge by description as opposed to acquaintance. As Russell's doctrine so clearly states, all our knowledge of phenomena external to the CNS is of the structural properties of the phenomena: we have no knowledge of the intrinsic or "first-order" properties of external objects at all. Our knowledge by acquaintance, which is knowledge of the intrinsic properties of the nervous system, although causally related to external events, is the result of the nature and functioning of that nervous system rather than the external events themselves. The only modeling of external, nonmental reality that thought can accomplish is of its structural properties.[25]

[25] But is the structure of reality, as it is disclosed by our best contingent (i.e.,

Paradoxically, the upshot of this is that we cannot regard the phenomenal world of immediate experiences as in any sense more 'real' or 'fundamental' than the world of science. This is so because *both* the phenomenal world and the scientific world are *equally* constructions in the mind of man. Indeed, both 'worlds' are the result of the primary activity of the nervous system: classification. But no purely 'phenomenalistic' interpretation of science can suffice: this is because the scientific image constantly contradicts the manifest phenomenal image. Phenomena per se are not subject to completely regular (i.e., invariant) classification — "Our knowledge of the phenomenal world raises problems which can be answered only by altering the picture which our senses give us of that world" (Hayek, 1952, p. 173). The primacy of what Sellars[26] has called the 'scientific image' over the 'manifest image'

scientific) theories, *really* a reflection of the structure of reality, or is it merely the result of our cognitive structure? Considering only the nature and functioning of the nervous system per se, the answer is very likely the latter alternative. For the nervous system is essentially a highly skilled *instrument of classification*: as noted in the discussion of habituation, neural activity must *judge* whether or not other neural activity (which we can only *assume* has its source in 'external' origins) is to count as input. That is, the classes of sensory order determine what counts as stimulation. It follows that if there were a stimulus which was not regular, i.e., not an instance of one of the appropriate kinds, we could not 'know' anything about it. The key concept here is *regularity*: what sensory perception can accept as input can never be unique properties of individual objects, but must always be properties which the objects have in common with other objects. The fundamental activity of the nervous system is classification. Perception is *always* classification, which is to say interpretation: interpretation of something (either 'object' or 'pattern of neural activity') as belonging to one or another class (of 'objects' or 'patterns of neural activity'). Thus, to reinforce again Russell's point, the qualities which we attribute to experienced external objects are not properties of those objects at all, but rather a set of *relations* by which our nervous system effects their classification. This is what knowledge by description means: all we can perceive of external events is their structural relation to each other and to our experience. We truly are theories of our environment: *all* we can know about the world is inherently theoretical (there is no direct awareness or factual bedrock), and all our 'experience' can do is modify our theories!

Return now to regularity. The line of reasoning just outlined indicates that we can 'know' (i.e., classify) only those kinds of events which show a degree of regularity in their occurrence in relation to other events. We could not know (classify) events which occurred in a completely irregular manner. The world as *we know it* remains a construction of the mind of man: "The fact that the world which we know seems wholly an orderly world may thus be merely a result of the method by which we perceive it" (Hayek, 1952, p. 176). This much of the thesis of idealism, it seems to me, is incontestable (if the arguments upon which they are based, which are really contingent theories, are sound). Empiricist theories of the mind are rendered so absurd that only fools or philosophers could have propounded them.

[26] See his essay "Philosophy and the Scientific Image of Man," in *Science, Perception, and Reality*, and *Science and Metaphysics* (London: Routledge & Kegan Paul, 1968). See also B. Aune, *Knowledge, Mind and Nature.*

is that science is better at the task of modeling reality — of approaching more closely the *reproduction* of the objective order of events — than is sensory experience. The phenomenal order, which reflects the (human) organism's initial theorizing or modeling of his environment, is at best a (very) rough approximation of the order of reality. The constructions of scientific thought are, *all* things considered, far superior in their degree of adequacy of representation of reality.

Parenthetically, this casts further light on the relation of consciousness to preconscious sensory perception. In both cases, the function of the nervous system is to model reality. But the model formed by the preconscious senses will often prove to be inadequate in that it will lead to falsified expectations. Thus, although the conscious, self-reflecting mind can know external reality only in terms of the classes that preconscious sensory perception has created, the 'data' of sensory experience form a basis for their own revision (i.e., reclassification). The conscious mind will reclassify the initial sensory experiences in order to model the structure of reality more adequately. "The experience that objects which individually appear as alike to our sense will not always behave in the same manner in relation to other classes of apparently similar objects, and that objects which to our senses appear to be different may in all other respects prove to behave in the same manner, will thus lead to the formation of new classes which will be determined by explicitly (consciously) known relations between their respective elements" (*ibid.*, p. 145). *If* anything can represent reality adequately, it will be the conscious mind, through the construction of scientific theories, rather than the preconscious sensory order with its initial classification of inputs.[27]

So the argument of classical idealism, that the world is a construction in the mind of man, leads from idealism to scientific or representational realism (specifically, to structural realism). If anything is going

[27] That is, ideally completed, Utopian science will be the most adequate representation of reality that human beings can aspire to. Such an ideal scientific account will, of course, have to make explicit the enormous amount of implicit knowledge (which Polanyi and Kuhn aptly term "tacit knowledge") that we all possess as highly evolved theories of our environment. The conscious activities of scientists are based upon almost wholly unconscious processes and principles of determination, and "science" must make explicit those processes and principles before its claim to adequacy can be taken seriously. But this is tangential to the main point in the text above: reflective thought refines and corrects the initial classifications effected by the preconscious sensory order.

to give us an adequate model of reality, it will be the scientific model rather than the manifest, sensory-based model. Since both images are on a par with equally 'infirm' foundations, i.e., since neither is in more direct contact with reality than the other, the choice between them must be made upon their relative efficiency and adequacy of performance upon their common task: modeling reality. The scientific image wins this contest hands down: science can beat the manifest image *at its own game.*

But the skeptic's question may now be repeated: "How do we know that there *is* an objective world, i.e., a reality external to the constructions of our own minds?" The answer depends upon the point of view adopted. Within the justificationist framework, the answer is that we *cannot* know there is an objective world, for we cannot *prove* its existence (or even render it probable). Further, it is not a question for logic alone to arbitrate. Nor can sense experience, the ultimate epistemological authority of empiricism, be other than impotent with regard to this issue.[28] Intellectualism fares no better: there are no a priori arguments that are compelling of a choice for realism or idealism. Solipsism, as Russell aptly remarked, is logically impeccable; it is only *psychologically* intolerable.

The skeptic's question can be answered by reflecting upon how we 'know' anything (other than formal systems) in a nonjustificational metatheoretical framework. For the nonjustificationist, knowledge is a matter of having warranted assertions, i.e., conjectural claims that can be defended by 'good reasons'. And 'good reasons' will consist of other knowledge claims that have been corroborated (as opposed to 'verified' or 'confirmed'). Such knowledge will be rational insofar as it is sub-

[28] Empiricism, pursued conscientiously to its inevitable conclusion, leads to *conceptual a priorism.* Hayek puts the self-stultifying nature of sensationalistic empiricism beautifully: "Precisely because all our knowledge, including the initial order of our different sensory experiences of the world, is due to experience, it must contain elements which cannot be contradicted by experience. It must always refer to classes of elements which are defined by certain relations to other elements, and it is valid only on the assumption that these relations actually exist. Generalization based on experience must refer to classes of objects or events and can have relevance to the world only in so far as these classes are regarded as given irrespective of the statement itself. Sensory experience presupposes, therefore, an order of experienced objects which precedes that experience and which cannot be contradicted by it, though it is itself due to other, earlier experience" (*ibid.*, p. 172).

Note that this reasoning becomes the strongest of possible *contingent* supports for realism rather than idealism: but note also that realism is *not* empiricism!

ject to criticism, where criticism is no longer fused with the attempt to prove or justify. Clearly there can be no justification of such knowledge claims, for there is no ultimate epistemological authority upon which to found them, i.e., ground them in certainty. Within a nonjustificational philosophy the problem of the justification of nondemonstrative inference (as genuine or rational knowledge) does not arise: since *all* knowledge claims are equally inferential rather than certain, equally nondemonstrative, the quest of justification is seen to be misguided. Such a quest can arise only if the ideal of Euclidean systems is taken to constitute genuine knowledge, and all other forms of knowledge (such as those of contingent science) are considered pale approximations to this genuine form. The dilemma the justificationist has created for himself, by defining his concepts in such a manner that genuinely rational contingent knowledge cannot be shown to exist, can be avoided only by stepping outside the justificationist metatheoretical framework.

The existence of an external nonmental world that is causally responsible for our perceptions of it in the mental world of direct acquaintance is a scientific claim, and it must be defended by adducing 'good reasons' for its truth (in exactly the manner that any other contingent proposition(s) must be defended). The good reasons that must be delivered (sooner or later) are an adequate theory of how the mind works which indicates that the mind could not work the way it in fact does unless there was a world external to our senses, a world having the properties that Utopian science attributes to it. To repeat, solipsism is only *psychologically* intolerable. That is, psychology must ultimately deliver a theory of the mind which delimits the a priori constraints upon human mental functioning (the scientifically determined Kantian categories of the understanding, if you will), and an indispensable ingredient of that theory must be a specification of how the mind could not function as it does indeed function without an external world. Nothing less will suffice for determining the warranted assertability, i.e., the *truth*, of this particular empirical claim. But one must never forget that explanations, even if true, are never justifications.

The a priori constraints upon human knowledge. Thus far we have been sketching (synthetic or 'scientific') a priori constraints upon the nature of our knowledge: by considering the nature of our knowledge and our biological nature, we have been adducing facts that a psychological theory of knowledge acquisition must be consonant with. That is, we have described facets of our knowledge and our thought that *any* theory of the acquisition of knowledge must acknowledge. There is one further such constraint that is important enough to warrant special mention.

Since our knowledge of reality is a product of the classificatory ability of the CNS, it is clear that a principle of causality is a presupposition of (the *acquisition* of) human knowledge. Classification must presuppose causal regularity. Further, perception requires for its very definition a concept of causality. Russell was very clear on this: "The conception of 'causal lines' is involved not only in the quasi-permanence of things and persons but also in the definition of 'perception'. When I see a number of stars, each produces its separate effect on my retina, which it can only do by means of a causal line extending over the intermediate space. . . . Generally, what is said to be perceived, in the kind of experience called a 'perception', is the first term in a causal line that ends at a sense organ" (1948, pp. 458–59). Science presupposes the causal theory of perception: human knowledge could not be what it in fact is unless the existence of what Russell aptly called "causal lines" were presupposed. "But for causal interconnectedness, what happens in one place would afford no indication of what has happened in another, and my experiences would tell me nothing of events outside my own biography" (*ibid.*, p. 162).

However, neither experience nor logic can prove the necessity or even the existence of causal lines. Yet without them, human knowledge and science as we know them to be would be impossible to achieve. For the 'scientific' justificationist, this is most perplexing. Russell had to make a *postulate of scientific inference* out of his notion of causal lines. He saw it as the only way to account for our knowledge transcending empirical particulars: such knowledge cannot be wholly based on experience. Speaking of the universal knowledge that science claims to know, Russell the justificationist asks *the* question: "But we most certainly do need *some* universal proposition or propositions, whether the five canons suggested in an earlier chapter or something different. And whatever these principles of inference may be, they certainly cannot be logically deduced from facts of experience. Either, therefore, we know something independently of experience, or science is moonshine" (*ibid.*, p. 505). In this he is caught up in the self-stultifying justificationist quest, and even his "postulates" cannot stave off the verdict that science is moonshine *on justificationist criteria*. But Russell had the right approach to inference and expectation — he considered them as biologically based mechanisms of adaptation given the world as it is disclosed by science. "The forming of inferential habits which lead to true expectations is part of the adaptation to the environment upon which biological survival depends" (*ibid.*, p. 507).

What we shall achieve, by diligently studying such inferences, is not knowledge of any postulates of scientific inference that are beyond justification, but rather a theory of the a priori constraints that give human knowledge the character that it does in fact have. What we shall wind up with is a synthetic or scientific theory — a fallible, conjectural one — about the nature, scope, and *limits* of human knowledge. What we need to know about scientific (and other) inference is not how to justify it but rather how it occurs and what constrains it. And it is precisely this sort of knowledge that is, at least in principle, available from the study of the inferring organism — from the psychology of inference and expectation.

The Acquisition of Human Knowledge

To the question of "How do we model reality?" the obvious answer is, "By forming concepts (of its structural properties)." To the question of "How is our knowledge represented?" the obvious answer is, "By our concepts." To the question of "What constrains the patterning of our inferences?" the obvious answer is, "The concepts that constitute our present knowledge." To the question "What is the end product of inference?" the obvious answer is, "Concepts." For these and similar reasons, explicating the nature of concept formation is perhaps the most important task faced by psychology. Although our knowledge of concept formation is quite sketchy and incomplete, we do know a number of things that cast light upon the problems of inference and expectation.

The most important thing we know about concepts is that their acquisition involves the interplay between experienced particulars and abstract rules. The concept literally is the interplay between instances and their rules of determination. All our knowledge consists in the repeated restructuring of classifications effected by the CNS: concepts are these classifications and their patterns of restructuring. The experienced particulars, i.e., the initial sensory classifications effected by the CNS, are alone not our concepts — our experience is not (the same as) our knowledge. *All* knowledge transcends experience: our concepts are a means by which we generate potentially infinite domains of experiences. The near infinite richness of content of the experienced particular is not the starting point from which the mind forms its concepts, but rather the *product* of a fantastic range of restructurings or classifications that the CNS effects. Human knowledge consists of rules of determination by means of which we construct the richness of the particular — all our knowledge is, as Hayek said, a system of rules of determination (or action patterns) which is evoked by the input to the CNS. What we learn in forming a new concept is essentially a new pattern or configuration (of experienced particulars) according to new rules of determination. Thus to *know* is always to be in a position to *infer*: to know is to generate an indefinitely extended domain of potential data points. Our concepts are inevitably inference tickets: to know a particular is always to be prepared to infer to an indefinite number of other particulars. What we know is inextricably bound up with what we expect. This being so, the answer to the question "What guides scientific life?" is obviously "What we (think we) know." But what *do* we know, and *how* do we know it?

Conceptual abstraction. The central problem for the 'psychology of inference' is to characterize *what is learned* and *how our memory utilizes prior knowledge*. If we can get clear on what is learned, or what is to say the same thing, *what our knowledge is*, then we will be in a position to discuss the role of experience in the determination of human knowledge and to see why induction-from-experienced-particulars is an inadequate theory of the nature of knowledge and its acquisition. Furthermore, if we can get clear on *how memory utilizes knowledge* in the formation of new concepts, then the 'guidance of life' problem will automatically be answered.

Most simply stated, the problem is to explicate how organisms classify particular instances as instances of thing-kind categories. For example, how do you know that $\triangle$ is a triangle? What is involved in learning that particular triangles are all instances of 'triangularity'? The most crucial thing to note is that it is impossible to exhaust the meaning of 'triangularity' with any list, no matter how long, of physical (referential) attributes: there are an infinitude of things that are all equally triangles. No induction-by-enumeration of referents theory can teach an organism the concept. The number of experienced particulars required to learn "what triangles are" (that is, that certain figures instantiate triangularity) would be indefinitely large. Generic concept formation is essentially 'productive' or 'creative' as the linguist uses the term: to know a concept is to be able to apply it appropriately to a totally novel instance, one that need not ever have occurred before in the history of the organism, or indeed the history of the world. The question comes to this: How can an organism recognize all the potential instances, on the basis of *no* prior exposure to them, as instances of *the same* concept? How can the organism construct a theory of what proper subset of its sensory presentations are instances of triangles, *and* then extrapolate from that given corpus of presentations the defining characteristics of 'triangularity' which would enable it to recognize *any* instance as a triangle?

One thing that is very clear is that an organism which has learned the concept of 'triangularity' has learned the *rules of determination* that enable the construction of triangles. Equally clear is the fact that those rules of determination must range over abstract entities (i.e., structures not found in the experienced surface structures of particular triangles). The only way known for organisms to "make infinite use of finite means" as they do in generic concept formation is to employ "grammars" of perception (or cognition, or behavior) that make use of abstract entities in their rules. And only grammars that allow indefinite recursion and which employ nonterminal symbols (for abstract entities) in their derivations of surface structures from underlying or 'deep' structures can do this. What these grammars show is how an abstract, underlying deep structural "meaning" (intension) can be mapped into indefinitely many distinct surface structure representations. A generic concept is the same as an underlying deep structure that can be characterized *only by the rules of determination which are its grammar*, rather than by listing the attributes or "experienced features" that are present in its potentially infinite number of surface structure instantiations. Indeed, your knowledge of the concept of triangularity is independent of any surface structure experiential representation: the concept is defined by purely "verbal" statements that characterize its rules of determination.

Experience per se, as in 'seeing' a triangle, is not the meaning of the concept.[29]

What is learned in concept formation, our knowledge of concepts, is the rules of determination that constitute the invariant relations in a group. Concepts are not copies or representations of particulars, they are the rules by which we construct particulars as instances of thing-kind classes. Our perceptual knowledge of reality seems to be based upon the possession of "rules of seeing" that determine the invariants in our environment and the group of *transformations* applicable to them. But no sensory experience alone is identical with *any* such concept: what the concept *is* consists of *both* the underlying rules of determination and its surface structure representation. That is, perceptual concepts are both deep *and* surface structures, and not either alone.

But there are other instances of concepts that are found in our knowledge which are not 'perceptual' the way 'triangularity' is and which are still contingent rather than a priori. How do we know, for example, that "all men are mortal" or that, as Newton proclaimed, "$F = ma$"? Predicates such as 'mortal', concepts such as force (or freedom), are experiential in the sense that our experiences are relevant to their determination, but they are not perceptual like 'triangular'. How do we recognize men, e.g., John Smith, *as* mortal? If recognition is restricted to sensory perception, to an iconic picturing notion, then it is obvious that we do *not* recognize men as mortal: mortality is not the sort of

[29] This was clear to philosophers such as Ernst Cassirer as far back as 1910: "The content of the concept cannot be dissolved into the elements of its extension, because the two do not lie on the same plane but belong in principle to different dimensions. The meaning of the *law* that connects the individual members is not to be exhausted by the enumeration of any number of instances of the law; for such enumeration lacks the generating *principle* that enables us to connect the individual members into a functional whole. If I know the relation according to which *a b* c . . . are ordered, I can deduce them by reflection and isolate them as objects of thought; it is impossible, on the other hand, to discover the special character of the connecting relation from the mere juxtaposition of a, *b*, c in presentation. . . . The unity of the conceptual content can thus be 'abstracted' out of the particular elements of its extension only in the sense that it is in connection with them that we become conscious of the specific rule, according to which they are related; but not in the sense that we construct this rule out of them through either bare summation or neglect of parts." E. Cassirer, *Substance and Function and Einstein's Theory of Relativity* (New York: Dover Publications, reprint, 1953), pp. 26, 17.

Cassirer's insight was lost to psychology during the fifty-odd year blight of behaviorism. Fortunately cognitive psychology is beginning to recover from behavioristic learning theory and its induction-by-enumeration model of knowledge acquisition; see U. Neisser, *Cognitive Psychology* (New York: Appleton-Century-Crofts, 1967).

thing that can be observed or photographed. Yet we do come to know that men are mortal, and that if John Smith is a man, he is mortal also. And the way in which we do this is not very different from the way in which we recognize a particular triangle as instantiating triangularity. In both cases a particular of a specific sort is classified as instantiating a generic relation that is not itself *any* experienced particular. Mortality as a concept is intimately linked, in our conceptual space, with the concept of man, and what we know in "John Smith will die" *is* that linkage. Our knowledge of such nonperceptual concepts is represented in some nonperceptual manner, specifically in whatever manner underlies our use of language. Whether this manner of representation is construed to be syntactic or semantic (there is considerable debate among linguists and psycholinguists here), it is clear that it is *abstract*, for no experiences that we can undergo are equivalent to it. We possess a conceptual schema that allows us to recognize (or perceive) instances of perceptual or linguistic formulations as relevant to that schema, but those latter formulations are *not* the schema itself. What we learn in learning such abstract concepts is to follow underlying rules of determination that can generate an indefinite number of surface structure manifestations.

Very similar considerations apply to the learning of consequential knowledge of nature, such as Newton's law $F = ma$. What we learn in doing science is how the basic law sketch $F = ma$ may be transformed and restructured in disparate instances of F's equaling ma's. We come to *see* (in the conceptual sense) that disparate surface structure manifestations share a common, underlying structural regularity that we may represent by the equation $F = ma$. The essence of concept formation is the abstraction of such rules of determination that generate surface structure instantiations. What the mind does is *construct* its concepts upon the basis of incomplete and often quite impoverished sensory experience. The psychology of concept formation requires a *constructive* theory of the mind rather than either an inductive or a deductive one.

Memory. In a sense, the heart of a constructive approach to concept formation and knowledge acquisition is a generative model of memory. Everyone knows that what constrains our inferences (or constructions) is, in some loose sense, our 'background knowledge', but virtually no philosophers have concerned themselves with characterizing the nature of human background knowledge. Thus a few words about the manner in which memory operates to provide that 'background knowledge' are appropriate. Once it is seen that memory does not "store" *particulars*, and that indeed the mind does not really "store" information *at all* (at least in anything like the customary sense), the appeal of an induction-from-particulars model of scientific concept formation vanishes.

The problem of prose recall is characteristic of the operation of memory in *meaningful* contexts (such as scientific concept formation, etc.). People just don't bother to memorize by rote in such situations: presented with an utterance pertaining to the "eminent arrival of the magistrate," they will remember that the gist of what was said was "Here come de judge." The latter utterance does contain the 'gist' of the former — it is indeed an adequate paraphrase — and yet there need not be a single 'surface structure' word shared in common by them. What the mind does with meaningful information is to represent it in terms of a semantically coherent schema, from which it can generate appropriate surface structure representations whenever the need arises. We do not store the given words of a sentence or even necessarily the paraphrase of a single sentence. Rather, what the head does with information is to *use* it, to construct an appropriate schema that can be used as the basis of future expectation. It is becoming quite clear from psychological research that sentences (or whatever is taken to be the appropriate unit in meaningful communication) do not 'carry' meanings in their surface structure manifestations — rather, they are 'triggers' to meanings that are already present in the (recipient) head. The structural representations that constitute human "understanding" are vastly richer than mere "words."[30]

[30] Cognitive psychology is beginning to acknowledge that comprehension involves more than the registration of words in sequence. For example, when presented with a series of "related" sentences, one usually "remembers" the 'gist' of the series as one syntactically complex but semantically integrated unit. Indeed confidence ratings indicate that we will be *more* confident that what was actually presented was the integrated unit than the actual simple sentences that were experienced. The semantic information is fitted into a plausible context by the hearer, and once the information is assimilated to such a schema the actual input is disregarded (see J. D. Bransford and J. J. Franks, "The Abstraction of Linguistic Ideas," *Cognitive Psychology*, 2 (1971), 331–50). If we can construct a meaningful context, virtually anything can be remembered, or judged comprehensible, etc. For example, paragraph passages that would be judged incomprehensible and all but impossible to recall by themselves can be recalled very well and are quite comprehensible when heard after exposure to a drawing which provides a context that is appropriate. J. D. Bransford and M. K. Johnson, Contextual Prerequisites for Understanding: Some Investigations of Comprehension and Recall," *Journal of Verbal Learning and Verbal Behavior*, 11 (1972),

To say that the head uses information to abstract "consequences" that are relevant to its expectations is not to say that the head stores information. The scientist's internal model of reality — his theory — is an abstract schema that is constantly being updated as the mind structures and restructures the information that sensory classification makes available. In modeling reality the mind functions as a generative concept forming device — it abstracts consequences from the input made available to it by relating particular surface structures to the sets of abstract rules of determination that constitute its "knowledge" or "memory." What human conceptual knowledge is, as Hayek said, is a set of patterns of action subsisting within the functioning CNS.

The upshot of this is clear: we should cease talking about science "learning by induction" and try to understand how the mind constructs its knowledge from its (tremendously impoverished) input. Experience per se, the perceived particulars determined by our initial sensory classifications, is not equivalent to our knowledge. *What we learn isn't facts* — instead we learn concepts which enable us to generate facts. Our concepts function as licenses to infer particulars, and those concepts are not just other particulars. What we remember is a conceptual schema, and experiences are only relevant to conceptualization insofar as they can be fitted into those conceptual schemata.

The role of experienced particulars: Learning from exemplars. The primary knowledge that we gain from the practice of science consists in the rules of determination that are the skeletal deep structure representation, the schema, of our concepts. Many concepts, such as energy, force, reinforcement, memory, evolution, etc., are purely conceptual rather than perceptual in the sense that, although sensory experience pertains to our usage of them, no sensory experience may be identified as a surface structure manifestation of the concept. What then is the role of experienced particulars in such abstract concept formation? The answer rein-

717–26. What we do with information is to use it, and that means to assimilate it to the knowledge structures we already possess and to the expectations we have for the future. If we can construct an appropriate context, we can understand or remember virtually anything; if we have no context, even information that is normally intelligible by itself is not understood or remembered. J. D. Bransford and N. S. McCarrell, "A Sketch of a Cognitive Approach to Comprehension: Some Thoughts about Understanding What It Means to Comprehend," in W. B. Weimer and D. S. Palermo, eds., *Cognition and the Symbolic Processes* (Washington, D.C.: Erlbaum Associates, 1974), pp. 189–229; J. J. Franks, "Toward Understanding Understanding," in *ibid.*, pp. 231–61.

forces the centrality of perception, as a *higher* mental process, in theory construction and knowledge acquisition.

What we learn from particular exemplars, or from exemplary puzzle solutions, has been well characterized by Kuhn: we *learn to* see disparate surface structure manifestations as similar in crucial respects to others. We learn to see particulars as instantiating the rules of determination that constitute the deep structural schema of a concept. By working with disparate surface structure configurations to the extent that we can see *their resemblance* to others, we tacitly become aware of the deep structural communality that underlies the particulars that we are familiar with and that can generate an indefinite range of further particulars. Scientific theories, as structural representations of reality, are ways of seeing that deep structure. Experienced particulars stand to that structure as instances stand to abstract rules of determination. What we learn in "doing" science are ways of perceiving. What we learn to see in "facts" are similarity relationships: we learn to see instances as instances of thing-kinds, and we learn how thing-kinds are structurally related. It is this perception of similarity relationships in the structural properties of reality that led Hanson to call theories *patterns*: they are patterns of constructions, according to which we come to conceive of reality. The interesting problems of inference and expectation that psychology must in the future address concern understanding the factors, such as our implicit background knowledge, that *constrain* the patterns of inference that we entertain.

To date, all that the historian of science or psychologist of knowledge acquisition can do is single out instances of "learning to see" similarity relationships in disparate structures. Thomas Kuhn has recently done this,[31] and it is not surprising that he has now held historical research in abeyance in order to study concept formation. And it is also not surprising that the psychology of perception is central to Kuhn's work: what the student of science learns in doing the exemplary problems of his discipline is a group-licensed way of seeing the interrelationships between various surface structures. Science is a matter of coming to recognize particular surface structure manifestations as being related to underlying deep structural rules of determination. And the purely structural knowledge of the underlying rules of determination constitutes the

[31] See "Second Thoughts on Paradigms," in F. Suppe, ed., *The Structure of Scientific Theories* (Urbana: University of Illinois Press, 1974), pp. 459–82.

knowledge by description which is our only knowledge of the nonmental realm.

Let us end this section with a slight digression. We are now in a position to understand a "fact of scientific life" that is all but inconceivable to the justificationist: that no scientist was ever persuaded to or dissuaded from a position by the "facts." It does not take much to see that Fries was correct, that science can neither prove nor disprove — contingent theories are as metaphysical in this regard as the wooliest of religious dogmas. But this has shocked the justificationist, because on his account *all that science can learn from experience is that a theoretical proposition is true or false*. Thus on justificationist criteria, science does not profit from, or make use of, experience.[32] This is but one more reason why science is irrational according to justificationist criteria.

But we can admit that science does not "induce" propositions from factual "experiences" without claiming that experience plays *no* role in scientific practice. Rather, we can relocate the role of experience in our learning from exemplars. What experienced particulars provide for the scientist is practice in perceiving — in perceiving instances as instances of thing-kind classifications specified by abstract rules of determination. In that sense experience remains the basis for our conceptual formulations, but the farther science progresses toward its quest for an adequate structural representation of reality, the sharper is the break between experience or knowledge by acquaintance and conceptual representation.[33] It is, as Hayek noted, precisely because all our knowledge is rooted

[32] Feyerabend loves to flaunt this in the face of the beleaguered justificationist: see his appendix "Science without Experience" in "Against Method," in M. Radner and S. Winokur, eds., *Minnesota Studies in the Philosophy of Science*, vol. 4 (Minneapolis: University of Minnesota Press, 1970), pp. 17–130.

[33] Cassirer put this beautifully when he wrote: "The goal of theoretical physics is and remains the universal laws of process. The particular cases, in so far as they are taken into account, serve only as paradigms, in which these laws are represented and illustrated. The further this scientific problem is followed, the sharper the separation becomes between the sytsem of our *concepts* and the system of the *real*. For all 'reality' is offered to us in individual shape and form, and thus in a vast manifold of particular features, while all *conception*, according to its function, turns aside from this concrete totality of particular features. . . . The direction of thought upon the 'concept,' and its direction upon the real, mutually exclude each other. For to the extent that the concept progressively fulfills its task, the field of perceptible facts recedes. . . . The final goal of the material sciences and of all other natural sciences is to remove empirical intuition from the content of their concepts. Science does not *bridge* the gap between "thoughts" and "facts," but it is science, which first creates this gap and constantly increases it" (1953, pp. 220, 221).

As Körner has indicated, the disconnection of scientific concepts from perceptual in experience that our principles of conceptual determination must transcend experience. This being so, what we can *learn* from experience cannot be what the justificationist desired; we are instead forced to regard experience as an incomplete surface structure that is compatible with, but not sufficient to lead unambiguously to, its underlying purely conceptual rules of determination. No scientist ever learned "the truth" from "facts," any more than he accepted a theory on the evidence of "facts": yet this does not mean that experience plays no role in scientific concept formation.

But the point of this section remains: the role of experienced particulars in scientific concept formation is that of surface structure manifestations of abstract, underlying rules of determination. As such, *what* we learn from "factual experiences" is a way of seeing: a way of seeing certain experienced particulars as similar in some respect to others.[34] *How* we learn this from "facts" it remains for the psychology of concept formation to tell us.

Another Look at the Guidance of Inquiry

In unprofessional moments at least, virtually no one believes that scientists 'induce' theoretical formulations from observed instances, and although a few still believe in the rationality of concept formation according to the H-D methodology, no one has proposed an account of how one gets the 'theory' from which the 'deductions' occur *in the first*

experiences is total and complete. We never see mere particulars, only particulars as instances of thing-kinds. "To describe a group of phenomena, then, means not merely to record receptively the sensuous impressions received from it, but it means to transform them intellectually" (*ibid.*, p. 264).

[34] Once again it must be emphasized that science cannot deal with truly unique events: our classificatory schemata always assimilate particulars to abstract principles of determination, and the knowledge that results is purely structural. According to Cassirer, "No content of experience can ever appear as something absolutely strange; for even in making it a content of our thought, in setting it in spatial and temporal relations with other contents, we have thereby impressed it with the seal of our universal concepts of connection, in particular those of mathematical relations. The material of perception is not merely subsequently moulded into some conceptual form; but the thought of this form constitutes the necessary presupposition of being able to predicate any character of the matter itself, indeed, of being able to assert any concrete determination and predicates of it. Now it can no longer seem strange, that scientific physics, also, the further it seeks to penetrate into the "being" of its object only strikes new strata of numbers, as it were. It discovers no absolute metaphysical qualities; but it seeks to express the properties of the body or of the process it is investigating by taking up into its determination new 'parameters'" (*ibid.*, p. 150).

place. It is all very well to say that science is a matter of conjectures and refutations, as *both* Popper and his sophisticated justificationist antagonists agree upon, but that slogan is impotent without a psychology of theory formation.[35] 'Conjectures and refutations' is a fine methodology for a completed (i.e., a *dead*) science or for a mythical realm such as Plato's Third World, where discovery is *by definition* not a psychological issue. But since mortal human beings live (and practice science) in the real world rather than in the 'Third World', the recent attempts to preserve a purely *philosophically* rational (and logical) reconstruction of scientific growth in a logician's nirvana can be ignored. Logic is not the guide to scientific life, and there is no use killing science (and then mounting it in the Third World history museum) in order to make its growth seem purely logical.[36]

[35] For example, Salmon, *The Foundations of Scientific Inference*, attempting to revive a Reichenbachian pragmatic justification of induction and couple it with Bayes's theorem, argues for "plausibility considerations" as determinants of the prior probability of scientific hypotheses. His model of acceptance is basically induction *by elimination*, which he admits is impotent in the face of an unlimited supply of potential hypotheses: unless there is a principled way to eliminate a priori all but a few of them, elimination as a practical approach cannot get off the ground. "There are, as I have emphasized repeatedly, infinitely many possible hypotheses to handle any finite body of data, but it does not follow that there is any superabundance of *plausible* ones" (Salmon, p. 129). His next move is obvious: it is claimed that "plausibility arguments" suffice to weed out the indefinitely large class of potential hypotheses. "If we put plausibility arguments — perhaps I should say 'implausibility arguments' — to the purely negative task of disqualifying hypotheses with negligible prior probabilities, falsification or elimination becomes a practical approach. This is, it seems to me, the valid core of the time-honored method of induction by elimination" (*ibid*., p. 129).

But Salmon has no theory of 'plausibility' *at all*. To be sure, he lists, in sketchy fashion, three classes of characteristics that "may be used as a basis for plausibility judgments" (*ibid*., p. 125), but all such characteristics are impotent to determine plausibility in a given case except in an ad hoc manner. The prior probability of an hypothesis, regardless of whether it is equated with plausibility or not, is only determinable a posteriori. That is, the *true* prior probability of a scientific hypothesis is exactly the same as its posterior probability: as Fries indicated $P(h,e) = 0$ *always*. And precisely the same argument holds for prior and posterior plausibility considerations: all scientific hypotheses are as *implausible* as they are unprovable. If Salmon had any idea whatsoever how scientists determine the plausibility of hypotheses, he would ipso facto have a theory of scientific concept formation, and he would then have a genuine 'guide to life'. As it stands, he has no theory of guidance at all, only a relabeling of the problem as one of 'plausibility considerations'.

[36] This recent attempt to 'kill the patient' in order to preserve the appearance of the corpse, i.e., to render the growth of scientific concepts rational *from the logician's standpoint of rationality*, to implement a true *logic* of discovery, is due to Popper, *Objective Knowledge* (London: Oxford University Press, 1972); Lakatos, "Falsification and the Methodology of Scientific Research Programmes"; and A.

A number of philosophers and scientists have commented upon the futility of the traditional approaches to the 'guidance of life' problem. One persistent theme that emerges in their presentations is the centrality of what Polanyi has called *tacit knowledge*.[37] Polanyi's point is that we know far more than we can ever tell: our knowledge of reality is tacit in the sense that we could never 'formalize' or 'objectivize' it in the form of a scientific theory. "Tacit knowing is the way in which we are aware of neural processes in terms of perceived objects" (1966, p. x). His example of our ability to recognize physiognomic features is sufficient to make this point: we can *know as faces* (i.e., recognize) an indefinitely extended number of particular configurations, despite the fact that we have never seen them before, and no one can say how he can do this (indeed, there was no psychological theory that could address the issue until very recently). Our knowledge is in this sense vastly richer than our theories of our knowledge: both *what* and *how* we know are beyond the pale of traditional epistemologies. As Polanyi points out, our knowl-

Musgrave, "Impersonal Knowledge," Ph.D. thesis, University of London, 1969. The 'critical fallibilists', having spent three decades vanquishing the disillusioned justificationists by showing that the latter's theory of assessment (inductive logic) could not possibly explain science, are now facing the collapse of their own metatheory of scientific rationality. That is, the Popperians, having shown that 'instant' rational assessment was a chimerical quest, that the rationality of science does not lie in 'inductive logic' or 'confirmation theory', proposed instead that the rationality of science was to be found in its (logical) growth. But now it is obvious that deductive logic does not guide knowledge acquisition either — the guidance of life, and hence of scientific concept formation, is psychological rather than logical. The psychology of inference will (ultimately) explain scientific concept formation, but it will not reconstruct concept formation as a logic of discovery. There remain two alternatives for the concerned philosopher who wishes to understand science. One can, like Paul Feyerabend (see "Against Method" and "Consolations for the Specialist," in Lakatos and Musgrave, eds., *Criticism and the Growth of Knowledge*, pp. 197–230), maintain that science is irrational (since the critical fallibilist metatheory has collapsed) and let it go at that. Feyerabend is now 'gleefully' destroying the last vestiges of the Popperian rationality of science (with great success), but he appears to have no 'positive' program with which to replace it (other than an irrational appeal to 'love thy neighbor'). On the other hand, one can admit the failure of Popperian *and* justificationist rationality, and search for a psychology of inference and a sociology of knowledge transmission. This is the way of Thomas Kuhn in *The Structure of Scientific Revolutions*, and most especially in "Second Thoughts on Paradigms" and, from a slightly different perspective, Jean Piaget as a representative genetic epistemologist (see e.g., B. Kaplan "Genetic Psychology, Genetic Epistemology, and Theory of Knowledge," in T. Mischel, ed., *Cognitive Development and Epistemology* (New York: Academic Press, 1971, pp. 61–81). Needless to say, this essay is a beginning exercise in this latter approach, as is my *Psychology and the Conceptual Foundations of Science*.

[37] See especially his books *Personal Knowledge* (New York: Harper & Row, 1958) and *The Tacit Dimension* (New York: Doubleday, 1966).

edge of scientific problems as *problems* is an instance of this tacit knowledge — neither inductive nor deductive methodologies of science can tell us how to perceive a scientific problem or puzzle as *such*. But having raised this fundamental problem (and implied its converse: how do we recognize a solution *as such*) Polanyi lets us down, claiming that his *labeling* of the problem as an instance of 'tacit knowing' is its solution.[38] The question of guidance in scientific life comes down to this: can we provide a psychological account of concept formation that can explain the all-pervasive phenomenon of 'tacit knowledge' in science and its practice? Specifically, can we explain how problems are seen *as such*?

I think that the constructive theory of the higher mental processes sketched above, with its generative approach to concept formation, can be elaborated to the point where it at least begins to solve the problem of tacit knowledge. If our knowledge is a system of highly abstract, generative rules of action (in Hayek's sense), then the rules of determination (of concepts) which constitute our understanding are always 'tacit.' Our expectations of the future are determined by our constructive memory: hence our theory of memory is ipso facto our theory of 'expectation' which in turn is ipso facto what guides our concept formation, which in turn is what guides scientific life. Just as our (physiological) memory creates the richness of our sensory experience, so our constructive conceptual memory creates the richness of our inferential patterns. And the knowledge which is instantiated by these acts of construction is almost totally tacit in the sense that our conscious awareness is virtually never able to articulate the rules of determination which we are following in inferring our concepts. (But this tacit knowledge need not remain so forever: there is no reason why a completed psychology could not render explicit the rules of action according to which our knowledge is acquired and structured. Insofar as Polanyi claims that tacit knowledge is *inevitably* tacit, he is arguing from ignorance, and his position is to that extent obscurantist.)

The guidance of life is the psychology of 'seeing.' In normal science

[38] For this claim see Polanyi (1966) where he says: "We have here reached our main conclusions. Tacit knowing is shown to account (1) for a valid knowledge of a problem, (2) for the scientist's capacity to pursue it, guided by his sense of approaching its solution, and (3) for a valid anticipation of the yet indeterminate implications of the discovery arrived at in the end" (p. 24). But a description of a problem, even if correct, is not its explanation. What is required is a theory of tacit knowing and its interrelation to concept formation.

practice, when an enormous amount of background knowledge is stored away, and the practitioners of a science are free to explore nature to great depths, learning from exemplars (in Kuhn's sense) seems to constitute the group-licensed 'way of seeing.' By working with particulars, the practitioner gradually comes to see the abstract rule of determination *that makes the particular the kind of particular that it is*. The scientist's consequential knowledge of nature is tacitly generated by his exploration of exemplary particulars and the rules that determine their particularity. The growth of knowledge during periods of normal science consists in the search for new applications of known rules of determination and for new rules that bear a determinate (transformational) relation to the old ones. In forming new concepts of the structure of reality we take for granted a given conceptual-classificatory point of view and generate novel instances (applications) within that point of view. We 'see' the world as consisting of a specifiable group and the determinate transformations that the group may undergo: our 'inferences' consist in applying those transformations that we think adequately characterize the structure of reality, and what we 'test' is the match of expectation (generated by our model) and reality. As Kuhn emphasizes, we do not aim to uncover either fundamental novelties of fact or anomalies during such exploration.

But no set of concepts is adequate to the task of modeling reality exactly: we *do* turn up anomalies and novelties. Revolutionary reconceptualization can then provide a new way of seeing (i.e., a new group and set of transformations) that renders anomalies and novelties as expectancies, and the whole game of modeling reality may begin afresh. But during periods of revolutionary science, concept formation is *entirely* tacit and consists of finding new ways of seeing the entire domain in question. We *restructure* our representation of reality. During such reconceptualizations the research community involved very literally comes to see the domain from a different perspective, much as one can come to see the alternative perspective of an ambiguous figure that had not appeared to be ambiguous until the 'switch' occurred. Our perception of reality is inherently ambiguous in this same manner: since all our knowledge of reality is a function of the innate and learned modes of classification of our nervous systems, since we have no direct and unambiguous contact with external reality, our perception (and conceptualization) can shift whenever it is determined by a new rule of determination. All we can 'perceive' of external reality are disparate surface structure particulars that bear some (unknown) relationship to the actual deep structural nature of reality. That external reality determines (in the sense of causally influences) our conceptualization of it, but it cannot do so in an unambiguous manner: that is, we are free to conceptualize it according to one or another of the sets of rules of determination that our nervous systems are capable of constructing. This is not to say that reality per se is inherently ambiguous; rather, it is our perceptual and conceptual knowledge that is ambiguous in the sense that reality does not uniquely and unambiguously specify our classification of that reality. But the point remains: since reality does not unambiguously specify our conceptualization of it, the possibility of a scientific revolution that institutes a new 'group-licensed way of seeing' always faces *any* domain — no matter how thorough and certain our 'knowledge' of it may appear. And the only way we will be able to understand what is occurring in such reconceptualizations is to understand how we perceive and conceive, and how future knowledge acquisition is guided by present concept formation. As it is in normal science practice, the psychology of inference and expectation is indispensable to the understanding of revolutionary science.

There is a moral for the methodology of science in the psychological nature of guidance. The justificationist conception of 'inductive logic' as the guide to *rational* or 'good' scientific practice must be abandoned because it fuses guidance to assessment (by equating both to 'confirmation theory'). Confirmation theory is impotent with respect to scientific practice: even if a proposition could be proven true it would say nothing about where to look for the *next* 'truth'. But the Popperian methodology, which pictures science as conjectures and refutations *guided* by deductive logic (in the H-D explanatory framework), although a significant improvement upon the justificationist approach, likewise fails completely when it addresses the problems of guidance. Regardless of how much we may be able to learn from our mistakes (in the Popperian tradition), one thing we cannot learn is where to look next. And the pious claims of Lakatos (1970) that his methodology of scientific research programs reconstructed the rational growth of science (in his "improvement" of Popperian methodology) overlook the fact that Lakatos himself could not tell how a 'research program' would originate or what direction its course will take. All these 'rational reconstructions' of scientific growth are after the fact analyses that are totally incapable of providing even the weakest of methodological directives for actual research practice.[39] All the problems of guidance wind up in the 'wastebasket' of *external* history of science.

The practice of science is more a matter of matching to standards than of conjectures and refutations. What scientists *do* in gaining knowledge of reality is to search for patterns of regularity that will lead to rules of determination that may be imposed upon their initial sensory classifications. Our attempts to model reality by thought are attempts to determine invariances of structural properties that obtain between our theories and our sensory experiences. The brain creates a standard of expectation against which it attempts to match its sensory classifications. The process of matching to standards *is* the process of our knowledge acquisition. All our sophisticated hypothetico-inferential techniques are merely variations upon the fundamental theme of matching inferences to expectations. He who understands the psychology of inference and expectation understands the genesis of human knowledge.

Inference and the rationality of science. But what of the *rationality* of scientific practice? When *is* 'inference' rational? For the justificationist, if *scientific* inference cannot be *proven* to be a genuine source of proven or probable knowledge, then science-as-a-means-of-inference is no more rational than voodoo-as-a-means-of-inference (see Salmon, 1967, p. 55). According to justificationist intellectual honesty, one might then as well "do" voodoo as "do" science, since the one would be as *irrational* as the other. The tacit assumption is that inference cannot be rational unless it leads to knowledge (as the justificationist defines the term): since voodoo does not lead to (justificationist) knowledge, it is deemed irrational. Science, as our "good" means of inference, must be rescued from Hume's skepticism or it will be no better than voodoo and witchcraft.

[39] It is instructive to trace the evolution of practical methodological directives. The justificationist slogan was "aim at indubitable truths disclosed by indubitable method." The neojustificationist said, "Aim at probable truths . . ." with nothing else changed. Popper, after arguing that justificationist 'method' was chimerical, said, "Aim at bold, highly falsifiable theories and test them severely." Lakatos, attempting to bring Popperian methodology in line with actual scientific practice, exhorted us to "treat budding research programs leniently, and remember that criticism cannot kill" (he added 'quickly' after kill, but it is obvious that criticism can *never* kill a research program). Regardless of the hortatory appeal of these phrases, it is manifestly obvious that they provide no guidance to research at all: they are all directives for assessment, not acquisition. Beside them even Kuhn's modest proposal of "learn from exemplary puzzle solutions" is a revelation — for it is at least directed at the *acquisition* of knowledge, which, after all, *is* the concern of 'guidance'.

But if one abandons the justificationist framework the situation is quite different. All inference is equally 'inferential', equally conjectural and uncertain, and therefore there is nothing in the nature of the inference per se that separates science from voodoo. What renders science a rational source of knowledge (as opposed to voodoo) is not the use of a special type or pattern of inference: for both share inference in common. It is not that scientific inference is logical and justifiable, while voodoo inference is illogical and unjustifiable: both are unjustifiable and psychological rather than logical. If science is to be deemed rational and voodoo irrational, it cannot be on the grounds that their inferential procedures differ: the rationality of science must lie elsewhere.

The problems of demarcation (of science from its rivals) are as manifold as the problems of inference. Its alleged possession of a theory of instant rational assessment (the computational formula $P(h,e) = c$ of inductive logic is an 'instant' assessment procedure) is what the justificationist thought demarcated science from its rivals. But there is no instant assessment in science, as Popper and Kuhn have unceasingly pointed out. Hence science is irrational according to justificationist criteria of rationality. While denying any theory of instant assessment, Popper followed the justificationists in locating the rationality of science in its assessment (i.e., acceptance) of statements as warrantedly assertible. For Popper rationality lies in criticism of already inferred statements — and the 'critical approach' unfuses criticism from the attempt to justify (see Bartley, 1962). Within this nonjustificationist framework, inference is rational to the extent that it produces knowledge — to the extent that its end products are warrantedly assertible statements. Inference is thus unfused from rationality: rationality lies in the critical approach to the problems of assessment rather than in the 'logic' of induction. Voodoo can be as rational as science as far as its inferences are concerned: if it is to be deemed irrational, it must be on other grounds (for Popper, on the grounds that it is not critical of the warrant of its assertions).

The problem, of course, is that it is not at all obvious that Popperian criteria of demarcation (including Lakatos's (1970) progressive and degenerating problem shifts and research programs) are sufficient to capture the rationality of science. Kuhn (in the Postscript to *The Structure of Scientific Revolutions* and in "Reflections on My Critics," in Lakatos and Musgrave, eds., *Criticism and the Growth of Knowledge*) has argued convincingly (to my mind) that Popperian demarcation criteria are in-sufficient and that indeed there may be no satisfactory demarcation of science from its rivals *on methodological grounds*. But regardless of whether or not we can succeed in capturing the rationality of science, the thing to note is that inference must be unfused from rationality in order to see the true problems of either. Regardless of the adequacy of demarcation attempts, there is nothing in inference per se that renders science rational and its rivals irrational. One could think that science was rational because it possessed *the* true inference procedure only if the justificationist metatheory were presupposed. The nature of inference is one thing, the nature of rationality is another.[40]

Nevertheless, explaining how scientific knowledge grows (i.e., how scientific concepts are formed) ought to obviate most of the traditional problems of acceptance and its 'rationality'. Inference is rational when it produces knowledge. (Thus the Popperian phobia of commitment is not warranted — being committed to a way of 'seeing', as in embracing a normal science research tradition, is not inherently irrational. So long as commitment produces knowledge claims that are knowledge claims, it is perfectly rational; it is only when commitment ceases to produce knowledge that it is irrational.) The problem of rationality must be relocated into explaining the nature of factual truth and the nature of knowledge claims: when is an inferential claim warrantedly assertible, and what does warranted assertibility depend upon? What is the nature (and conditions) of factual truth? So far as rationality is related to inference, it has to do with the warrantability of inferences, i.e., with the *products* of inference, rather than with the processes per se. Discussing these issues is beyond our present scope.

The moral for the understanding of scientific *practice* that results from the study of knowledge acquisition comes down to this: no matter what philosophical reconstructions of the problems of assessment may

[40] This being so, defenses of the scientific 'game' *cannot connect methodology*, even synthetically through a contingent theory, *to epistemology*. This is because all sentient higher organisms, including all men (scientists and nonscientists alike), share the same inferential means of knowledge acquisition. Thus it is futile to connect verisimilitude to science (as did Lakatos in "Popper on Demarcation and Induction," in P. A. Schilpp, ed., *The Philosophy of Sir Karl Popper*, pp. 241–73) in a quasi-inductive justification of science over its alternatives. It is futile because all thought models reality, not just "scientific" thought (and because the empirical content of a theory cannot be equated with Popperian truth content). Such quasi-inductive "justifications" justify common sense, voodoo, etc., just as much as science, because they all share the same inferential processes. If science qua game is a better game to play than voodoo, it cannot be because of its method of inference.

yield, the rationality of scientific concept formulation is ultimately psychological (and to a certain extent sociological). The growth of scientific knowledge is rational in the sense that an ideally complete psychology of inference and expectation could explain the principles according to which human concept formation occurs. No purely philosophical methodology, which fails to base its account upon the psychological principles of knowledge acquisition, can be other than a reconstruction of the warrantability of knowledge claims *once they are put forward*. The philosopher's task is to worry about the assessment of theoretical formulations once they are at hand — he must surrender the acquisition of knowledge to the psychological sciences. This means, simply put, that there is no unitary problem of growth in science: there are at least two problems (assessment and acquisition), probably more. Purely philosophical approaches to a 'unitary' theory of growth, such as those of Popper and Lakatos, are bound to fail. And lest the philosopher feel complacent with retaining the problem of assessment in science, he need only be reminded that all extant methodologies are glaringly inadequate in one or another respect, and if he wants to preserve assessment as a philosophical problem he had better hurry, for the psychologist can say a few things about that problem, too. The philosopher had better come up with something more adequate than more volumes like this one on 'confirmation theory' or the scientists that he is supposed to be "educating" will cease listening entirely.