THE EPISTEMOLOGICAL ENGINE

Joseph C. Pitt

By now it is common to acknowledge that observation is theory-laden. Even observations made in the context of normal daily routine, i.e. non-scientific observations, are colored by background assumptions, proto-theories and the metaphysics of ordinary belief. As a result of this relation between the theoretical dimension of our knowledge and our observations it has been argued that theories and their associated observation-sets form independent units and, hence, rational comparison between theories is impossible. Further, it has been suggested if theories are not comparable then the growth of knowledge is a chimera, relativism is unavoidable and the priority of science as our epistemological engine is in jeopardy.

Fortunately the actual situation is not as portrayed. Theories *are* compared, many scientists *are* convinced that they are actually dealing with the "real" nature of matter, and the success of the scientific enterprise alone argues for its priority. Now it might be argued that what goes on in the world is irrelevant to such conceptual difficulties as those noted above. But, it should also be noted that the philosophy of science has a responsibility to be true to the actual practice of science in addition to providing the services of conceptual analysis and logical explication. That scientists do compare data and theories is a fact we have to account for. We can't have it both ways, arguing, on the one hand, that science is our best example of rational activity, and providing, on the other hand, analyses that fit no particular science. Our goal must be historically accurate and philosophically perspicuous results.

Where then is the slip between theorizing about science and the nature of the activity as practiced ? The problem lies not in flawed reasoning. Once theory and observation were accepted as the basic ingredients of scientific practice it almost seemed inevitable

that logic would lead us to incommensurability. Rather, our mistake lay in neglecting to include a third basic ingredient along with theory and observation : technology. As I will argue, the ramifications of this addition are enormous, for if introduced fully and not merely attached as an after-thought, we will be forced to rework our analysis of science. In this paper I attempt the beginning of such an effort. I concentrate on providing an alternative to the established view of the epistemology of science where technology is introduced as a factor with which to be reckonned. In the established view there are a number of assumptions about the nature of knowledge in general and, more specifically, its characterization in science which bear on my general theme; these are :

(1) the Aristotelian distinction between theoretical and applied knowledge, with science represented on the side of pure knowledge.

(2) the hierarchical account of knowledge with "pure" scientific knowledge presented as superior to applied knowledge.

(3) the characterization of technology as applied knowledge, hence inferior to science[1].

The result of operating in the context of an atmosphere created by implicitly accepting these assumptions is a highly rarefied and unresponsive philosophy of science. To challenge the established view I propose a pragmatic account of knowledge which overtly takes technology as primary in any account of the nature of knowledge-producing activities. I first examine a number of outmoded and cherished epistemological presumptions which have been primarily responsible for our failure to give technology its due. In part 2, a model of technology is proposed. In part 3 some of the consequences of accepting the model are considered by looking at what this view means for our analysis of rationality. As will be seen many traditional problems such as incommensurability fall by the wayside.

1. Pragmatic Epistemology

Let us begin by examining some highly visible conceptual albatrosses. The first of these is the venerable distinction between "pure" and "applied" with respect to science and technology. As noted above, rumor has it that science is pure and technology is applied. The problem here is to determine what is supposed to be pure or applied in either area. If the proposed answer is "knowledge", then the view that science is pure knowledge and technology is applied knowledge is surely false, for science is not pure

knowledge. Likewise, if technology is supposed to be applied knowledge, this view must be rejected, for technology does not necessarily require prior grounding in the theoretical work of science. Without trying to define "knowledge" we can, nevertheless, agree that the product of science is knowledge. To maintain this requires invoking (rather than attacking) a different distinction, this time between the process whereby we produce knowledge and the product of that process.

Science itself is a process composed of a large number of diverse activities undertaken by a variety of individuals. If science is a process and science is supposed to be pure knowledge, then it would appear that science as pure knowledge would be pure process. This view, however, can't be maintained. Everything we have learned about the conceptualized nature of knowledge argues against it. Science cannot generate undefiled pictures of the way the world really is; for this is what it would be like for pure knowledge to be pure process; i.e. unmediated access to the flow of events in undifferentiated space and time.

If pure knowledge as pure process can't be made intelligible, what about "pure knowledge" as "knowledge for its own sake"? That battle cry must be abandoned also. There is a basic incompatibility between the idea of knowledge and the idea of knowledge for its own sake. That incompatibility stems from the fact that the production of knowledge is a community enterprise. The process/product distinction used above can also be of assistance here. One of the more restrictive features of the epistemology championed by Locke, Berkeley and Hume was its emphasis on the means whereby an *individual* acquired the basic material from which he formulated his beliefs and the process by which he transformed that material into beliefs. Coupled with this insistence on the role of the individual was a confusion of the means individuals employ to develop their beliefs with the product of that process represented as knowledge.

There is no problem in allowing that individuals inquire, discover and offer candidates for inclusion in the body of accepted and integrated claims we recognize as knowledge. Let us refer to these as candidate-claims. Trouble begins when it is also assumed that the individual producing that candidate-claim is also solely responsible for determining whether or not that candidate-claim is to be admitted into the knowledge corpus. To make this assumption is to take the path of traditional empiricist epistemology and to

commit oneself to dealing with its attendant difficulties. The way out of those difficulties is to make the determinant of knowledge the community, not the individual. In this view the ultimate status of a knowledge candidate-claim is determined by the community, not the individual. This is the insight of C. S. Pierce and it characterizes the major break with traditional epistemology which pragmatism represents[2].

If the community is the arbiter of knowledge, action is its criterion. For the community the ultimate test of what is to count as knowledge is determined in terms of action. If the world is reported to be a certain way, the final test of that candidate-claim will be the success of a group acting as if the world were in fact that way. Not only is knowledge determined by the limits of action in this fashion, its purpose lies in action. We seek to discover or uncover the way the world really is to better make our way around in it[3].

Now it may be that certain individuals involved in the process called science do whatever they do simply because of a certain innocent curiosity or because they derive personal delight from these activities. But we should not confuse the collective activity called science with the personal delights of one person. Whatever candidate-claim the individual proposes, it is the conclusions of the community of investigators that makes the difference as to whether or not that candidate-claim is to count as knowledge. The community decides in accordance with its criterion of capable of being successfully acted upon[4].

If we turn now to the other side of the coin and consider the idea that technology is merely applied science and, hence, applied knowledge, we don't fare much better. To begin with there are technological items which have been used without a sufficient theoretical knowledge base to explain how they work, as examples consider Galileo's telescope and the roads of Rome. When Galileo built his telescope there was no scientific explanation based on a proven theory of what made the telescope do the things it did; likewise for gun powder, the fact that boats floated and so on. If we turn to Roman roads, there simply were no scientific principles which guided their construction or that of aquaducts or catapults.

Second, the history of science shows clearly how such quick accounts of the relation between science and technology fail. For the facts of the matter are vastly more complicated than the pure/applied distinction would suggest. If anything, the relation is symbiotic and mutually nurturing with theory and mechanism

feeding on and fueling each other. One would think that the analysis of this relationship would be exactly the sort of problem that would attract historians and philosophers of science and technology. And while historians have addressed some of the issues here, philosophers have been extremely slow to rise to the challenge.

2. Technology

Traditionally, philosophers of science have been concerned with th concept of a theory, the nature of explanation, and the problem of confirmation. Although current research in the philosophy of science has gone beyond this somewhat restrictive set of concerns, only recently has the range of issues explored been expanded to include questions such as the rationality of change in science. Until this expansion there was little apparent need for philosophers to worry about the details of the technological mechanisms used in various experiments or the social, political and economic considerations surrounding the process of on-going scientific activity. Rather, it was assumed that these matters had little to do with reconstructing the logic of the concepts with which they were concerned.

In taking notice of factors beyond those of the logic of concepts, philosophers of science were simultaneously provided the opportunity to integrate philosophical discussions of technology into the general dialogue. Kuhn forced us to recognize that factors other than those of pure reason or logic play a serious role in determining the options scientists select and the avenues of research they develop; science does not just roll on unimpeded towards the truth. An examination of those factors becomes increasingly important today when we try to understand the structure and development of the contemporary scientific establishment. If many features crucial to the past development of science were a function of technological considerations, such indeed is turning out to be doubly the case today. Not only are specific instrumentalities involved, as in the past where cognitive advances sometimes waited for the development of items such as the telescope, but today another type of "machine" is intimately involved in the doing of science in a way which critically determines what can and cannot be done. This is the bureaucracy of the institution of science. It includes not only federal funding agencies, but journals and the educational process. In the context of the doing of science, the bureaucracy is a tool to be used and manipulated to achieve specifiable results.

That the support system of science can be viewed as a tool should come as no surprise. Although we tend to think of tools in more simple-minded terms, there is nothing in principle that keeps us from recognizing more complicated items that are as much tools as hammers and wheels are. (See Mumford, 1966). To do so, however, requires that we run a bit roughshod over some territory that we will return to later with a more fine grained analysis.

To accommodate this social dimension of technology we need a revised account of technology. This requires redefining "technology" and if we are to engage in the process of redefining "technology" then it seems worthwhile to include as many dimensions of technology as possible. One way is to expand our account beyond the standard view of tool-as-mechanical-mechanism to tool as mechanism in general. Thus if a tool is conceived of as a means to facilitate accomplishing tasks, then it would appear that governments, organizations and hierarchies represent tools just as much as hammers and nails do.

For too long technology has been construed as mechanical tools. But not all tools man has used to secure himself against the elements, are mechanical devices. Some of the earliest and most important of man's technologies are the social structures he devised for establishing order and protection. But extending our understanding of tools to incorporate these social phenomena must be done carefully. There is a danger that if construed too broadly "technology" may become a concept without content. Nevertheless, to understand technology and how it relates to science and society in general, we have to see it in broader terms than have previously been used.

In his short brief for common sense approaches to understanding technology, Emmanuel Mesthene characterizes technology as "the organization of knowledge for the achievement of practical purposes" (1970, p. 25). And while this approach to technology permits the latitude we require, it still needs some critical attention. For instance, on just about any contemporary epistemology, the phrase "the organization of knowledge" would be redundant. However "knowledge" ends up being refined, it eventuates in a structured set of claims, for which organization is a necessary condition; gone are the days of unique knowledge claims standing alone from all influence. But, if that is the case, then Mesthene's account of technology comes out as knowledge for the sake of practical purposes. Not only is this a perfectly acceptable account

for a pragmatist, but it also provides the initial structure for unpacking a model of technology which both exhibits the insight characterized here, and permits us to see the nature of the similarities in the varieties of technology we wish to consider.

Common to tools and institutions and decision-making procedures alike is the simple process of transforming some input into an output. In making a decision, knowledge is as an input and achieving a practical purpose is an output. Following this line of thought leads to the view of technology as an input/output transformation process. The basis for this model can be found in Glendon Schubert's analysis of the structure of decision-making in the Supreme Court (Schubert, 1965). Consistent with Schubert's view on the nature of the decision-making/policy-making process is the consequence that there are input/output transformation processes whose function is to develop other input/output transformation processes. Let us refer to these as first-order and second-order transformations. Decisions are first-order transformations. The result of a first-order transformation may be either another first-order transformation or a second-order transformation.

A second-order transformation involves a constructed device. The construction of an oil refinery is a second-order transformation. It is the result of a first-order transformation in which a company made a decision using available knowledge etc. to build a refinery. Thus, decision-making procedures are first-order transformation processes or first-order transformers. In the case of the building of a refinery, we have a nice complicated example because the decision to build the refinery actually amounts to authorizing another series of first-order transformers which are the planning, designing and construction of the project. The construction of the refinery involves decisions as well as manipulation of material. The completed refinery is itself a second-order transformer since it deals with the manipulating of raw materials and their transformation by mechanized means. So we distinguish between mechanized processes and decision-making processes thereby allowing institutionalized decision-making processes such as bureaucracies (or science) to be characterized as technologies, albeit of a different kind.

But to talk of technology as first and second-order input/output transformations doesn't mean that we can't analyze it further. The model proposed here recognizes three components : first-order transformations, the second-order transformations and assessment feedback.

(a) a first-order transformation process is a set of deliberations wherein using an already established knowledge base, or starting from a given state of cognitive and social development, we confront a decision-forcing situation generated by what is perceived to be a problem (or a set of problems). Some of the solutions to these problems may require the creation of new machines. Other problems may require for solution the implementation or elaboration of a specific kind of social machinery such as a new organization or bureaucracy or legal system. Whatever the nature of the particular solution to these kinds of problems, the setting proposed for those deliberations raises a number of other issues. Among these are the nature of the deliberative process, the structure of practical reason and the nature of rationality, and the definition of "knowledge". I will return to some of these below.

(b) The second part of the model consists of secondary processes which exhaust what is normally called "technology" in ordinary usage, i.e. the machines. But it is also a broader category, allowing as it does, social structures. Thus, an oil refinery can be discussed in terms of the transformation of crude oil into petroleum or a legal system in terms of transforming conflict into resolution or science in terms of knowledge, theory and data into more knowledge. It is important to note here that whatever the end product of this kind of process, it is not an end in itself. These products have specific purposes and, therefore, further uses. At this point we might be tempted to invoke the intrinsic/extrinsic distinction to chracterize items which are ends in themselves as opposed to items whch are expressly concerned with furthering other ends. But like other distinctions under attack in this paper, this one also won't serve us well since it too will prove limiting. The limitations become apparent when we realize that social machinery offers us a different kind of product. It may be argued, for example, that the end product of a decision-making process such as the legal system is justice. But justice can be seen both as an intrinsic end and as furthering social productivity by providing a reliable means for mediating conflict. If we see science as a social process whose product is knowledge we see that given our earlier discussion we can't argue for an intrinsic end for science.

(c) The third and final component of technology is assessment. Technology assessment is a special kind of decision-making in which the effects of implementing decisions of the first kind [(a) above] are illuminated by means of a feedback mechanism that

makes possible the upgrading of knowledge base for further decision-making. The scientific process may be its best general example. In many respects one of the most important aspects of contemporary technology is the large extent to which it formally incorporates an assessment feedback mechanism into the decision-making surrounding the development and implementation of plans for new ventures. A full-scale discussion of this phenomenon (not to be undertaken here) would reveal the means by which changing goals and values affect the development and implementation of new and innovative techniques for transforming raw materials into usable results in both the physical and social domains.

At this point it should be apparent that the proposed model is intended only to schematize the complexity and pervasiveness of technology rather than as a definitive description of its structure. For, on any more detailed analysis it will become clear that wherever a decision is involved, so too are a variety of other considerations, among them the assessment process, the nature of the second order transformer, i.e., mechanical or social, the social circumstances, the goals of the individuals as well as the goals of the institutions, and so on down a virtually unclosed spiral. The success of this model, then, should not be judged in terms of whether it simplifies our view, it doesn't; nor was it intended to. The merits of the proposed account are to be found in the manner in which it exhibits the complexity of the actual situation, while still providing the means by which to isolate and analyze the relevant components and their interaction in the changing face of science.[6]

3. Technology and Rationality

While briefly describing the first component of the social process model of technology, four items were enumerated which need amplification. These are : the structure of the deliberative process, the structure of practical reason and the nature of rationality, and the definition of "knowledge". A great deal of work on the deliberative process has been done by a number of economists and general theorists of decision-making to which little can be added at this time. We have already dealt somewhat with the problem of knowledge and propose temporarily to leave that aside, preferring to return at the end for a final word. We turn now to the structure of practical reason. A proper view of practical reason places discussions of rationality in a very different light and should mitigate somewhat

against absolutist judgments on the merits of science and technology.

There is a standard account of rational man as *homo economicus* — economic man (genetically ?) programmed by the maximin principle. The main complaints with this account are two : (1) on an individual basis, that is for the purposes of explaining the behavior of a single individual, the *homo economicus* model is too successful — it explains everything. (2) On the other hand, at the level of group decision-making, it can't readily explain very much. Group interests are hard to identify without generating self-satisfying criteria. But the *homo economicus* model, nevertheless, comes fairly close to representing the way things actually are.

The difficulty with the individual case is that an account of why an individual acted as he or she did can always be constructed after the fact. To predict individual behavior on the basis of principles more precise than maximize utilities and minimize losses is more difficult, since predicting individual utilities ahead of time is hopeless; preference rankings and the reasons for change in such orderings are not necessarily accountable for on the basis of rational principles. *After* the fact, self-sacrifice by religious fanatics can be seen to be as rational as unloading stock at the end of a week of rising prices before the market closes on Friday afternoon.

But to a pragmatist these sorts of difficulties are bothersome only to the extent that they call for correction, not elimination. To shore up the predictive ability of the individual model requires some means by which individuals can be categorized according to membership in social groups, thereby facilitating the isolation of relevant governing *social* norms which function as overriding principles for individual choice.

The assumption here is that whatever goals and objectives an individual may have, those goals are obtainable only within a social context, thereby requiring that actions designed to achieve those goals accord with socially accepted norms. This is where we can register the point made earlier about knowledge as the product of group endorsement. To elaborate, it is essential to realize that social norms are not merely utilities to be weighted and figured into some formula. Rather, social norms function as setting the context in whch further deliberations occur. This is true even for the extreme case such as the rugged individualist. For example, once having identified oneself as fitting in the mold, or wanting to fit the mold of the rugged individualist, the actor has identified a group whose behavior can be characterized in terms of certain general principles

of action, however antisocial or non-social that may be. Under these circumstances if one is to continue one's membership in the group, only certain options for action are appropriate. So group membership narrows the set of possible options in any given decision. How one then proceeds to select some particular option is a different story.

This general account can be made more precise by distinguishing between the reasoning pattern which sets the context and the reasoning which generates or precipitates a particular action.[7] The context is set with the selection of an appropriate policy to guide one's actions. The general form of such a policy statement is :

When in circumstances C do X.

Having determined that one indeed is in circumstances C, then one is entitled to believe that he or she is warranted in deciding to do X. Deciding whether or not one is in the appropriate circumstances is not usually deemed to be a problem. But on the model being elaborated here, it is *the basic problem* — the one over which most difficulty arises. For once having determined what the circumstances are, the best action ought to be clear.

It is at the level of deciding what the situation is that we do best to invoke the typical decision-theoretic structure. It is rarely clear that we know what circumstances we are in. In fact, the risk involved in deciding under uncertainty is the risk of misidentifying the circumstances, not the risk of choosing the "wrong" or less optimal option. For, depending on the circumstances, the options will vary. Deciding what the circumstances are involves all the relevant factors invoked in standard decision-making such as the correctness of the information, the degree of risk given the misidentification of the circumstances and the action appropriate to it, the value of this being circumstance X as opposed to Y, etc.

The *homo economicus* model can then be seen as best suited not for the determination of the rationality of the choice of action, but rather for the determination of circumstances. The question of the rationality of the choice of the action is a question not to be directed to the individual so much as to the community of actors that decided that in circumstances C it is appropriate to do X. These decisions are rarely straightforward; if anything, the factors that bear on group decisions are seldom matters of pure reason or logic alone. Thus, if we draw an analogy between the scientific community

and a board of directors, given that the board of directors agrees that its job is to maximize profits, little follows as to how this is best accomplished. Indeed, in the minds of the members of the board, even what constitutes maximizing profits may vary. The bottom line of the Annual Report may not be the bottom line. As an example, consider the current debate in the U.S. over the possible adverse effects of the policy of rewarding top executives with bonuses based on yearly profits. It is suggested that this policy encourages short term profits at the expense of the long term viability of the company. How a board of directors is supposed to resolve this issue on rational grounds is unclear. Maybe here too we are stuck with an outdated notion. Anyone who has ever engaged in group decision-making knows that the stark ontological primacy of compromise is the furnace in which all general policies for action are cast. Why then should we assume that such results are rational ?

It still may be possible however, to salvage rationality as a label for a category of actions. Although not usually viewed this way, a strong case can be made for the claim that the most important advance in technology today, however you want to read it, comes not from the machines, or even the social changes they have helped secure, but rather from the methods we have developed to assess the effects of the machines and changes. In this sense technology monitors its own progress. It incorporates feed-back mechanisms whch have major effects on the determination of the nature of the initial circumstances in the next stage of action. These feed-backs force reassessments of basic assumptions. Failure to heed the warnings of the assessments, if anything counts, is tantamount to acting irrationally — in the sense of irresponsibility. So, to the extent that rationality has a role to play, it is best teamed with evaluation rather than logical method; that is, it is not the failure to calculate the utilities correctly that marks an individual as irrational; rather, he is irrational if he is irresponsible, if he consciously fails to reconsider the nature of the circumstances given new data and its effect on the *status quo*. The overall picture of technology that finally emerges essentially ties technology to people in decision-making, hence, action-oriented situations. In other words, technology is best understood as man at work. Science then becomes part of this process. Contrary to Feyerabend (1970), rationality is seen not in steadfast defense of theory, but in the constant reappraisal of background assumptions in the light of new results.

There is much to recommend this view. There also looms a

major objection. "Is it not the case", it might be argued, "that this social conception of technology is so broad that it makes every feature of human activity technological ? If so, doesn't the thesis suffer from the same defects as the *homo economicus* model it explains everything, thereby explaining nothing ?"

Despite its emphasis on the social, the thesis here is primarily epistemological. As such it captures two essential features of knowledge characteristic of not only specialized efforts such as science, but common sense as well, thus strengthening the view that science is an extension of common sense reasoning (Quine, 1969; Sellars, 1963; Pitt, 1981). The first of these concerns the changing nature of knowledge, the second is its teleological character.

Knowledge is a tool we use to get around in the world. We use other tools as well. To use those tools successfully — i.e. to achieve one's goals, requires knowing how to use them. This requires experience. In acquiring experience the knowledge of what the tool is and how it can be used changes. That is, as we learn how to use the tool we learn its limits and how it interacts with the world and other tools. This changes the character of our knowledge by extending it and shaping its parameters. This feature of human experience, that we learn from experience in this way is common to all purposeful human action. The fact of acting changes the body of knowledge and, *ipso facto*, how we act. The formal incorporation of a feedback mechanism in the model captures this essential feature.

In addition to acknowledging certain basic characteristics of knowledge, this feedback feature of the model also provides a means of understanding how the size and complexity of tools affect our actions. It allows that every purposeful action is essentially tied to knowledge. If that were not the case, then purpose would not be distinguishable from instinct.

Now, given these two features of knowledge, that it changes in the face of being used and that it is purposeful, it is clearly not fair to charge that the social model of technology does too much. In fact, the case *for* the model can be made even stronger. To do so requires attacking yet another treasured distinction, this time between two different types of knowledge, theoretical and practical. This will be our final point.

There are two different types of attempted justifications for the conclusion that there must be different types of knowledge. Each has flaws. The first account begins with an assumed difference between types of activity and argues backwards to the conclusion

that there must therefore be different types of knowledge. Here it is first noted that there are different human activities. It is also acknowledged that different methods are employed by those engaged in those activities. Thus, it is concluded, physics must be fundamentally different from art since what physicists do is different from what painters do. This conclusion, however, is hasty. For it is equally clear that, depending on how we characterize what is fundamental, what physicists do is not fundamentally different from what artists do. Both artist and physicist work with preconditions and presuppositions concerning what it is they are trying to achieve. Both test those presumptions and adjust them in the light of experience, learning as they go along. A distinction between method and content is not needed to see that in a very basic sense the process in which physicist and painter are engaged remains fundamentally the same; it is the process captured by our model.

But since the activities and the products of the activities can be differentiated, this differentiation has been used as a justification for claiming that what it takes to use different methods to achieve these different ends must also be differnt, hence there must be different kinds of knowledge. But this doesn't follow either. The same form for knowledge can be used to generate different products. It seems reasonable, and this is one insight the economist provides, that the same structure given different input can generate different output. That doesn't explain *too much*, it simply allows for the systematic explanation of difference. Where the economist went wrong was in assuming that there is always some of the *same* input present in every deliberation, i.e. that the same objectives rule decision-making for all men — to maximize profit. That *does* attempt to explain too much by oversimplifying. If we loosen that restriction, we are onto a powerful explanatory model.

Let us now turn to the second objectionable attempted justification for maintaining there are two different types of knowledge. Aristotle did us all a disservice with his praise of the contemplative life and his assumption that abstract knowledge of basic metaphysical principles is superior to more mundane sorts of knowledge. (Nicomachean Ethics, Book X, Chapter 6—7). One must also not forget Aristotle's belief that in seeking higher and more general principles we are in fact aiming for knowledge of the underlying structure of the universe and the accompanying view, equally strange, that the way to find that structure is to generalize to principles of ever increasing universality [Whether Aristotle can sing-

ly be held responsible for warping the nature of the search for know-
ledge is dubious; Plato must share some of the blame too.] In these
two assumptions we have the foundation of what is today called
scientific realism. Recently scientific realism has come under intense
scrutiny (Laudan, 1981). Not only has the assumption about the
knowable structure of ultimate reality been challenged, but the
methodological principles underlying the realist's interpretation of
the results of science have been attacked, Consequently   the
structure of knowledge enthroned by the realists has also fallen, even
though they may not realize it yet. With no standing case for the
viability of an epistemology with different kinds of knowledge, it
should be easy to make plausible an argument for one basic form for
knowledge.

But while that move is tempting, for present purposes it is un-
necessary. It is enough to argue for a single criterion for knowledge :
action. Knowledge is for the sake of action. And while we may want
to admit that different steps are needed to transform information
or hypotheses into action, that does not disallow action as our
criterion. A hypothesis that leads us to believe that certain results
would follow *if* it were correct, ought not to be accepted as correct
if those results fail to obtain. Failure in the realm of action is the
final challenge. Whether it results in rejection of the hypothesis or the
modification of auxiliary hypotheses, to the extent that failure
motivates revision we must concede its power. In the face of action
all explanation and knowledge claims are born equal. It then follows
that to the extent that there is a hierarchy wherein knowledge/
candidate-claims are ranked according to their degree of
theoreticality, that hierarchy is formed *after* the test of action. The
positivists understood this when they characterized the efforts of
their research programs in terms of logical *reconstructions.* If there
is a hierarchy of types of knowledge, then it is an indicator of
transformability into concrete results, not an *a priori* assessment
of intrinsic merit.

The deficiencies of past accounts of science lie in the
assumption that one begins with current theory and proceeds to
examine the logical difficulties it encounters. The model offered here
suggests we begin with current theory *and* the current reservoir of
tools, institutions, goals, values, etc. and see how they affect one
another at a given time. It allows (intuitively in accord with history)
that there is no one factor which dictates how things change, but
many factors in constantly changing configurations. And finally,

the model stresses the point that success is determined by the amount of control we can exert over nature, hence its insistence that the means to accomplish that control is primary. That is why we must begin first with an understanding of technology if we are to finally give science its appropriate analysis.

Virginia Polytechnic Institute and State University

NOTES

[1] This is such a popular claim that in some senses it is hard to document. It does form a background, however, to arguments about funding basic research, as well as providing the basis for distinguishing within science between different forms of activity by scientists, e.g. experimentalists versus theoretical physicists, field versus lab biologists.

[2] Pierce (1955, p. 28, 39) emphasizes the individual's beliefs as opposed to what is produced by the community of investigators. His convergence view of the development of knowledge takes knowledge *per se* out of the hands of the individual and sees it as the result of long term investigation.

[3] I take this to be a practical matter of some substance. If this were not the case I would doubt that we would find ourselves constantly having to defend basic research which has no immediate promise of paying off with usable products.

[4] In many respects the view expressed here echoes some of what Kuhn has said (1970, pp. 237–238). It also raises the spectre of a robust relativism. Robust relativism is to be distinguished from other more debilitating forms of relativism. And while this is not the place to argue this case, I would suggest that in the rush to avoid some of the unwanted consequences of relativism associated with sociology and ethics, some critics may have thrown the baby out with the bath water. Any account of the changing nature and content of scientific knowledge is going to have to admit some form of relativism. The

usual move is to attempt to blunt this with an appeal to realism and some form of convergence theory. For some of the same reasons as Laudan advances in his (1981), I find this tactic unsatisfactory. The model proposed below is an attempt to provide an epistemologically satisfactory way out of the worries associated with the more narrow-

[5] There is a considerable amount of work that has been done on the question of the so-called influence of technology on society and the attendant difficulties. In addition to the more ideologically motivated pieces, which are of little concern here, there is some serious work that has been done in technology assessment — not to be confused with assessment in general, more of which later. See for example : Rossini (1979, 1980); Hetman (1973); Brooks and Bowers (1970); Carpenter (1977); Coates (1976); Arnstein (1977); Arnstein and Christakis (1975); Porter, Rossini, Carpenter and Roper (1980); Churchman (1971);

[6] The lasting value of Kuhn's model of scientific revolutions should be seen in this light. That is, despite his intentions, the structure outlined in his (1970) should be seen as a point of entry into the complexity of the interrelations of the components involved in change in science, not as a statement of fact.

[7] This is the line of attack Wilfrid Sellars takes in his (1964).

REFERENCES

ARISTOTLE, (1941), *Collected Works*. Edited by R. McKeon. New York : Random House.

ARNSTEIN, S.R., (1977), "Technology Assessment : Opportunities and Obstaclies", *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC—7, pp. 571—582.

ARNSTEIN S. R., and CHRISTAKIS, A., 91975), *Perspectives on Technology Assessment*. Jerusalem : Science and Technology Publishers.

BROOKS, H. and BOWERS, R., (1970), "The Assessment of Technology", *Scientific American*, Vol. 222, pp. 13—21.

CARPENTER, S. R. (1977), "Philosophical Issues in Technology Assessment, " *Philosophy of Science*, Vol. 44, pp. 574—593.

CHURCHMAN, C. W., (1971), *The Design of Inquiring Systems*. New York : Basic Books.

COATES, J. F., (1976), "The Role of Formal Models in Technology Assessment", *Technological Forecasting and Social Change*, Vol. 9, pp. 139—190.

CONRAD, J. (ed.), (1980), *Society, Technology and Risk Assessment*. New York : Academic Press.

FEYERABEND, P. K., (1970), "Against Method", *Minnesota Studies in the Philosophy of Science*, Vol. IV, pp. 17—130.

HEMPEL, C., (1965), *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. Glencoe : Free Press.

HETMAN, F., (1973), *Society and the Assessment of Technology; Premises, Conceptions, Methodology, Experiments, Areas of Application*. Paris : Organization for Economic Cooperation and Development.

JARVIE, I. C., (1980), "Technology and the Structure of Knowledge", *Philosophy and Technology*, Carl Micham and Robert Mackey, editors. New York : Free Press.

KUHN, T., (1970), "Reflections of My Critics", *Criticism and the Growth of Knowledge*. Cambridge : Cambridge University Press.

KUHN, T. (1970a), *The Structure of Scientific Revolutions*, 2nd edition. Chicago : The University of Chicago Press.

LAUDAN, L., (1981), "A Confutation of Convergent Realism", *Philosophy of Science*, Vol. 48, pp. 19—49.

MESTHENE, E., (1970), *Technological Change*. New York : Mentor.

MUMFORD, L., (1966), *Technics and Human Development*. New York : Harcourt Brace Jovanovich.

PEIRCE, C. S., (1955), *Philosophical Writings of Peirce*. Edited by J. Buchler. New York : Dover.

PITT, J. C., (1981), *Pictures, Images and Conceptual Change*. Dordrecht : Reidel.

PORTER, A. L. and ROSSINI, F. A., (1980), "Technology Assessment/Environmental Impact Assessment; Toward Integrated Impact Assessment", *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC—10, pp. 417—424.

PORTER, A.L.; ROSSINI, F. A.; CARPENTER, S. R.; ROPER, A. J., (1980), *A Guidebook for Technology Assessment and Impact Analysis*. New York : North Holland.

QUINE, W. V. O., (1953), "Two Dogmas of Empiricism", *From a Logical Point of View*. Cambridge, Mass. : Harvard University Press, pp. 20—46.

QUINE, W. V. O., (1969), "Natural Kinds", *Ontological Relativity and Other Essays*. New York : Columbia University Press, pp.

114—138.

ROSSINI, F. A., (1979), "Technology Assessment : A New Type of Science ?", *Research in Philosophy and Technology*, Vol. 2, pp. 341—355.

SCHUBERT, G., (1965), *Judicial Policy-Making.* Glenview : Scott, Foremann and Company.

SELLARS, W., (1963), "Empiricism and the Philosophy of Mind", *Science, Perception and Reality.* London : Routledge and Kegan Paul.

SELLARS, W., (1964), "Induction as Vindication", *Philosophy of Science*, Vol. 31, pp. 197—231.

VAN FRAASSEN, B. C., (1980) *The Scientific Image.* Oxford : The Clarendon Press.