

PHILOSOPHICAL  
PROBLEMS OF  
*SCIENCE &*  
*TECHNOLOGY*

Edited by

ALEX C. MICHALOS

University of Guelph  
Guelph, Ontario, Canada

ALLEN & BACON, INC. ◆ BOSTON

© 1974

## READING 2

Towards a Philosophy  
of Technology*Mario Bunge*

The application of the scientific method and of scientific theories to the attainment of practical goals poses interesting philosophical problems, such as the nature of technological knowledge, the alleged validating power of action, the relation of technological rule to scientific law, and the effects of technological forecast on human behavior. These problems have been neglected by most philosophers, probably because the peculiarities of modern technology, and particularly the differences between it and pure science, are realized infrequently and cannot be realized as long as technologies are mistaken for crafts and regarded as theory-free. The present paper deals with those problems and is therefore an essay in the nearly non-existent philosophy of technology.

## SCIENCE: PURE AND APPLIED

The terms "technology" and "applied science" will be taken as synonymous, although neither is adequate: in fact, "technology" suggests the study of practical arts rather than a scientific discipline, and "applied science" suggests the application of scientific ideas rather than that of the scientific method. Since "technique" is ambiguous and "epitechique" unborn, we shall adopt the current lack of respect for etymology and go over to more serious matters.

The method and the theories of science can be applied either to increasing our knowledge of the external and the internal reality or to enhancing our welfare and power. If the goal is purely cognitive, pure science is obtained; if primarily practical, applied science. Thus, whereas cytology is a branch of pure science, cancer research

is one of applied research. The chief divisions of contemporary applied science are the physical technologies (e.g., mechanical engineering), the biological technologies (e.g., pharmacology), the social technologies (e.g., operations research), and the thought technologies (e.g., computer science). In many cases technology succeeds a craft: it solves some of the latter's problems by approaching them scientifically. In other cases, particularly those of the social and thought technologies, there is no antecedent prescientific skill because the problems themselves are new. But in every case a distinction must be made between pure research, applied research, and the applications of either to action.

The division of science into pure and applied is often challenged on the ground that all research is ultimately oriented toward satisfying needs of some sort or other. But the line must be drawn if we want to account for the differences in outlook and motivation between the investigator who searches for a new law of nature and the investigator who applies known laws to the design of a useful gadget: whereas the former wants to understand things better, the latter wishes to improve our mastery over them. At other times the difference is acknowledged, but it is claimed that applied science is the source of pure science rather than the other way around. Clearly, though, there must be some knowledge before it can be applied, unless it happens to be a skill or know-how rather than conceptual knowledge.

What is true is that action—industry, government, warfare, education, etc.—often *poses problems* that can be solved only by pure science. And if such problems are worked out in the free and lofty spirit of pure science, the solutions to them eventually may be applied to the attainment of practical goals. In short, practice is one of the sources of scientific problems, the other being sheer intellectual curiosity. But giving birth is not rearing. A whole cycle must be performed before anything comes out from practice: Practice → Scientific Problem → Scientific Research (statement and checking of hypotheses) → Rational Action. Even so, this is far from being the sole way in which scientific research and action mingle. Ever since theoretical mechanics began, in the eighteenth century, to shape industrial machinery, scientific ideas have been the main motor and technology their beneficiary. Since then, intellectual curiosity has been the source of most, and certainly of all important, scientific problems; technology has often followed in the wake of pure research, with a decreasing time lag between the two.

This is not to debase applied science but to recall how rich its conceptual background is. In applied science a theory is not only the summit of a research cycle and a guide to further research: it

is also the basis of a system of rules prescribing the course of optimal practical action. On the other hand, in the arts and crafts, theories are either absent or instruments of action alone. In past epochs a man was regarded as practical if, in acting, he paid little or no attention to theory or if he relied on worn-out theories and common knowledge. Nowadays a practical man is one who acts in obedience to decisions taken in the light of the best technological knowledge—not pure scientific knowledge, because this is mostly remote from or even irrelevant to practice. And such a technological knowledge, made up of theories, grounded rules, and data, is in turn an outcome of the application of the method of science to practical problems. Since technology is as theory laden as pure science, and since this either is overlooked or explicitly denied by most philosophers, we must take a closer look at technological theories and their application.

#### TECHNOLOGICAL THEORIES: SUBSTANTIVE AND OPERATIVE

A theory may have a bearing on action either because it provides knowledge regarding the objects of action, for example, machines, or because it is concerned with action itself, for example, with the decisions that precede and steer the manufacture or use of machines. A theory of flight is of the former kind, whereas a theory concerning the optimal decisions regarding the distribution of aircraft over a territory is of the latter kind. Both are *technological theories* but, whereas the theories of the first kind are *substantive*, those of the second kind are, in a sense, *operative*. Substantive technological theories are essentially applications, to nearly real situations, of scientific theories; thus, a theory of flight is essentially an application of fluid dynamics. Substantive technological theories are always preceded by scientific theories, whereas operative theories are born in applied research and may have little if anything to do with substantive theories, this being why mathematicians and logicians with no previous scientific training can make important contributions to them. A few examples will make the substantive-operative distinction clearer.

The relativistic theory of gravitation might be applied to the design of generators of antigravity fields (i.e., local fields counteracting the terrestrial gravitational field), which in turn might be used to facilitate the launching of spaceships. But, of course, relativity theory is not particularly concerned with either field gener-

ators or astronautics: it just provides some of the knowledge relevant to the design and manufacture of antigravity generators. Palaeontology is used by the applied geologist engaged in oil prospecting, and the latter's findings are a basis for making decisions concerning drillings; but neither palaeontology nor geology is particularly concerned with the oil industry. Psychology can be used by the industrial psychologist in the interests of production, but it is not basically concerned with production. All three are examples of the application of scientific (or semiscientific, as the case may be) theories to problems that arise in action.

On the other hand the theories of value, decision, games, and operations research deal directly with valuation, decision-making, planning, and doing; they even may be applied to scientific research regarded as a kind of action, with the optimistic hope of optimizing its output. (These theories could not tell how to replace talent but how best to exploit it). These are operative theories, and they make little if any use of the substantive knowledge provided by the physical, biological, or social sciences: ordinary knowledge, special but non-scientific knowledge (of, e.g., inventory practices), and formal science are usually sufficient for them. Just think of strategical kinematics applied to combat or of queuing models: they are not applications of pure scientific theories but theories on their own.

What these operative or non-substantive theories employ is not substantive scientific knowledge but the *method* of science. They may be regarded, in fact, as scientific theories concerning action, in short, as theories of action. These theories are technological in respect of aim, which is practical rather than cognitive, but apart from this they do not differ markedly from the theories of science. In fact, every good operative theory will have at least the following traits characteristic of scientific theories: (1) they do not refer directly to chunks of reality but to more or less idealized models of them (e.g., entirely rational and perfectly informed contenders or continuous demands and deliveries); (2) as a consequence they employ theoretical concepts (e.g., "probability"); (3) they can absorb empirical information and in turn can enrich experience by providing predictions or retrodictions; and (4) consequently they are empirically testable, though not as toughly as scientific theories.

Looked at from a practical angle, technological theories are richer than the theories of science in that, far from being limited to accounting for what may or does, did or will *happen* regardless of what the decision-maker does, they are concerned with finding out *what ought to be done* in order to bring about, prevent, or just change the pace of events or their course, in a preassigned way. In a conceptual sense, the theories of technology are definitely poorer

than those of pure science; they are invariably *less deep*, and this because the practical man, for whom they are intended, is chiefly interested in net effects that occur and are controllable on the human scale: he wants to know how things within *his* reach can be made to work *for him*, rather than how things of any kind really are. Thus, the electronics expert need not worry about the difficulties that plague the quantum electron theories; and the researcher in utility theory, who is concerned with comparing people's preferences, need not burrow into the origins of preference patterns—a problem for psychologists.

Consequently, whenever possible the applied researcher will attempt to schematize his system a *black box*; he will deal preferably with external variables (input and output), will regard all others as at best handy interverting variables with no ontological import, and will ignore the adjoining levels. This is why his oversimplifications and mistakes are not more often harmful—because his hypotheses are superficial. (Only the exportation of this externalist approach to science may be harmful.) Occasionally, though, the technologist will be forced to take up a deeper, representational viewpoint. Thus, the molecular engineer who designs new materials to order, that is, substances with prescribed macroproperties, will have to use certain fragments of atomic and molecular theory. But he will neglect all those microproperties that do not show up appreciably at the macroscopic level: after all, he uses atomic and molecular theories as tools—which has misled some philosophers into thinking that scientific theories are *nothing but* tools.

The conceptual impoverishment undergone by scientific theory when used as a means for practical ends can be frightful. Thus, an applied physicist engaged in designing an optical instrument will use almost only ray optics, that is, essentially what was known about light toward the middle of the seventeenth century. He will take wave optics into account for the explanation in outline, not in detail, of some effects, mostly undesirable, such as the appearance of colors near the edge of a lens; but he will seldom, if ever, apply any of the various wave theories of light to the computation of such effects. He can afford to ignore these theories in most of his professional practice because of two reasons. First, the chief traits of the optical facts relevant to the manufacture of optical instruments are adequately accounted for by ray optics; those few facts that are not so explainable require only the hypotheses (but not the whole theory) that light is made up of waves and that these waves can superpose. Second, it is extremely difficult to solve the wave equations of the deeper theories save in elementary cases, which are mostly of a purely academic interest (i.e., which serve essentially

the purpose of illustrating and testing the theory). Just think of the enterprise of solving a wave equation with time-dependent boundary conditions such as those representing the moving shutter of a camera. Wave optics is scientifically important because it is nearly true; but for most present-day technology it is less important than ray optics, and its detailed application to practical problems in optical industry would be quixotic. The same argument can be carried over to the rest of pure science in relation to technology. The moral is that, if scientific research had sheepishly applied itself to the immediate needs of production, we would have no pure science, hence no applied science either.

### DOES PRACTICE VALIDATE THEORY?

A theory, if true, can be employed successfully in applied research (technological investigation) and in practice itself—as long as the theory is relevant to either. (Fundamental theories are not so applicable because they deal with problems much too remote from practical problems. Just think of applying the quantum theory of scattering to car collisions.) But the converse is not true, that is, the practical success or failure of a scientific theory is no objective index of its truth value. In fact, a theory can be both successful and false; or, conversely, it can be a practical failure and nearly true. The efficiency of a false theory may be due to either of the following reasons. First, a theory may contain just a grain of truth, and this grain alone is employed in the theory's applications. In fact, a theory is a system of hypotheses, and it is enough for a few of them to be true or nearly so in order to be able to entail adequate consequences if the false ingredients are not used in the deduction or if they are practically innocuous. Thus, it is possible to manufacture excellent steel by combining magical exorcisms with the operations prescribed by the craft—as was done until the beginning of the nineteenth century. And it is possible to improve the condition of neurotics by means of shamanism, psychoanalysis, and other practices as long as effective means, such as suggestion, conditioning, tranquilizers, and above all time are combined with them.

A second reason for the possible practical success of a false theory may be that the accuracy requirements in applied science and in practice are far below those prevailing in pure research, so that a rough and simple theory supplying quick correct estimates of orders of magnitude very often will suffice in practice. Safety coefficients will mask the finer details predicted by an accurate and

deep theory anyway, and such coefficients are characteristic of technological theory because this must adapt itself to conditions that can vary within ample bounds. Think of the variable loads a bridge can be subjected to or of the varying individuals that may consume a drug. The engineer and the physician are interested in safe and wide intervals centered in typical values rather than in exact values. A greater accuracy would be pointless since it is not a question of testing. Moreover, such a greater accuracy could be confusing because it would complicate things to such an extent that the target—on which action is to be focused—would be lost in a mass of detail. Extreme accuracy, a goal of scientific research, not only is pointless or even encumbering in practice in most cases but can be an obstacle to research itself in its early stages. For the two reasons given above—use of only a part of the premises and low accuracy requirement—infinitely many possible rival theories can yield “practically the same results.” The technologist, and particularly the technician, are justified in preferring the simplest of them: after all, they are interested primarily in efficiency rather than in truth, in getting things done rather than in gaining a deep understanding of them. For the same reason, deep and accurate theories may be impractical; to use them would be like killing bugs with nuclear bombs. It would be as preposterous—though not nearly as dangerous—as advocating simplicity and efficiency in pure science.

A third reason why most fundamental scientific theories are of no practical avail is not related to the handiness and sturdiness required by practice but has a deep ontological root. The practical transactions of man occur mostly on his own level; and this level, like others, is rooted to the lower levels but enjoys a certain autonomy with respect to them, in the sense that not every change occurring in the lower levels has appreciable effects on the higher ones. This is what enables us to deal with most things on their own level, resorting at most to the immediately adjacent levels. In short, levels are to some extent stable: there is a certain amount of play between level and level, and this is a root of both chance (randomness due to independence) and freedom (self-motion in certain respects). One-level theories will suffice, therefore, for many practical purposes. It is only when a knowledge of the relations among the various levels is required in order to implement a “remote-control” treatment, that many-level theories must be tried. The most exciting achievements in this respect are those of psychochemistry, the goal of which is, precisely, the control of behavior by manipulating variables in the underlying biochemical level.

A fourth reason for the irrelevance of practice to the validation of theories—even to operative theories dealing with action—is

that, in real situations, the relevant variables are seldom adequately known and precisely controlled. Real situations are much too complex for this, and effective action is much too strongly urged to permit a detailed study—a study that would begin by isolating variables and tying some of them into a theoretical model. The desideratum being maximal efficiency, and not at all truth, a number of practical measures will usually be attempted at the same time: the strategist will counsel the simultaneous use of weapons of several kinds, the physician will prescribe a number of supposedly concurrent treatments, and the politician may combine promise and threats. If the outcome is satisfactory, how will the practitioner know which of the rules was efficient, hence which of the underlying hypotheses was true? If unsatisfactory, how will he be able to weed out the inefficient rules and the false underlying hypotheses?

A careful discrimination and control of the relevant variables and a critical evaluation of the hypotheses concerning the relations among such variables is not done while killing, curing, or persuading people, not even while making things, but in leisurely, planned, and critically alert scientific theorizing and experimentation. Only while theorizing or experimenting do we *discriminate* among variables and *weigh* their relative importance, do we *control* them either by manipulation or by measurement, and do we *check* our hypotheses and inferences. This is why factual theories, whether scientific or technological, substantive or operative, are empirically tested in the laboratory and not in the battlefield, the consulting office, or the market place. (“Laboratory” is understood here in a wide sense, to include any situation which, like the military maneuver, permits a reasonable control of the relevant variables.) This is, also, why the efficiency of the rules employed in the factory, the hospital, or the social institution can be determined only in artificially controlled circumstances.

In short, practice has no validating force; pure and applied research alone can estimate the truth value of theories and the efficiency of technological rules. What the technician and the practical man do, by contrast to the scientist, is not to *test* theories but to *use* them with noncognitive aims. (The practitioner does not even *test things*, such as tools or drugs, save in extreme cases; he just uses them, and their properties and their efficiency again must be determined in the laboratory by the applied scientist.) The doctrine that practice is the touchstone of theory rests on a confusion between practice and experiment and an associated confusion between rule and theory. The question “Does it work?”—pertinent as it is with regard to things and rules—is impertinent in respect of theories.

Yet it might be argued that a man who knows how to do some-

thing is thereby showing that he knows that something. Let us consider the three possible versions of this idea. The first can be summed up in the schema "If  $x$  knows how to do (or make)  $y$ , then  $x$  knows  $y$ ." Yet we are unable to make them, and we know part of one million years, man has known how to make children without having the remotest idea about the reproduction process. The second thesis is the converse conditional, namely, "If  $x$  knows  $y$ , then  $x$  knows how to do (or make)  $y$ ." Counterexamples: we know something about stars, yet we are unable to make them, and we know part of the past, but we cannot even spoil it. The two conditionals being false, the biconditional " $x$  knows  $y$  if and only if  $x$  knows how to do (or make)  $y$ " is false, too. In short, it is false that knowledge is identical with knowing how to do, or know-how. What is true is rather this: knowledge considerably *improves* the chances of correct doing, and doing *may* lead to knowing more (now that we have learned that knowledge pays), not because action is knowledge, but because, in inquisitive minds, action may trigger questioning.

It is only by distinguishing scientific knowledge from instrumental knowledge, or know-how, that we can hope to account for the coexistence of practical knowledge with theoretical ignorance and the coexistence of theoretical knowledge with practical ignorance. Were it not for this the following combinations hardly would have occurred in history: (1) science without the corresponding technology (e.g., Greek physics); (2) arts and crafts without an underlying science (e.g., Roman engineering and contemporary intelligence testing). The distinction must be kept, also, in order to explain the cross-fertilizations of science, technology, and the arts and crafts, as well as to explain the gradual character of the cognitive process. If, in order to exhaust the knowledge of a thing, it were sufficient to produce or reproduce it, then certain technological achievements would put an end to the respective chapters of applied research: the production of synthetic rubber, plastic materials, and synthetic fibres would exhaust polymer chemistry; the experimental induction of cancer should have stopped cancer research; and the experimental production of neuroses and psychoses should have brought psychiatry to a halt. As a matter of fact, we continue doing many things without understanding how, and we know many processes (such as the fusion of helium out of hydrogen) which we are not yet able to control for useful purposes (partly because we are too eager to attain the goal without a further development of the means). At the same time it is true that the barriers between scientific and practical knowledge, pure and applied research, are melting. But this does not eliminate their differences, and the

process is but the outcome of an increasingly scientific approach to practical problems, that is, of a diffusion of the scientific method.

The identification of knowledge and practice stems not only from a failure to analyze either but also from a legitimate wish to avoid the two extremes of speculative theory and blind action. But the testability of theories and the possibility of improving the rationality of action are not best defended by blurring the differences between theorizing and doing, or by asserting that action is the test of theory, because both of these are false and no program is defensible if it rests on plain falsity. The interaction between theory and practice and the integration of the arts and crafts with technology and science are not achieved by proclaiming their unity but by multiplying their contacts and by helping the process whereby the crafts are given a technological basis and technology is entirely converted into applied science. This involves the conversion of the rules of thumb peculiar to the crafts into grounded rules, that is, rules based on laws. Let us approach this problem next.

### SCIENTIFIC LAW AND TECHNOLOGICAL RULE

Just as pure science focuses on objective patterns or laws, action-oriented research aims at establishing stable norms of successful human behavior, that is, rules. The study of rules—the grounded rules of applied science—is therefore central to the philosophy of technology.

A rule *prescribes* a course of action; it indicates how one should proceed in order to achieve a predetermined goal. More explicitly, a rule is an instruction to perform a finite number of acts in a given order and with a given aim. The skeleton of a rule can be symbolized as a string of signs, such as 1-2-3- . . . - $n$ , where every number stands for a corresponding act; the last act,  $n$ , is the only thing that separates the operator who has executed every operation, save  $n$ , from the goal. In contrast to law formulas, which say what the shape of possible events is, rules are norms. The field of law is assumed to be the whole of reality, including rule-makers; the field of rule is but mankind; men, not stars, can obey rules and violate them, invent and perfect them. Law statements are descriptive and interpretive, whereas rules are normative. Consequently, while law statements can be more or less true, rules can be only more or less effective.

We may distinguish the following genera of rules: (1) *rules of conduct* (social, moral, and legal rules); (2) *rules of prescientific*

*work* (rules of thumb in the arts and crafts and in production); (3) *rules of sign* (syntactical and semantical rules); (4) *rules of science and technology* (grounded rules of research and action). Rules of conduct make social life possible (and hard). The rules of prescientific work dominate the region of practical knowledge which is not yet under technological control. The rules of sign direct us how to handle symbols—how to generate, transform, and interpret signs. And the rules of science and technology are those norms that summarize the special techniques of research in pure and applied science (e.g., random sampling techniques) and the special techniques of advanced modern production (e.g., the technique of melting with infrared rays).

Many rules of conduct, work, and sign are *conventional*, in the sense that they are adopted with no definite reasons and might be exchanged for alternative rules with little or no concomitant change in the desired result. They are not altogether arbitrary, since their formation and adoption should be explainable in terms of psychological and sociological laws, but they are not necessary either; the differences among cultures are largely differences among systems of rules of that kind. We are not interested in such groundless or conventional rules but rather in founded rules, that is, in norms satisfying the following *definition*: A rule is *grounded* if and only if it is based on a set of law formulas capable of accounting for its effectiveness. The rule that commands taking off the hat when greeting a lady is groundless in the sense that it is based on no scientific law but is conventionally adopted. On the other hand, the rule that commands greasing cars periodically is based on the law that lubricators decrease the wearing out of parts by friction; this is neither a convention nor a rule of thumb like those of cooking and politicking—it is a well-grounded rule.

To decide that a rule is effective it is necessary, though insufficient, to show that it has been successful in a high percentage of cases. But these cases might be just coincidences, such as those that may have consecrated the magic rituals that accompanied the huntings of primitive man. Before adopting an empirically effective rule we ought to know *why* it is effective; we ought to take it apart and reach an understanding of its *modus operandi*. This requirement of rule foundation marks the transition between the prescientific arts and crafts and contemporary technology. Now, the sole valid foundation of a rule is a system of law formulas, because these alone can be expected to correctly explain facts,—for example, the fact that a given rule works. This is not to say that the effectiveness of a rule depends on whether it is founded or groundless but only that, in order to be able to *judge* whether a rule has any chance

of being effective, as well as in order to *improve* the rule and eventually *replace* it by a more effective one, we must disclose the underlying law statements, if any. We may take a step ahead and claim that the blind application of rules of thumb has never paid and, second, to try to transform some law formulas into effective technological rules. The birth and development of modern technology is the result of these two movements.

But it is easier to preach the foundation of rules than to say exactly what the foundation of rules consists in. Let us try to make an inroad into this unexplored territory—the core of the philosophy of technology. As usual when approaching a new subject, it will be convenient to begin by analyzing a typical case. Take the law statement “Magnetism disappears above the Curie temperature (770° C for iron).” For purposes of analysis it will be convenient to restate our law as an explicit conditional: “If the temperature of a magnetized body exceeds its Curie point, then it becomes demagnetized.” (This is, of course, an oversimplification, as is every other ordinary-language rendering of a scientific law: the Curie point is not the temperature at which all magnetism disappears but, rather, the point of conversion of ferromagnetism into paramagnetism or conversely. But this is a refinement irrelevant to most technological purposes.) Our nomological statement provides the basis for the nomoprismatic statement “If a magnetized body is heated above its Curie point, then it is demagnetized.” (The pragmatic predicate is, of course, “is heated.”) This nomoprismatic statement is, in turn, the ground for two different rules, namely, R1: “In order to prevent demagnetizing a body, do not heat it above its Curie point.” Both rules have the same foundation, that is, the same underlying nomoprismatic statement, which in turn is supported by a law statement assumed to represent an objective pattern. Moreover, the two rules are equiefficient, though not under the same circumstances (changed goals, changed means).

Notice, first, that unlike a law statement a rule is neither true nor false; as a compensation it can be effective or ineffective. Second, a law is consistent with more than one rule. Third, the truth of a law statement does not insure the efficiency of the associated rules; in fact, the former refers to idealized situations which are not met with in practice. Fourth, whereas given a law we may try out the corresponding rules, given a rule we are unable to trace the laws presupposed by it; in fact, a rule of the form “In order to attain the goal *G*, employ the means *M*” is consistent with the laws “If *M*, then *G*”, “*M* or *G*”, and infinitely many others.

The above has important consequences for the methodology

of rules and the interrelations between pure and applied science. We see there is no single road from practice to knowledge, from success to truth; success warrants no inference from rule to law but poses the problem of explaining the apparent efficiency of the rule. In other words, the roads from success to truth are infinitely many and consequently theoretically useless or nearly so, that is, no bunch of effective rules suggests a true theory. On the other hand, the roads from truth to success are limited in number, hence feasible. This is one of the reasons why practical success, whether of a medical treatment or of a government measure, is not a truth criterion for the underlying hypotheses. This is also why technology—in contrast to the prescientific arts and crafts—does not start with rules and ends up with theories but proceeds the other way around. This is, in brief, why technology is applied science whereas science is not purified technology.

Scientists and technologists work out rules on the basis of theories containing law statements and auxiliary assumptions, and technicians apply such rules jointly with groundless (prescientific) rules. In either case, specific hypotheses accompany the application of rules, namely, hypotheses to the effect that the case under consideration is one where the rule is in point because such and such variables—related by the rule—are in fact present. In science such hypotheses can be tested; this is true of both pure and applied research. But in the practice of technology there may not be time to test them in any way other than by applying the rules around which such hypotheses cluster—and this is a poor test indeed, because the failure may be blamed either on the hypothesis or on the rule or on the uncertain conditions of application.

### SCIENTIFIC PREDICTION AND TECHNOLOGICAL FORECAST

For technology knowledge is chiefly a means to be applied to the achievement of certain practical ends. The goal of technology is successful action rather than pure knowledge, and accordingly the whole attitude of the technologist while applying his technological knowledge is active in the sense that, far from being an inquisitive onlooker or a diligent burrower, he is an active participant in events. This difference of attitude between the technologist in action and the researcher—whether pure or applied—introduces certain differences between technological forecast and scientific prediction.

In the first place, whereas scientific prediction says what will

or may happen if certain circumstances obtain, technological forecast suggests how to influence circumstances so that certain events may be brought about, or prevented, that would not normally happen; it is one thing to predict the trajectory of a comet, quite another to plan and foresee the orbit of an artificial satellite. The latter presupposes a choice among possible goals, and such a choice presupposes a certain forecasting of possibilities and their evaluation in the light of a set of desiderata. In fact, the technologist will make his forecast on his (or his employer's) estimate of what the future *should* be like if certain desiderata are to be fulfilled; contrary to the pure scientist, the technologist is hardly interested in what would happen anyway; and what for the scientist is just the final state of a process becomes for the technologist a valuable (or disvaluable) end to be achieved (or to be avoided). A typical scientific prediction has the form "If  $x$  occurs at time  $t$ , then  $y$  will occur at time  $t'$  with probability  $p$ ." By contrast, a typical technological forecast is of the form "If  $y$  is to be achieved at time  $t'$  with probability  $p$ , then  $x$  should be done at time  $t$ ". Given the goal, the technologist indicates the adequate means, and his forecast states a means-end relationship rather than a relation between an initial state and a final state. Furthermore, such means are implemented by a specified set of actions, among them the technologist's own actions.

This leads us to a second peculiarity of technological forecast: whereas the scientist's success depends on his ability to separate his object from himself (particularly so when his object happens to be a psychological subject)—that is, on his capacity of detachment—the technologist's ability consists in placing himself, within the system concerned, at the head of it. This does not involve *subjectivity*, since after all the technologist draws on the objective knowledge provided by science; but it does involve partiality, a *parti pris* unknown to the pure researcher. The engineer is part of a man-machine complex, the industrial psychologist is part of an organization, and both are bound to devise and implement the optimal means for achieving desiderata which are not usually chosen by themselves; they are decision-makers, not policy-makers.

The forecast of an event or process that is not under control will not alter the event or process itself. Thus, for example, no matter how accurately an astronomer predicts the collision of two stars, the event will occur in due course. But if an applied geologist can forecast a landslide, then some of its consequences can be prevented. Moreover, by designing and supervising the appropriate defense works the engineer may prevent the landslide itself; he may devise the sequence of actions that will refute the original forecast.



Similarly, an industrial concern may forecast sales for the near future on the (shaky) assumption that a given state of the economy, say prosperity, will continue during that lapse. But if this assumption is falsified by a recession, and the enterprise had accumulated a large stock which it must get rid of, then instead of making a new sales forecast (as a pure scientist would be inclined to do), the management will try to *force* the original forecast to come true by increasing advertisement, lowering sale prices, and so on. As in the case of vital processes, a diversity of means will alternatively or jointly be tried to attain a fixed goal. In order to achieve this goal any number of initial hypotheses may have to be sacrificed: in the case of the landslide, the assumption that no external forces would interfere with the process and, in the case of the sales, that prosperity would continue. Consequently, whether the initial forecast is *forcefully falsified* (as in the case of the landslide) or *forcefully confirmed* (as in the case of the sales forecast), this fact cannot count as a *test* of the truth of the hypotheses involved; it will count only as an efficiency test of the rules that have been applied. The pure scientist, on the other hand, need not worry about altering the means for achieving a preset goal, because pure science *has* no goals external to it.

Technological forecast, in sum, cannot be used for controlling things or men by changing the course of events perhaps to the point of stopping them altogether, or for forcing the predicted course even if unpredictable events should interfere with it. This is true of the forecasts made in engineering, medicine, economics, applied sociology, political science, and other technologies: the sole formulation of a forecast (prognosis, lax prediction, or prediction proper), if made known to the decision-makers, can be seized upon by them to steer the course of events, thus bringing about results different from those originally forecasted. This change, triggered by the issuance of the forecast, may contribute either to the latter's confirmation (self-fulfilling forecast) or to its refutation (self-defeating forecast). This trait of technological forecast stems from no logical property of it; it is a pattern of social action involving the knowledge of forecasts and consequently is conspicuous in modern society. Therefore, rather than analyzing the logic of causally effective forecast, we should start by distinguishing three levels in it: (1) the conceptual level, on which the prediction  $p$  stands; (2) the psychological level—the knowledge of  $p$  and the reactions triggered by this knowledge; and (3) the social level—the actions actually performed on the basis of the knowledge of  $p$  and in the service of extra-scientific goals. This third level is peculiar to technological forecast.

This feature of technological forecast sets civilized man apart

from every other system. A nonpredicting system, be it a jukebox or a frog, when fed with information it can digest will process it and convert it into action at some later time. But such a system does not purposely produce most of the information, and it does not issue projections capable of altering its own future behavior. A predictor—a rational man, a team of technologists, or a sufficiently evolved automation—can behave in an entirely different way. When fed with relevant information  $I^t$  at time  $t$ , it can process this information with the help of the knowledge (or the instructions) available to it, eventually issuing a prediction  $P^t$  at a later  $t'$ . This prediction is fed back into the system and compared with the present goal that controls the whole process (without either causing it or supplying it with energy). If the two are reasonably close, the system takes a decision that eventually leads it to act so as to take advantage of the course of events. If, on the other hand, the prediction differs significantly from the goal, this difference will again trigger the theoretical mechanism, which will elaborate a new strategy: a new prediction,  $P^{t''}$ , will eventually be issued at time  $t''$ , a forecast including a reference to the system's own participation in the events. The new prediction is fed back into the system and, if it still disagrees with the goal, a new correction cycle is triggered, and so on until the difference between the prediction and the goal becomes negligible, in which case the system's predicting mechanism comes to rest. Henceforth the system will gather new information regarding the present state of affairs and will act so as to conform to the strategy it has elaborated. This strategy may have required not only new information regarding the external world (including the attitudes and capabilities of the people concerned) but also new hypotheses or even theories which had not been present in the instruction chart originally received by the predictor. If the latter fails to realize it or to obtain and utilize such additional knowledge, his or its action is bound to be ineffective. Moral: the more brains the better.

### TECHNOLOGICAL FORECAST AND EXPERT PROGNOSIS

The preceding account of technological forecast is based on the assumption that it relies on some theory, or rather theories, whether substantive or operative. This assumption may be found wanting by anyone knowing that the forecasts issued by experts in medicine, finance, or politics are often successful and yet involve no great deal of theorizing. True, most often *expert prognosis* relies on inductive

(empirical) generalizations of the form "*A* and *B* occur jointly with the observed frequency *f*," or even just "*A* and *B* occur jointly in most cases," or "Usually, whenever *A* then *B*." The observation that a given individual, say a human subject or an economic state of affairs, has the property *A* is then used to prognose that it has, or will acquire, the property *B*. In daily life such prognoses are all we do, and the same applies to most expert prognoses. Occasionally such prognoses made with either ordinary knowledge or specialized but nonscientific knowledge are more successful than predictions made with full-fledged but false or rough theories; in many fields, however, the frequency of hits is not better than the one obtained by flipping a coin. The point, though, is that expert forecast using no scientific theory is not a scientific activity—just by definition of "scientific prediction."

Yet it would be wrong to think that experts make no use of *specialized knowledge* whenever they do not employ scientific theories; they always judge on the basis of some such knowledge. Only, expert knowledge is not always explicit and articulate and, for this reason, it is not readily controllable: it does not learn readily from failures, and it is hard to test. For the progress of science, the failure of a scientific prediction is by far preferable to the success of an expert prognosis, because the scientific failure can be fed back into the theory responsible for it, thereby giving us a chance to improve it, whereas in the case of expert knowledge there is no theory to feed failure into. It is only for immediate practical purposes that expert prognoses made with shallow but well-confirmed generalizations are preferable to risky scientific predictions.

Another difference between expert prognosis and technological forecast proper would seem to be this: the former relies more heavily on *intuition* than does scientific prediction. Yet the difference is one of degree rather than of kind. Diagnosis and forecast, whether in pure science, in applied science, or in arts and crafts, involve intuitions of a number of kinds: the quick identification of a thing, event, or sign; the clear but not necessarily deep grasp of the meaning and/or the mutual relations of a set of signs (text, table, diagram, etc.); the ability to interpret symbols; the ability to form space models; skill in realizing analogies; creative imagination; catalytic inference, that is, quick passage from some premises to other formulas by skipping intermediate steps; power of synthesis or synoptic grasp; common sense (or rather controlled craziness), and sound judgment. These abilities intertwine with specialized knowledge, whether scientific or not, and are reinforced with practice. Without them theories could neither be invented nor applied—but, of course, they are not suprarational powers. Intuition is all right as long

as it controlled by reason and experiment; only the replacement of theorizing and experimenting by intuition must be feared.

A related danger is that of *pseudo-scientific projection tools*, so common in applied psychology and sociology. A number of techniques have been devised to forecast the performance of personnel, students, and even psychologists themselves. A few tests, the objective ones, are somewhat reliable; this holds for intelligence and skill tests. But most tests, particularly the subjective ones, (the "global evaluation" of personality by means of interviews, the thematic apperception test, the Rorschach, etc.) are in the best of cases inefficient and in the worst of cases misleading. When they have been subjected to the test of prediction—that is, when their results have been checked with the actual performance of the subjects—they have failed. The failure of most individual psychological tests, and particularly of the subjective ones, is not a failure of psychological testing in general; what is responsible for such failures is either the total absence or the falsity of the underlying psychological theories. Testing for human abilities without first establishing *laws* relating objective indices of abilities or personality traits is as thoughtless as asking a tribesman to test an aircraft. As long as no theoretical foundations of psychological tests are secured, their employment as predictive instruments is not better than crystal gazing or coinflipping: they are practically inefficient and, even if they succeeded, they would not contribute to psychological theory, because they are unrelated to theory. The limited success of psychological testing has led many to despair of the possibility of finding a scientific approach to human behavior, but the right inference is that such an attempt has been tried only after a large number of alleged tests invaded the market. What is wrong with most of "applied" (educational, industrial, etc.) psychology is that it does *not* consist in the application of scientific psychology at all. The moral is that practical wants—such as personnel training and selection—should not be allowed to force the construction of "technologies" without an underlying science.

Technological forecast should be maximally reliable. This condition excludes from technological practice—not, however, from technological research—insufficiently tested theory. In other words, technology will ultimately prefer the old theory that has rendered distinguished service in a limited domain and with a known inaccuracy to the bold new theory that promises unheard-of forecasts but is probably more complex and therefore partly less well tested. It would be irresponsible for an expert to apply a new idea in practice without having tested it under controlled conditions. (Yet this is still done in pharmacy: recall the affair of the mutagenic drugs

in the early 1960's.) Practice, and even technology, are bound to be more conservative than science. Consequently, the effects of a close association of pure research with applied research, and of the latter with production, are not all of them beneficial: while it is true that technology challenges science with new problems and supplies it with new equipment for data-gathering and data-processing, it is no less true that technology, by its very insistence on reliability, standardization (routinization), and speed at the expense of depth, range, accuracy, and serendipity, can slow down the advancement of science.

#### OTHER PROBLEMS

We have looked into a few problems of the philosophy of technology. Many other challenging problems have been left out, for example, the logic of technological rules; the test of technological theories; the patterns of technological invention; the reason textile, aircraft, and other industries are still largely based on crafts; and the power of technology to bring together previously separate fields (cases of cybernetics, nuclear engineering, computer science, space science, and bioengineering). These and many other problems are waiting to be discovered and worked out by philosophers attentive to their own times. Why should the waiting time be so long?

#### SUGGESTIONS FOR FURTHER READING

- Campbell, N. *What Is Science?* New York: Dover Publications, Inc., 1952, pp. 1-36.
- Feigl, H., and Brodbeck, M. (ed.) *Readings in the Philosophy of Science*. New York: Appleton-Century-Crofts, Inc., 1953, pp. 3-20.
- Kemeny, J. G. *A Philosopher Looks at Science*. Princeton, N.J.: D. Van Nostrand Co., Inc., 1959, pp. 174-186.
- Krimmerman, L. I. (ed.) *The Nature and Scope of Social Science*. New York: Appleton-Century-Crofts, Inc., 1969, pp. 7-40, 139-204.
- Kuhn, T. S. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press, 1962.
- Martin, M. *Concepts of Science Education*. Glenview, Ill.: Scott, Foresman and Co., 1972, pp. 6-43.
- Nagel, E. *The Structure of Science*. New York: Harcourt, Brace and World, Inc., 1961, pp. 1-15.
- Popper, K. R. *Conjectures and Refutations*, 4th ed., London: Routledge and Kegan Paul, 1972, pp. 253-293.
- Ravetz, J. R. *Scientific Knowledge and Its Social Problems*. Oxford: Clarendon Press, 1971, pp. 7-241.
- Storer, N. W. *The Social System of Science*. New York: Holt, Rinehart & Winston, Inc., 1966.