TOWARDS A TECHNOETHICS *

Gone are the days of the divine right of kings—or of anyone else, whether owner, manager, labor leader, politician, bureaucrat, technologist or scholar. Absolute and groundless authority is being contested all over the world: ours is an iconoclastic age. Moreover nobody recognizes rights without duties, privileges without responsibilities. Everyone is rightly held responsible for what he does and even for what he fails to do when he ought to act. And the responsibilities are not to some conveniently distant deity or sovereign, or even to the anonymous people, but to definite individuals—superiors, peers, subordinates, neighbors, the public, and even possible future humans.

Moreover the old separations among different kinds of responsibility—moral, professional, social, etc.—are being lifted. We are beginning to realize that the separation of responsibilities is just a retreat from total responsibility, hence a cloak for wrongdoing. A person in charge of something, be it a machine or another individual, is not composed of a number of moral entities but is a single person, now acting in one capacity, now in another. All these various capacities should combine harmoniously. Being an affectionate parent does not exculpate any crimes; being a competent engineer does not confer rights of piracy on the environment; being an efficient manager does not entitle him to oppress others. Every human being has a number of intertwined responsibilities and each of them is as personal and intransferable as a joy or a grief.

This paper examines some of the special responsibilities of the technologist in our age of pervasive—and, alas, all too often perverse—technology. We shall defend the thesis that the technologist, just like anyone else, is personally responsible for whatever he does, and that he is responsible to all mankind not just to his employers. We shall also claim that the technologist has the duty to face, ponder over, and solve his own moral problems. And we shall submit that he is singularly privileged to do so, as he can tackle moral problems, and even the theory of morality, i.e. ethics, with the help of an approach and a set of tools alien to most philosophers and yet promising to deliver the technoethics that philosophers have not deigned to work on. To this end we shall put forth a value theory allowing one to weigh

*Paper discussed at the International Symposium on Ethics in an Age of Pervasive Technology held at the Technion-Israel Institute of Technology, December 21–25, 1974.

means and ends, and to conceive of moral norms in the image of technological rules.

1. The Inescapable Responsibilities of the Technologist

Look around and you will easily recognize the professions that have most contributed to shaping industrial society. They are the scientists, the engineers, and the managers (including the politicians). The former have supplied the basic knowledge, the engineers have used the latter to design their works, and the managers have organized the manpower that has implemented those designs. The outcome of their labors is well known: a new kind of society, one that may carry mankind either to a higher evolutionary level or to quick extinction.

And yet, paradoxically enough, the sociologist tells us that, by and large, scientists, engineers and managers—that is, the main artificers of modern society—do not feel constrained or fired by any extraprofessional moral and social responsibilities. In particular the technologists seem to be indifferent, when not outright callous, in the face of large-scale yet avoidable tragedy—unemployment, poverty, inequity, oppression, war, the mutilation of nature, the squandering of natural resources, or the debasement of culture.

Whether or not they have a hand in eliciting any calamities, most scientists, engineers and managers disclaim any responsibilities and shut their eyes to suffering and squalor. Worse, their position in society is such that they must play blind if they wish to perform efficiently. In fact no professional can do his job efficiently if distracted and upset by cries of distress: he must shut himself up in his office or laboratory if he is to get on with his specific work, be it investigating, designing, or organizing. Unless of course it is part of his job to relieve that very distress or at least avert future sources of distress.

True, some professional bodies have imitated the physicians and adopted moral codes ruling their own work. But most of these codes are limited to professional responsibilities, so they leave ample margin for irresponsibility. The minutiae are taken care of, the great issues pushed aside. Hence the scientist feels free to go on with his research come what may; the engineer, to go ahead with his projects go what or who may; and the manager, to push production or sales come or go what or who may.

Thus there are hardly any external constraints capable of being internalized and preventing the scientist, the engineer and the manager from engaging in morally wrong or socially injurious activities qua professionals. Whether or not he behaves well towards his fellow human beings is left to his moral conscience or, worse, to that of his superiors. Unfortunately conscience, usually wakeful in private and professional matters, is rather sleepy when it comes to affecting the anonymous lives of others.

Let me hasten to declare that I am not siding with the enemies of science and technology. There is nothing inherently wrong with science, engineering, or managing. But there can be much evil in the goals which either of them is made to serve, as well as in some of the side effects accompanying the best of goals. If the goals are evil—as is the case with genocide, the oppression of minorities or nations; the cheating of consumers, the deception of the public, or the corruption of culture—then of course whoever serves them engages in evildoing, even if not legally sanctioned as wrongdoing. In this case the scientist, engineer or manager becomes a mere instrument.

Instruments are morally inert and socially not responsible. Hence, when acting as a tool, the scientist, engineer or manager will refuse to take any blame except when failing to deliver the goods—though, inconsistently enough, he is not above accepting praise when successful. If taken to task he is apt to plead innocent or claim to have acted under duress (*Befehlnotstand*); he may go so far as to react indignantly. Obviously, he does this because of either excessive humility or excessive arrogance. In the former case he crawls under his superiors, in the latter he climbs above ordinary mankind; in either case he places himself beyond common decency.

The scientist, engineer or manager may well wash his hands but this will not free him from moral duties or social responsibilities—not only qua a human being and a citizen, but also as a professional. And this because, let us recall, they more than any other occupational group are responsible for the shape the world is in. You cannot manipulate the world as if it were a chunk of clay and at the same time disclaim all responsibility for what you do or refuse to do, particularly since your skills are needed to repair whatever damages you may have done or at least to forestall future such damages. In short the engineer and the manager, precisely because of the tremendous power they wield or contribute to building up, have a greater not a diminished moral and social responsibility. This being so, they had better face it.

2. *The Technologist Is Torn Between Conflicting Interests*

Suppose a ream of engineers is in charge of the design and construction of an industrial plant. What is expected from them? A lot, to wit:

M: The management expects an efficient and profitable plant.
W: The workers expect satisfactory working conditions.
N: The neighbors expect a pollution-free operation.
T: The professional colleagues expect a technologically advanced design, execution and operation of the project.
C: The consumers expect useful and reasonably priced products.

In addition, the unemployed expect a new source of jobs; the suppliers, substantial orders; the banks, a new client; and the government, another source of revenue or perhaps a new funnel for subsidy.

The engineer can ignore some of these expectations and demands but not all of them, the more so since they are not all mutually compatible. For example, if the management demands and obtains a minimization of costs together with a maximization of profits, then all the other groups affected by the project will be disappointed. Consequently the engineer will ignore some groups, favor others, and try to compromise with still others. Obviously in making decisions of this kind he poses and solves moral problems. And he does it by adopting, tacitly or explicitly, some moral code or other.

Every moral code boils down to a ranking of interests—or, to put it in a more polite way, moral codes rank values. To abbreviate, write '$A > B$' to designate the proposition "A is preferable to B", or "There exists an individual or a group to whom the value of A is greater than the value of B". Then our engineers are faced with, among others, the following moral codes:

Private interest morality: M takes precedence over all others.

Professional interest morality: T takes precedence over all others.

Public interest morality: $C > N > W > M > T$.

The choice among these possibilities will depend in turn upon the overall moral code of the decision makers—and who the decision makers are depends in turn upon the kind of firm and the kind of society. In the ideal society—which, alas, is nonexistent—public interest morality prevails, so that engineers and managers are, together with politicians, servants of the community. But let us not argue about this point: all that matters to our present concern is that every technologist, in any society, faces conflicting interests and makes moral decisions that agree with certain moral codes while disagreeing with others. In short the technologist—whether engineer or manager—is a moral agent even if his decisions and actions are regarded as immoral by those who are hurt by them. And, as we all know, the technologists can harm, either by misusing good technology or by employing technology that is inherently evil. But this last concept deserves a separate section.

3. Not All Technology Is Good

I submit that all pure science is good or at worst worthless since, by definition, it is concerned only with the improvement of our models of the world, and knowledge is a good in itself. On the other hand technology is concerned with human action upon things and men. That is, technology gives power over things and men—and not all power is good to everyone. Just think of thanatology or the technology of killing: the design of tactics and strategies of aggression, of weaponry and defoliants, of extermination camps and so on. All this is inherently evil by any moral code except of course that of the mass murderers. And whatever spinoff it may have by way of useful gadgetry this is by far outweighed by its evil effects—the destruction of human lives, the disruption of family and friendship bonds, the increase in aggressivity, violence and callousness, and the mutilation of the environment.

Surely individuals, groups and nations have the right to defend themselves from aggressors and oppressors, if need be with the force of arms. But if they leave defence, which is a political matter, in the hands of the thanatologist, he is apt to counsel attack as the best defense. And nowadays, with the emergence of a tightly knit international system, any local war can engulf a whole area and even the entire world—in which case, given the nuclear bomb stockpile, the human species could be wiped out. These platitudes bear repetition not only because we should do something about the danger of any war but because modern warfare is eminently technological and therefore a reminder that not all technology is good.

Being morally ambivalent, technology should be controlled instead of being allowed to develop unbridled in the interest of whatever group can afford it. In other words, the technologist must be held not only technically but also morally responsible for whatever he designs or executes: not only should his artifacts be optimally efficient but, far from being harmful, they should be beneficial, and not only in the short run but also in the long term. Don't tell me that only a free agent can be held morally responsible, so that a technologist acting under orders from above must be exempted from any blame: this was Adolf Eichmann's line of defense. If ordered to do something evil the technologist is free to refuse to take the order; if necessary he can leave his job, or he can sabotage his project, or he can fight it. Of course he may be punished for refusing to obey. But this is part of the game of human life—any kind of human life—in society—any society. The more responsible a position the more risky it is—but also the more rewarding.

The technologist is morally responsible for his professional actions because these ensue from deliberate and rational decisions in the light or the darkness of some moral code or other. The technologist is responsible *for* his professional work and he is responsible *to* all those affected by it, not only to his employer. A technologist intent on pleasing his employer only, with a total disregard for the interests of everyone else, is an accomplice or an instrument rather than a whole professional facing his entire set of responsibilities. Just as the good politician (whether or not he is successful) makes the right use of power, so the good technologist makes the right use of his knowledge and know-how, which is their use for the good of mankind. And this is not a piece of rhetoric, for, if we are to survive, we must try to avert the disasters, of increasing magnitude and in increasing numbers, brought about with the help of technology. I do not mean just the effects of inherently evil technology but also the morally wrong and the technically shortsighted use of potentially good technology. Suffice it to mention the covering of black soil by highways, the unchecked burning of fossil fuels, the degrading of forests into the shopping catalogues called newspapers, and the Great Air and Water Robberies.

Every large-scale technological project is apt to have a strong impact on both nature and society. (Just think of the biological and social changes brought about by the building of a dam.) Therefore, if the obnoxious side effects of any such project are to be minimized, its design should not be left in the hands of engineers alone, particularly if they are anxious to please their employers. The community affected by the project has the right to keep it under the control of other specialists—applied social scientists, public health officials, city planners, conservationists, etc—to the point of vetoing the whole thing if its negative effects are likely to outweigh its social benefits. It is a question not of slowing down technological progress but of preventing progress in one respect (e.g. engineering design) from blocking advancement in other respects.

Because of the close relationships among the physical, biological and social aspects of any large-scale technological project, advanced large-scale technology cannot be one-sided, in the service of narrow interests, shortsighted, and beyond moral control: it is many-sided, socially oriented, farsighted, and morally bridled. But none of this is possible as long as the technologist regards himself as a mere employee and hides behind the economic or political leadership. The technologist, to be a good technologist must regard himself as a trustee and a leader. In other words, competent, socially beneficial and morally inspired technology calls for global technocracy—the rule of experts in all fields of human action. But this is another story.

4. *Technology As a Source of Inspiration for Ethics*

It is easier to scold the scientist or the technologist for failing to live up to his moral responsibilities, than to recommend him the reading of a treatise on ethics for his moral edification. Indeed moral philosophy is underdeveloped, so much that it has ignored the special problems posed by science and technology. Moreover it won't be able to tackle these problems unless it learns a thing or two from science and technology. I shall proceed to indicate three lessons ethics has yet to learn from contemporary technology.

A first lesson is that the classical distinction between what is and what ought to be can no longer be maintained in the face of postbehaviorist psychology and of cybernetics. In fact we have learned that every control system, be it a furnace with thermostat or an organism endowed with a nervous system, has an *ought* built into it in the form of a set of final or goal states which the system tries to attain or to keep. Any such system behaves in such a way that its *is* approaches its *ought*, the size of the misalignment thus being reduced. Similarly motivation studies on higher animals have shown that they select a goal and proceed to attain it by successive trials, a successful sequence of such attempts bringing about the coincidence between what is and what ought to be. The lessons for ethics are that not all oughts are beyond reach, and not all of them are lofty ideals.

A second lesson ethics ought to learn from science and technology is that fact and value, far from being at odds, become blended in action, and that this synthesis is consecrated by certain action theories such as statistical decision theory. But before ethics can learn this lesson it must purge itself from the idealistic view that value, or the good, is an autonomous entity. A fresh look at the behavior of higher animals shows that values are not entities but properties the organism assigns certain entities. Thus there is no such thing as The Good, but rather a whole set of things and events that we, or other organisms, value as good. (In other words, every value is the value of a valuation function the domain of which is the set of things or events.) Once values are recognized as an outcome of the valuation activity of an organism, they cease to be disjointed from facts: they become aspects of certain facts. Moreover, in decision theory values join with another property of facts, namely their probability. In fact the very definition of a rational decision as that which maximizes the expected utility (or subjective value) involves such a synthesis of fact and value.

A third lesson moral philosophers can learn from scientists and technologists is the way to conceive of moral norms. Traditionally these have been regarded not as statements but as prohibitions or exhortations: as such they would both hover above the world of fact and creep under the world of reason. So much so that they would have to be handled by a special logic—deontic logic. This view is unacceptable to a naturalist who regards reason and value as so many hypostatizations of certain organismic activities, namely reasoning and valuing respectively. It is unacceptable to a rationalist or an empiricist, both of whom would like moral rules to be subject to some control other than authority or tradition. Indeed an imperative such as *Thou shalt do x*! Or *Thou shalt not do y*! looks impregnable to reason and experience.

We have had enough groundless orders: it is high time to treat ourselves as rational and responsible beings capable of adopting, discussing and rejecting grounded rules of conduct. Which is exactly what technology does. Thus, when an electrical engineer is assigned the task of designing a power plant, he does not issue a command such as 'Let electric power be!'. Instead, he gets hold of his scientific knowledge and his wits, and produces a design that can resist a critical examination. His final recommendations or norms are not blind starting points but thoughtful outcomes of his work. Accordingly his conclusions are not categorical imperatives but conditionals of the forms *If A produces B, and you value B, choose or do A*, and *If A produces B and C produces D, and you prefer B to D, choose or do A rather than C*. In short, the rules he comes up with are based on both fact and value. I submit that this is the way moral rules ought to be fashioned, namely as rules of conduct deriving from scientific statements and value judgements. In short ethics could be conceived as a branch of technology. How this project can be implemented will be outlined in the following section.

5. *Ethica More Technico*

A technological rule boils down to a formula of the form *To get G do M*, or *To avoid G refrain from doing M*, where '*M*' stands for a means to a goal abbreviated '*G*'. In either case the rule, far from being arbitrary, is founded on some natural or social law of the form "*If M then G* [either always or with a fixed probability]". Because there are no isolated variables there are no isolated events either; therefore every means and every goal is actually a conjunction of items. In particular so are goals: because every action has side effects, every goal is accompanied by side effects, some of which are disvaluable. Consequently it is convenient to include the side effect(s) *S* in the preceding formulas, namely thus:

Law: If M then G and S.

Rules: To get G and S do M; To avoid G and S refrain from doing M.

Such rules, let us repeat, combine scientific knowledge with explicit valuation. The former consists in the underlying law. And the axiological component consists in that *M*, *G* and *S* are not just facts but facts evaluated by somebody. I submit that a *rational moral rule* has exactly the same structure as a technological rule, in that both rest on scientific laws and explicit valuations.

Drop the requirement that a rule should be founded on a scientific law and you get a rule of thumb, whether technological or moral, detached from the body of scientific knowledge and therefore hardly defensible or criticizable except as regards its efficiency. And drop the requirement of explicit valuation, and you may underrate side effects and even confuse means and ends. In particular, fail to assess the means and you get either a costly technology or an inhumane morality. Furthemore I submit that the very setting up and use of technological rules is inseparable from ethical considerations. Indeed it would be just as technically mistaken as morally wrong to aim for goals that are outweighed by either harmful means or negative side effects. Many a technical mistake is then a moral slip. In the age of technology, erring and sinning are becoming equivalent.

The expected rejoinder from the antiscientific quarter is that no matter how much science and technology are injected into decision making, there will remain an arbitrary and irrational residue, namely valuation. This objection rests on the false presupposition that values transcend facts (recall section 4) and on ignorance of the fact that valuations are made explicitly and daily in the fields of advanced technology and business administration. Moreover decision theory, which involves values, is occasionally used in those fields. What is true is that the values there in question are subjective values or utilities, while a fully rational action theory ought to use objective values as well. If a theory of objective values were at hand we would be in a better position to make rational decisions and argue about them. That such a theory is feasible will be shown presently.

6. Towards a Value Theory of Means, Goals, and Side Effects

We proceed to outline a value theory that can serve as a basis for weighing means, goals, and side effects, and thus help to make or adopt rules of conduct both technically feasible and morally right.

Let S be a set of objects that can be evaluated by someone in some respect. The members of S shall be things or states of things or events—in particular human actions. Some members of S shall have unique inverses, others won't. For example, if 'b' stands for giving, its inverse $\bar{b}$ stands for taking away—not for not-giving, which is no action at all. And if 'b' stands for writing, '$\bar{b}$' will stand for erasing a piece of writing. When an element x of S happens to have a unique inverse $\bar{x}$, then either x followed by $\bar{x}$ or $\bar{x}$ followed by x will equal the neutral or inane element e. Moreover, as has just been intimated, some members of S combine to form compound objects. If x and y are members of S, and they actually combine, the object resulting from their combination will be denoted as '$x + y$'. We shall assume that this operation, where defined, is associative, i.e. that $x + (y+z) = (x+y) + z$. Furthermore we assume that all members of S are idempotent, i.e. $x + x = x$ for any x in S. But $+$ is not defined for every pair of objects in S. That is, some compounds may not exist. For example, if b is giving birth, then $\bar{b}$ is killing, so that $b + \bar{b}$ is equal to e; but in this case $\bar{b} + b$ is not defined: it is not a member of S, hence not a valuable object or even a disvaluable one. In short, + is, like $^{-}$, a *partial* operation on S. As a consequence + is not commutative, i.e. $x + y$, even if it exists, is not necessarily equal to $y + x$. When $x+\bar{x}$ happens to exist, it equals the neutral element e, which is assigned zero value. If the operations $^{-}$ and + were total, not partial, i.e. if S were closed under inversion and composition, the structure $\langle S, ^{-}, +, e\rangle$ would be a group of idempotents. As it is, it will be said to be a *partial Boolean group*.

Surely this is too coarse an algebraic structure for the set of objects of valuation, if only because (a) it makes no provision for the case when an object has more than one inverse (e.g. several antidotes to a given poison), and (b) it lumps disjunction and conjunction into the single operation +. A partial lattice is a more suitable structure and the one I adopt in a more detailed investigation in course, because it distinguishes options from conjunctions and it incorporates idempotence automatically. However, the simpler structure will do for our present purpose, which is just to show how value theory can serve as one of the legs of a rational ethics—the other leg being scientific knowledge.

Next we introduce a valuation function V that assigns every object x in S a value $V(x)$ that, for the sake of definiteness, may be taken to be a real number. (The converse is not true: several objects may be assigned the same value. That is, V is not 1:1.) This function is defined as follows:

(i) *the neutral element has zero value:* $V(e) = 0$;

(ii) *good and evil neutralize one another:* if both x and its opposite $\bar{x}$ are in S, then $V(x) + V(\bar{x}) = 0$;

(iii) unlike utility, *value is additive*: if x and y are different objects of valuation, and $x + y$ is defined, then $V(x + y) = V(x) + V(y)$.

It is easily seen that $V(x + x) = V(x)$ and $V(x + e) = V(x)$ follow. Also $V(\bar{x} + y) = - V(x) + V(y)$ is a theorem, and $V(\bar{x} + x) = 0$ a corollary of the latter. For our present purposes no further consequences are needed.

Let us now apply our value theory to ethics, i.e. to the evaluation of means, goals, and side effects. To this end we adopt the convention that, if p and q are any propositions, then $\bar{p}$ is the negation of p, and $p + q$ the disjunction of p and q. Further, we assume that propositions obey the ordinary propositional calculus, and that the valuation function V applies to them, so that '$V(p) = v$' means that p is worth v.

Consider now an arbitrary rational rule, whether technological or moral, from our value-theoretic point of view. It will fit the schema "To get G and S do M", the ground or foundation of which is the law schema "If M, then G and S". Call f the value or effectiveness of this cognitive means and g the value of the goal G. That is, set

$$V(M \Rightarrow G\&S) = f,\ V(G) = g.$$

Our task is to unpack the separate values of M and S, and to relate them to f and g. This is done with the help of logic. In our notation,

$$M \Rightarrow G\&S = \bar{M} + (G\&S), \text{ and } G\&S = \bar{\bar{G}}\ \&\ \bar{\bar{S}} = \overline{(\bar{G} + \bar{S})},$$

whence

$$M \Rightarrow G\&S = \bar{M} + \overline{(\bar{G} + \bar{S})}.$$

By our calculus

$$f = V(M \Rightarrow G\&S) = -\ V(M) - V(\bar{G} + \bar{S}) = -\ V(M) + V(G) + V(S).$$

Hence finally, calling $V(M) = m$ and $V(S) = s$, we obtain a central theorem:

$$m + f = g + s.$$

That is, the practical means combined with the cognitive means balances the goal combined with the side effect. Shorter: the total input or cost equals the total output or benefit. Thus valuable knowledge (of the means-end relationship) can offset the cost of the practical means, whereas defective knowledge calls for a greater investment in practical means. (Moral: Support research.) For example, if $f = 1$ then $m = g + s - 1$, whereas if $f = -1$ then $m = g + s + 1$, which may be ruinous. And in all cases the goal is worthy just in case the means too is worthy. Indeed $g > 0$ if and only if $m + f > s$. Good breeds good, and evil, evil. The better the means, the more valuable the goal.

Our final formula suggests the following rules of conduct:

R1: To assess a goal evaluate it jointly with the side effect—i.e. estimate the total value $g + s$.

R2: Match the means to the goal both technically and morally: employ only worthy practical means and optimal knowledge.

R3: Eschew any action where the output fails to balance the input, for it is either inefficient or unfair.

In sum, instead of accepting rules of thumb in the realm of morals we can and should try to form them in the image of technological rules, i.e. on the strength of factual knowledge and objective valuation.

Summary and Conclusions

(i) Unlike pure science, which is intrinsically valuable, technology can be valuable, worthless or evil, according to the ends it is made to serve. Consequently technology must be subjected to moral and social controls.

(ii) Evil technology can be eliminated only by discarding evil goals as well as by avoiding obnoxious means. And the misuses of good technology are corrected and averted not by slowing down all technological research but by promoting good technology and rendering it morally and socially sensitive.

(iii) The technologist, just as everybody else, is personally responsible for whatever he designs, plans, recommends, or executes. Hence he is just as subject to praise and blame as anyone else—actually even more so because of the rational character of his decisions.

(iv) The technologist is responsible not only to his employer and his profession but also to all those likely to be affected by his work. And his primary concern should be the public good.

(v) The technologist who contributes to alleviating any social ills or to improving the quality of life is a public benefactor. But he who knowingly contributes to deteriorating the quality of life, or deceives the public, by devising worthless or dangerous products or disseminating false information, is a public criminal.

(vi) Because no single specialist can cope with all of the many-sided and complex problems posed by large-scale technological projects, these should be entrusted to teams of experts in various fields, including applied social scientists, and should be placed under public scrutiny and control.

(vii) Because of the strong impact technology is having on both society and the environment, the technologist should share power with the manager and the politician. Global technocracy, or the rule of experts in all fields of human action, is not a threat but a promise, particularly if answerable to the public.

(viii) Technologists should tackle their own moral problems rather than pretend that they can be transferred to managers and politicians. Moreover they should contribute to the overhauling of ethics, attempting to construct a technoethics as the science or right and efficient conduct.

Mario Bunge

McGill University
May 21, 1975