

PHILOSOPHICAL SPECULATIONS ON SYSTEMS DESIGN

by

C. West Churchman

Working Paper No. CP-355

Center for Research in Management Science

University of California

October 1973

# PHILOSOPHICAL SPECULATIONS ON SYSTEMS DESIGN

by

C. West Churchman

University of California

Berkeley, California

## 1. Introduction

This paper is a philosophical discussion of a very old idea, namely, that humans can consider before they act and that potentially their consideration can produce a better act than the one that would have occurred without any prior consideration. In other words, why should you think before you leap?

The philosophical problem is relevant to the many efforts that are being made today to improve our human condition, by planning, operations research, management science, systems science, information systems, technological assessment, educational evaluation, and so on. Since I am not concerned about the specific characteristics of each of these efforts, but rather about their general philosophical justification, I need a generic label, which I'll call "systems design" (SD). I should emphasize that "systems design" denotes any plausible effort to prepare for decisions by means of prior inquiry and, therefore, also includes applied social science, applied psychology, applied theory of altered states of consciousness, and so on.

The basic stipulation of the philosophical inquiry of this paper was given by Descartes in his Discourse on Method: all issues are open to discussion, and none shall be taken as settled just because "everyone" agrees. The phrase "few can doubt that . . ." is to be taken as a signal, not of acceptance, but of the obligation to listen to the few.

## 2. Systems design

My purpose here will best be served by first defining "systems design" in a manner which will then enable me to discuss its philosophy: systems design is implementing improvement in social systems by means of the best available method of inquiry.

This definition has the relevant advantage of highlighting the philosophical issues of SD. Taking the ideas in reverse, the meaning of "best method of inquiry" is derived from the idea of how knowledge is acquired, or what philosophers call epistemology. The meaning of a "social system" is derived from the meaning of social reality, or what philosophers call social ontology. The meaning of "improvement" is derived from the meaning of the good, or what philosophers call ethics. The meaning of "implementing" is derived from the meaning of translating knowledge into action, or what philosophers call pragmatics.

## 3. "Scientific method"

It will be appropriate to abbreviate "best method of inquiry" by "scientific method" because of the historical tradition that "science" is mankind's body of knowledge however acquired. Hence "scientific method" is not the special province of any discipline, but rather is the appropriate method of gaining knowledge in any area of human concern.

We begin with a philosophical discussion of scientific method, and proceed to the other ideas, but we'll see that the route is not linear: what we may learn about social reality, improvement, or implementation will make us revisit the meaning of scientific method. Since this non-linear journey may be an unusual one for you, it may help to describe at the outset how we'll proceed. We'll start with a fairly obvious and old idea about the meaning of scientific method, and then, by considering one meaning of social reality, we'll find that this first definition of scientific method is inadequate; and we'll try to develop another, richer one. This second definition will lead to a revision of the definition of social reality, which will carry us into the domain of improvement where a first definition will be tried, which pushes us back to reconsidering scientific method and social reality again. If you guess that the journey has no end, you're a good guesser: learning philosophy is science learning a restless process of thought.

Scientific method is variously defined in the literature of science, often in terms of the way the writer was trained to do his Ph.D. research. But, as I said, here scientific method is taken to refer to any of the ways in which human knowledge can be produced. In the Western world, two such ways have been identified as especially appropriate: reason, or the thought process of inferring consequences from already known principles, and observation, or the empirical process of using our senses to tell us what is happening. In the seventeenth and eighteenth centuries, philosophers battled over which process dominated; those who said it was reason were called rationalists, those who said it was observation were called empiricists. In this century there are few pure rationalists

or empiricists; the common tendency is to regard scientific method as some judicious combination of these processes. But the important point is that each process must follow certain rules which legitimize the results. Thus, reason must follow orderly rules of deduction, must not indulge in blatant contradictions, must start with plausible or even highly defensible assumptions. Observation must be conducted carefully, according to well-established procedures where the observer does not inject his own biases into the results.

We might sum this up by saying that both reason and observation can be deceptive, but that there are clear and explicit ways of minimizing deception and that these ways are available in principle to all humans, excluding perhaps a subclass who are "mentally deficient."

#### 4. Social reality

All the above sounds very fine and should have us nodding in agreement, which like most agreeable statements is already an indication to the philosopher that it says very little and what little it says may well be wrong. I've bored you with it because it has so often been said with no philosophical justification. The truth is that "scientific method" is an extremely elusive concept. To begin to understand its elusiveness, especially in connection with SD, we should examine the second on the list of SD concepts: social reality. Here again the literature abounds in definitions, of which we can select one: "social reality" consists of a group of humans who are taken to be goal seekers. Their goals really exist, and their actions can be evaluated in terms of whether they do or do not attain these goals. Hence the group is taken to have real choices among a set of possible actions. Finally,

and this is the crux of the matter, it is possible to consider an action without taking the action; the consideration consists of evaluating the action in order to estimate its likelihood of attaining the group's goals. The evaluation is "valid" to the extent that what it asserts about the outcome of the action would really occur. It should be pointed out that the "evaluation" can be made of actions already taken, e.g., of existing government programs. Thus so-called "cost-benefit" studies of educational or health programs all assume this version of social reality.

It can readily be seen that these evaluations are "counterfactuals" in the sense that the results are in the form: if the social group were to do so-and-so, then such-and-such an outcome would occur. For example, if this program were to be funded (or eliminated), then such-and-such benefits and/or costs would occur for society.

Now one might try to "test" these counterfactuals by conducting "social experiments" in which we systematically change certain sectors of society in various ways, e.g., by experimental design techniques, and attempt to use traditional hypothesis-testing methods. But it is highly questionable whether such social experimentation is feasible, or, if so, desirable. It often is far more advantageous to determine the validity of a counterfactual without disturbing social reality. Hence, direct observation and reason may not constitute the appropriate scientific method for systems design.

But is there really any difficulty here? After all, engineering and indeed all applied sciences (medicine, agricultural science, etc.)

are in the business of trying to assert valid counterfactuals. The methodology seems to be fairly straightforward. One strives to develop a theory about a specific sector of nature; the theoretical assertions say that if such-and-such is the state of a system at a given time, then so-and-so will be the state at another time. The theory is tested by observing how closely it maps "reality," by using the judgment of a community of investigators who develop standards for accepting or rejecting theoretical assertions. Once a theory is accepted, one can then make "counterfactual" assertions of the sort that SD requires. Of course, there are many technical problems involved in the process; logicians, for example, have worried about the epistemological status of counterfactuals. Furthermore, no matter how often and thoroughly a theory has been tested, it may yet fail under new circumstances. But on the whole, relying on counterfactual assertions of applied science seems to work out fairly well, as anyone who steps onto a large jet must feel.

The above account is not so obvious a description of scientific method, but is widely accepted as an adequate description by applied scientists. But from the philosophical point of view, the account is unacceptable for SD. The reason it is unacceptable is what philosophers call the "self-reflective" paradox. You will note in the account that the "community" of applied scientists is said to establish the standards of acceptable theory. But the community is a "group" and furthermore a group that is seeking a goal: to establish reliable counterfactuals. Such a community of applied scientists is therefore one example of SD's "social reality." The SD question addressed to such a group is whether

its choice of actions (i.e., its "scientific method") is to be taken as the best available, or at least as adequate as any other choice. Since the method chosen is very dependent on group agreement among trained applied scientists, the method is certainly suspect from an SD point of view; for example, a great deal of planning and operations research consists of going beyond organizational agreement to policies which no one in the organization has even imagined.

What is "paradoxical" about all this? Why, we apparently need to know what SD is in order to define SD, because SD is the application of "scientific method," and the correct "scientific method" is a legitimate subject matter for SD. I should hasten to point out that the self-reflective paradox is a philosophical issue and need not concern all SD practitioners; after all, it is often sound SD to ignore some problems of an organization, even though they exist and are felt to be critical by many people in the organization. Thus, we have arrived at a definition of a philosopher: someone who is concerned about the self-reflective paradox and attempts to understand it and possibly "solve" it.

I should also point out that SD is not the only area of investigation in which the self-reflective paradox appears. Psychology is another example, since there is one viewpoint of the world in which everything has a psychological base, including the sciences, and therefore including psychology itself. Indeed, this characteristic of an area of inquiry, its susceptibility to the self-reflective paradox, should be regarded as the characteristic which makes the area "basic" or "fundamental."

What is the "solution" of the self-reflective paradox? The soundest reply seems to be that the paradox does not belong in the world view of social reality which assumes that there are problems and solutions. In other words, the paradox does not fit well into the original definition of SD's idea of social reality. What was deficient in the first viewpoint of social reality was SD's own reality. According to this viewpoint, SD stood aside, observed and evaluated the possible actions of the social group, much as a laboratory scientist observes the objects on his bench. But this viewpoint doesn't work because SD is a social reality itself and cannot possibly "pull aside" from itself.

##### 5. Improvement

We can gain a better sense of how to modify the earlier concept of social reality if we extend our travels into the third domain of inquiry, "improvement," and the question of how to judge that one action is better than another. This is the age-old question of ethics, and the domain has many confusing pathways. Certainly no one should try to enter it who has not read in the history of ethics since he needs all of the wise guidance he can get, and much of the wisdom of ethics comes to us from our past; it's not an area where only recent citations are relevant.

Ethics is a very practical subject; indeed, it can be considered as the investigation of the practical. Thus, SD must consider the practical questions of who its clients should be, what should be done when the client is seeking the wrong goals, how much to charge for its services, how much of the results should be revealed, and so on.

All of these are ethical issues which are essentially practical: they are issues about the practice of SD.

In this story I intend to restrict the conversation about ethics to one central theme, namely, the scope of concern about improvement. An SD team may be studying the advertising policy of a firm, say, one that sells cigarettes or machine guns. The ethical issue of scope is whether "improvement" means "increased sales per dollar of advertising" and "increased kill capability," or something much broader like "contribution to health or peace."

#### 6. The paradox of SD data

The issue has a technical base which can help us understand its complexity. Consider, for example, a typical operations research study of inventory policy. The purpose of such a policy is to minimize the "costs" of maintaining inventory against various kinds of demand. Too large an inventory runs the risk of tied-up capital, high taxes, obsolescence, and so on. Too little inventory runs the risk of shortages and hence serious delays or loss of patronage. The proper policy is therefore a balance between these (and other) opposing costs. In operations-research practice, measuring the relevant costs is at least as important as modeling the system mathematically. The clue as to how to measure costs can be seen if we consider one such cost, the cost of funds tied up in inventory, sometimes called the cost of capital. The idea is that a dollar spent on an item that sits on a shelf waiting for future demand is a dollar that cannot be used for other opportunities, like investment in bonds, where it would yield an annual interest. Hence, the cost

of capital tied up in inventory is an "opportunity cost," and its measurement consists in estimating the maximum return that would occur if the inventory dollar were to be used in the best alternative opportunity.

But how are we to determine the best alternative? The answer lies in the operations of another subsystem of the organization, the financial system which controls the flow of cash. This system can also be studied by SD, and one result of the study might be an estimate of the optimal use of a dollar which has been released from the inventory system; in this case, the financial system would have provided a measure of the cost of capital in the inventory system. The italics are supposed to emphasize that the cost of capital is not a general cost operating across a whole firm or industry, as is sometimes erroneously assumed, but rather a cost that is specific to each subsystem. What should happen to a dollar released from inventory need not be what should happen to a dollar gained by reducing the work-force, for example.

But now another paradox arises. In order to study the inventory system, we need to measure costs, and specifically the cost of capital. To measure the cost of capital we need to study the cash-flow system. But to study the cash-flow system we must also measure some costs (transaction costs, risk costs, etc.) which will inevitably require us to study some other subsystems. The point is that SD is restless; it cannot legitimately measure costs in terms of historical outlays because these outlays may have been based on wrong policies. If so, and SD uses historical data, it is measuring incorrectly.

The same remarks apply to demand. It is very tempting to use past records of demand to forecast demand that will be made on the inventory system. But SD should never find itself in such a forecasting business because if the subsystem which has historically generated demand has been operating incorrectly (and very often it has), then the forecast based on history is an incorrect measure of demand, no matter how elegant the forecasting statistics may be.

This very practical yet philosophical point is often missed in many forecasting studies, where the purpose of the investigator seems to be one of predicting future events on the basis of past patterns of events, or--what is apparently the same thing--on the basis of expert opinion. From the point of view of SD, one has to ask what subsystems are influencing future events and whether these subsystems are judged to be properly designed. It is ridiculous, for example, to forecast a specific technological breakthrough, like the feasible hover-craft or cure for cancer, without considering the linkages of these system sectors to other subsystems of the world.

The same remarks apply to all proposals for "information systems" which store data of various kinds. An information system which is to be used for improving social systems must contain a total subsystem-linkage image, or else the designers cannot judge which "data" really constitute information.

One might be tempted at this point to say that SD is not really in the business of improving social systems by identifying real improvements, but rather its business is to provide "if, then" estimates: "if we were to develop a cheap hover-craft, then it

would take 10 years and 30 billions of dollars." The idea is that SD might then avoid the paradox of its requiring a knowledge of all relevant subsystem linkages in order to acquire any meaningful data. But this attempt to make SD's task more modest doesn't resolve the paradox at all, since all estimates like "ten years" and "30 billion dollars" must be based on the way in which the subsystems (finance, management, research, etc.) ought to behave. An "if, then" statement based on a very ineffective management subsystem would put SD in the business of supporting bad policies, i.e., into the devil's business.

The philosophical paradox of SD arises because we apparently can see no end to the "linkages" between subsystems which are the sources of the data which SD requires. In the inventory example, the cost of capital is a linkage between the inventory system and the cash-flow system. The cost of liquidity is a linkage between the cash-flow system and the investment system. A firm's social responsibility is a linkage between the firm's total financial system and the social system in which the firm exists. And so it goes--apparently either into an infinite regress or into a vicious circle.

It is important to point out the difference between the technical problem of an interlinked system and the paradoxical aspect of interlinking systems. It is no news that many machines must be designed by considering the interplay of their parts: there is no such thing as a generally optimal generator, but rather for each design we can estimate the specific generator that does the trick. But all such machine designs assume that certain principles of design

are appropriately based on past tests. If we think of the whole world as a machine--or a brain, then we can use brain or machine design principles to help us solve the technical problem of linking together subsystems. But such a technical solution by no means resolves the paradox because we need to know some important things about the system which decides that the world is a machine or a brain. Is such an inquiring system operating optimally? The paradox consists of pointing out that there can be no place on which one can stand to answer such a question without again raising the same question.

#### 7. Implementation

Paradox, if it is meaningful, always carries us to other considerations, in this case to yet another domain which I called implementation. If we had allowed the first part of our travels to extend this far, we would have said that implementation would consist in transforming a thought-out solution of a problem into action so that the social group would do what the SD team would propose and, most important, it would do so because the SD team had thought out the solution and influenced the group's behavior. So put, implementation would then be another aspect of SD's social reality, in which SD sees itself as really causing a group to veer towards the correct solution. More generally, SD's idea of social reality includes the ability of appropriate inquiry to influence action in appropriate ways.

But this happy viewpoint won't do in the light of what has just been said. For one thing, there must always be a cost of implementation, which is the opportunity foregone for the social group and SD to have done something else. Thus an SD team may devise a

wonderful new inventory system, and enlist the assistance of managers and workers in installing it. But all of them might have been engaged in inventing a new product, or analyzing prices, or whatever.

Hence, "implementation" for SD cannot simply mean causing a group to put an SD thought into action, because this may not be the best activity for the group to follow. We are being frustrated again by SD on SD.

#### 8. Back to "scientific method"

To lessen the frustration, let's try another theory of pragmatics which avoids the trap of problem-solution. This theory starts with the premise that all scientific inquiry--and therefore SD--must begin partially blind and that this blindness in no way detracts from the "science" of the inquiry. Thus, SD may study inventory, but admits that it doesn't know whether this is the "best" area for inquiry, admits that it isn't sure about the appropriate costs and demands, but uses its best judgment to make its first estimates. SD then admits that the "solution" may not be a solution at all and need not, for example, be better than the existing policy. But still implementation will be tried, i.e., the possibly erroneous thought process will be put into action if possible, not because the resulting social change is an improvement, but because thereby SD may learn, be less blind, see more of its surroundings and understand itself better.

Again, this is not a new idea. It is quite reasonable to view the whole history of the experimental sciences in this manner. In the early stages, e.g., when Galileo wrote his Dialogues, both theory

and fact were very uncertain. Indeed, Galileo tested his very plausible theory of unobstructed rolling bodies with a defective clock. Had the history of experimental science stopped there, perhaps we would never know whether to believe his test or not. But the process did continue and some theory and some facts became quite solid, though never unshakeable, as the story of Einstein shows. Every finding is subject to further investigation; those findings that continuously "work out" and, more important, both support and are supported by other findings become the lessons that have been learned well and the basis for further learning.

It is important to notice that this "learning" account of scientific method shifts the picture of social reality away from the eventful problem-solving imagery to a wholistic-process imagery. Thus, "improvement" is no longer to be conceived as a specific change for the better in one sector of society, but as a property of the whole system. It follows that we must search for ways of identifying whole-system improvement if we are to make any philosophical sense out of SD's enterprise.

#### 9. "Improvement" as Progress

Here again, history comes to our aid. We are searching for a concept which connotes the gradual maturation of the human system, and the nineteenth century in particular offers us "progress." Having introduced the word, I should hasten to add that it does not mean "growth," nor is it linear. It stands for the continual process of change of the human condition and for the assumption that the overall property of this change can be from the worse to

the better. Progress is not inevitable; its existence requires the existence of a human will, the will for betterment. The concept of progress, in various forms, underlies all efforts to understand how the plight of the human world can be lessened; to assume that it can be lessened is to assume that progress is really possible. What is needed to make it possible is "global SD."

Until recent decades, it seemed possible to consider science as a prototype of such whole-system development, and even to suggest that the precision of certain physical constants--like the velocity of light in vacuo--could be used as surrogates for the real progress of all science from, say, 1500A.D. to the present time. The experience and attitudes of more recent times indicate strongly that "advancement" in science and technology does not necessarily imply progress in the total human system because other ideals of peace and freedom must also be considered. Thus, we are no longer very much impressed by the precision and reliability of physical measurements since this so-called improvement in a subsystem of society can be used by other subsystems for purposes of exploitation and destruction. Naturally, a physicist could respond that the fruits of his inquiry must be used well and that "it's up to the political or managerial system to make sure that they are." But SD can make no such facile reply: it's up to SD to show that a proposed subsystem improvement, like greater precision of measurement or more accurate theory of natural phenomena, is a real whole-system improvement.

10 "Social reality" and "implementation" revisited

Thus our philosophical inquiry into the meaning of improvement leads us back to the meaning of social reality: is mankind an ideal seeker, a living form which can strive for the betterment of its condition?

In the philosophy of SD, the answer must be "yes," else the SD enterprise becomes a sardonic joke, a struggle to help improve sections of society when the struggle as a whole is a grand illusion. But the answer leads us again to "implementation," which now has a dramatically different meaning from the earlier one. Now the central problem of implementation is to design a society in which the possibility of progress, i.e., ideal seeking, is made secure.

At this point our speculations could readily become reactionary, reacting either to Descartes in the seventeenth century or to St. Paul in the first. One often hears these days that global-systems simulations are only "first approximations," crude and wrong in many ways. But both Descartes and Paul realized that if you believe you have a "first approximation," then you also have an idea about the second, the third, and the "limit," else your "first" is not an "approximation" at all: "Now we see through a glass darkly, but then face to face. Now I know in part; then I shall understand fully..." To Paul, the "basic" aspects of mankind's progress lie in faith, hope and love. To Descartes, they lie in the proof of the existence of a benign Supreme Being. Most global simulators do not care to resort to such reactions to our past and indeed ignore them completely and without reflection. None of them seem to wish to

forecast what society would be like in the year 2020 if mankind's faith, hope and love with respect to God were restored, although this may be the most important forecast of all. The philosophical point is that none of us who concern ourselves with SD have a sound basis for choosing between Paul and an MIT simulator. If we could only have faith and hope in a Supreme Being, as did the seventeenth-century rationalists, then the implementation problem of SD would seem to be solved: God helps those who help themselves. But the spirit of questioning which characterizes scientific method does not permit any such facile response to the problem of guaranteeing our destiny.

Thus it seems safe to say that today we have no satisfactory theory of implementation from the philosophical point of view. And until we can arrive at a more satisfactory idea of mankind and progress, we must admit we do not have a sound basis for evaluating social change or for evaluating social studies aimed at improving sectors of our society. No university faculty can claim an ability to judge the real quality of an SD research by either students or themselves.

Such a "conclusion" might be taken to be negative, especially by the cynics, but I take it to be altogether very exciting and hopeful. If it were to be more widely recognized, then we might turn away from niggardly criticisms of inadequate research techniques to the more central issues of human progress. Among other things, we'd encourage Ph.D. candidates not to be specific and narrow but broad and comprehensive.

11. "Social reality": the individual

But such preaching needs to be checked by reflection in this philosophical journey. The need arises from the reflection that the whole account of progress and social reality leaves out a very important aspect of social reality: the individual. We have been speaking of larger and larger worlds extending to the globe and all future and past globes. These larger images may be real, but so are you, in your here-and-now life.

In the first version of social reality, the individual was seen as a goal seeker; he was, in fact, the limiting case of a group, a group consisting of one person, whose goals presumably were unified since he had no one to combat. In economics, the individual becomes an entity who carries about a utility function for the outcomes of his action. But this "utilitarian psychology" is utterly naive and philosophically suspect, just as is the idea that each individual is an "information processing" entity. The naivete is well documented in depth psychology, which claims that the unconscious mind has "values" that influence a person's life in critical ways and which conflict with his conscious values. The philosophical suspicion arises with the very plausible speculation that the greatest values of all are two ineffables: the lonely and unique individual and the unique relationship between two such individuals, which we call love.

It is really remarkable how difficult, subtle and exciting the philosophical issues become once we try to view them in terms of the individual. Kant's "never-ending awe" extended outward to the starry

heavens and inward to the soul. Nor does the immensity of the world or the universe in the least detract from the importance of the inner man; after all, immensity itself is a creation of man's psyche. Most living creatures don't see an immense world; and, after all, it was Kant's inner psyche which was filled with awe.

Of course, you can stop your philosophical inquiry into the psyche anywhere you wish, at the surface in positivistic behaviorism, a little deeper in attitudinal or motivational psychology, deeper still in analytic psychology, and so on. The philosopher wishes to travel as deeply into the inner psyche as he wishes to travel broadly into the whole system, else "social reality" is only partially explored. This is why SD is not a "discipline" from the philosophical point of view.

Our ignorance of the inner world is immense, and our clues for enlightenment must be found in the most obscure writings of the past, including mysticism and myth. But even in the midst of obscurity we can discern one theme which makes us doubt the philosophical exposition of SD given earlier. It is that individuals, as individuals, should not be classified with respect to society's decisions regarding their lives. The individual is a sacred value; as Kant says, each individual is an end-in-itself which ought not to be used merely for the sake of other ends.

I am less interested here in defending this psychological speculation than in pointing out its consequences, if true, for the philosophy of SD. But I should mention that the speculation is and has been widely held by humans of all walks of life; it is a

cornerstone of many religions and doctrines of morality. What it implies is that classifications and comparisons of individuals for the sake of rewarding them, punishing them or otherwise changing them in ways that they themselves do not will to happen are all immoral. One consequence, thoroughly devastating with respect to our earlier speculations about SD, is that the so-called principle of trade-offs of values is absolutely inappropriate with respect to the sacred value of the individual. The trade-off principle states that in determining the decision or policy to be made, we should determine how much of one value (e.g., speed on highways) to trade off for another value (e.g., safety). The principle is the basis for a great deal of SD, e.g., in cost-benefit or cost-effectiveness analysis, mathematical programming, and a whole host of similar techniques. The philosophical speculation says that if an individual is used as a means only, so that his sacred value is sacrificed, then the loss of this value cannot be atoned by trading it off by another value. For example, so runs the speculation, it is immoral to compare the value of a soldier's life versus the value of defending the country. Or, if you deliberately murder another human for the sake of your own goals, you cannot "trade off" this sin by donating a million dollars to your favorite charity. One way to remove sin is through forgiveness; but the methods of attaining forgiveness are lacking in much of SD practice today.

Again, I need to remind you that this speculation is a part of the philosophy of SD, which takes on the joy and frustration of questioning the axioms of SD practice. There is no need as yet to

go out and shake fingers at SD practitioners who are busy trading off values; there is a need, on the other hand, to question the foundations of such practice if SD is to remain healthy. Also, it should be recognized why many people reject and are even horrified by many SD recommendations; these critics are not stupid, but rather believe in a totally different picture of social reality. Nevertheless, in my mind's eye I can see the unreflective SD man standing on the right of the king on judgment day and being completely puzzled by the divine pronouncement, "as you did it to one of the least of these my brethren, you did it to me." Surely, he will say to himself, the optimal depends on how many and how important.

Now if you were a well-trained SD type, you might be tempted to sneak around the difficulties presented by the values of the inner life by simply agreeing to write them into the constraint equations: maximize net social benefit subject to never treating anyone as a means only. You can do this, of course, but you'll lack philosophical justification if you do. For one thing, in most models we can readily compute the cost of a constraint and, hence, the cost of avoiding immoral policies. But, as I indicated earlier, morality is never based on tradeoffs and constraint and, hence, the cost of benefits.

But belief in the existence of the inner life, the psyche, introduces a deeper philosophical speculation which will end my story of the endless philosophical journey. It is the speculation that from the point of view of the inner life, progress and the

closely associated ideas of development and maturation have no relevance and cannot be sensibly used to describe the process of the inner life. To be sure, there has been a strong temptation on the part of depth psychologists to transfer the viable ideas about progress in the outer, social world to the life-process of the inner world, but the philosophical speculation argues that this all-too-human intellectual step was erroneous. Even the labels psychopathology, psychotherapy, and mental disease may be inappropriate from the point of view of the inner life, although they clearly have sensible meanings in the outer, social life.

Also, so runs the speculation, the inner life is ineffable and the usual means of communication, e.g., comparison, are not appropriate. That the inner life is essentially ineffable does not imply that we can't talk about it, but rather that we must speak in myths, poems, stories, allegories, in short, in a language far different from that so "effable" a language, mathematics. Thus there may be "progress" in the inner life, but its meaning is essentially ineffable.

Does this philosophical speculation help us better understand SD? It does at least one important thing: it indicates that inquiry into the inner life, through meditation, psychoanalysis, group therapy, religious devotion, etc., is also as much "systems design" as is a large global simulator.

But if we accept the speculation, then we must realize that we evaluators live in two worlds, the outer world of practice and the inner world of the self, and that both worlds are real. But they contradict each other, do they not? The reply may be that the two

worlds are dialectical, by which is meant that through contrast we can learn a great deal about each of them.

So it looks as though the outer world dominates after all because there is a purpose to the clash between the two worlds, and the purpose is to learn which is an aspect of progress. But not necessarily. For there are two opposite versions of learning (of course). In the outer world, learning in SD means a shared experience and reasoning to help us better understand the consequences of our actions in the future. Hence, outer learning stretches its significance into the future or into different problem settings. It is learning "from a distance" and, hence, is what people sometimes call Apollonian, since Apollo always shot his arrows at a distance. But learning in the inner world means a direct and unique occasion, not for the sake of anything, but for itself. Such learning touches the essence of the present and is Dionysian, since Dionysius "touches."

## 12. Pause in the Story

At this point the philosophical story can stop itself by its own volition. The nice part about philosophical inquiry is that you can always turn it off simply by introducing a common-sense axiom which you have no intention of examining further. For example, the last section of this article can be successfully ignored by simply denying the existence and/or relevance for SD of the inner life of the type postulated by the philosophical speculation. If you believe that the psyche is a brain or is like a large information-processing computer, then there is no "inner" life other than a set of activities

that in principle can be examined and explained. You simply stop philosophizing and start doing other things: designing systems to do such-and-such in the outer world. You'll know you've left the philosophical journey once you hear yourself say: "Well, obviously," or "it's at least clear that . . . ." To a neurologist friend of mine, it's obvious that all psychology in principle is based on neurology; to me it's obvious that his conviction arises out of his own unique psyche.

As for the philosopher, he can end with a joke, to show that he's not all that deadly serious: isn't it marvelous that in order to defend its right to be so precise, SD must be so frustratingly vague?

11. References: a personal note of appreciation.

Over the many years I've acted as an editor of several academic journals, I've become more and more annoyed by our ritualistic style of introducing quotations and proliferating footnotes and references. The reader must constantly shift from the author's style and theme to the name of a book or some other author's style and theme. Too often, the quotation is out of context, or repetitive, or embarrassingly obvious (my ego cringes when I read: "as Churchman says, it is important to consider the larger system").

Nevertheless, I do appreciate one function of references: the acknowledgement that so many others were talking to me as I wrote this piece. E.A. Singer's Experience and Reflection (University of Pennsylvania Press, 1958) is as always a starting point, counterbalanced to some extent by Carl Jung's writings. Both Singer and Jung acknowledge their debt to the historical writers in philosophy and science.

I've enjoyed immensely the contrasts in SD approaches among the practitioners in the outside world. It's been immensely enlightening to compare Russell Achoff, Stafford Beer, J. Forrester and George Dantzig. Of course my personal bias goes toward Achoff's very realistic view of the managerial world.

As for the inner world, the contrasts are even greater. Besides Jung, other Jungians (e.g., Maria Von France and James Hillman) have inspired my remarks about the "no-progress" nature of inner life. Eric Jantsch has been on a trip similar to the one of this paper, of extending SD into other lands of human imagination. Above all, Thomas Cowan, in many articles and countless more conversations has so nicely kept me from drawing conclusions.