SDC

SP-877

# On Inquiring Systems

C. West Churchman

13 July 1962

# Inquiring Systems

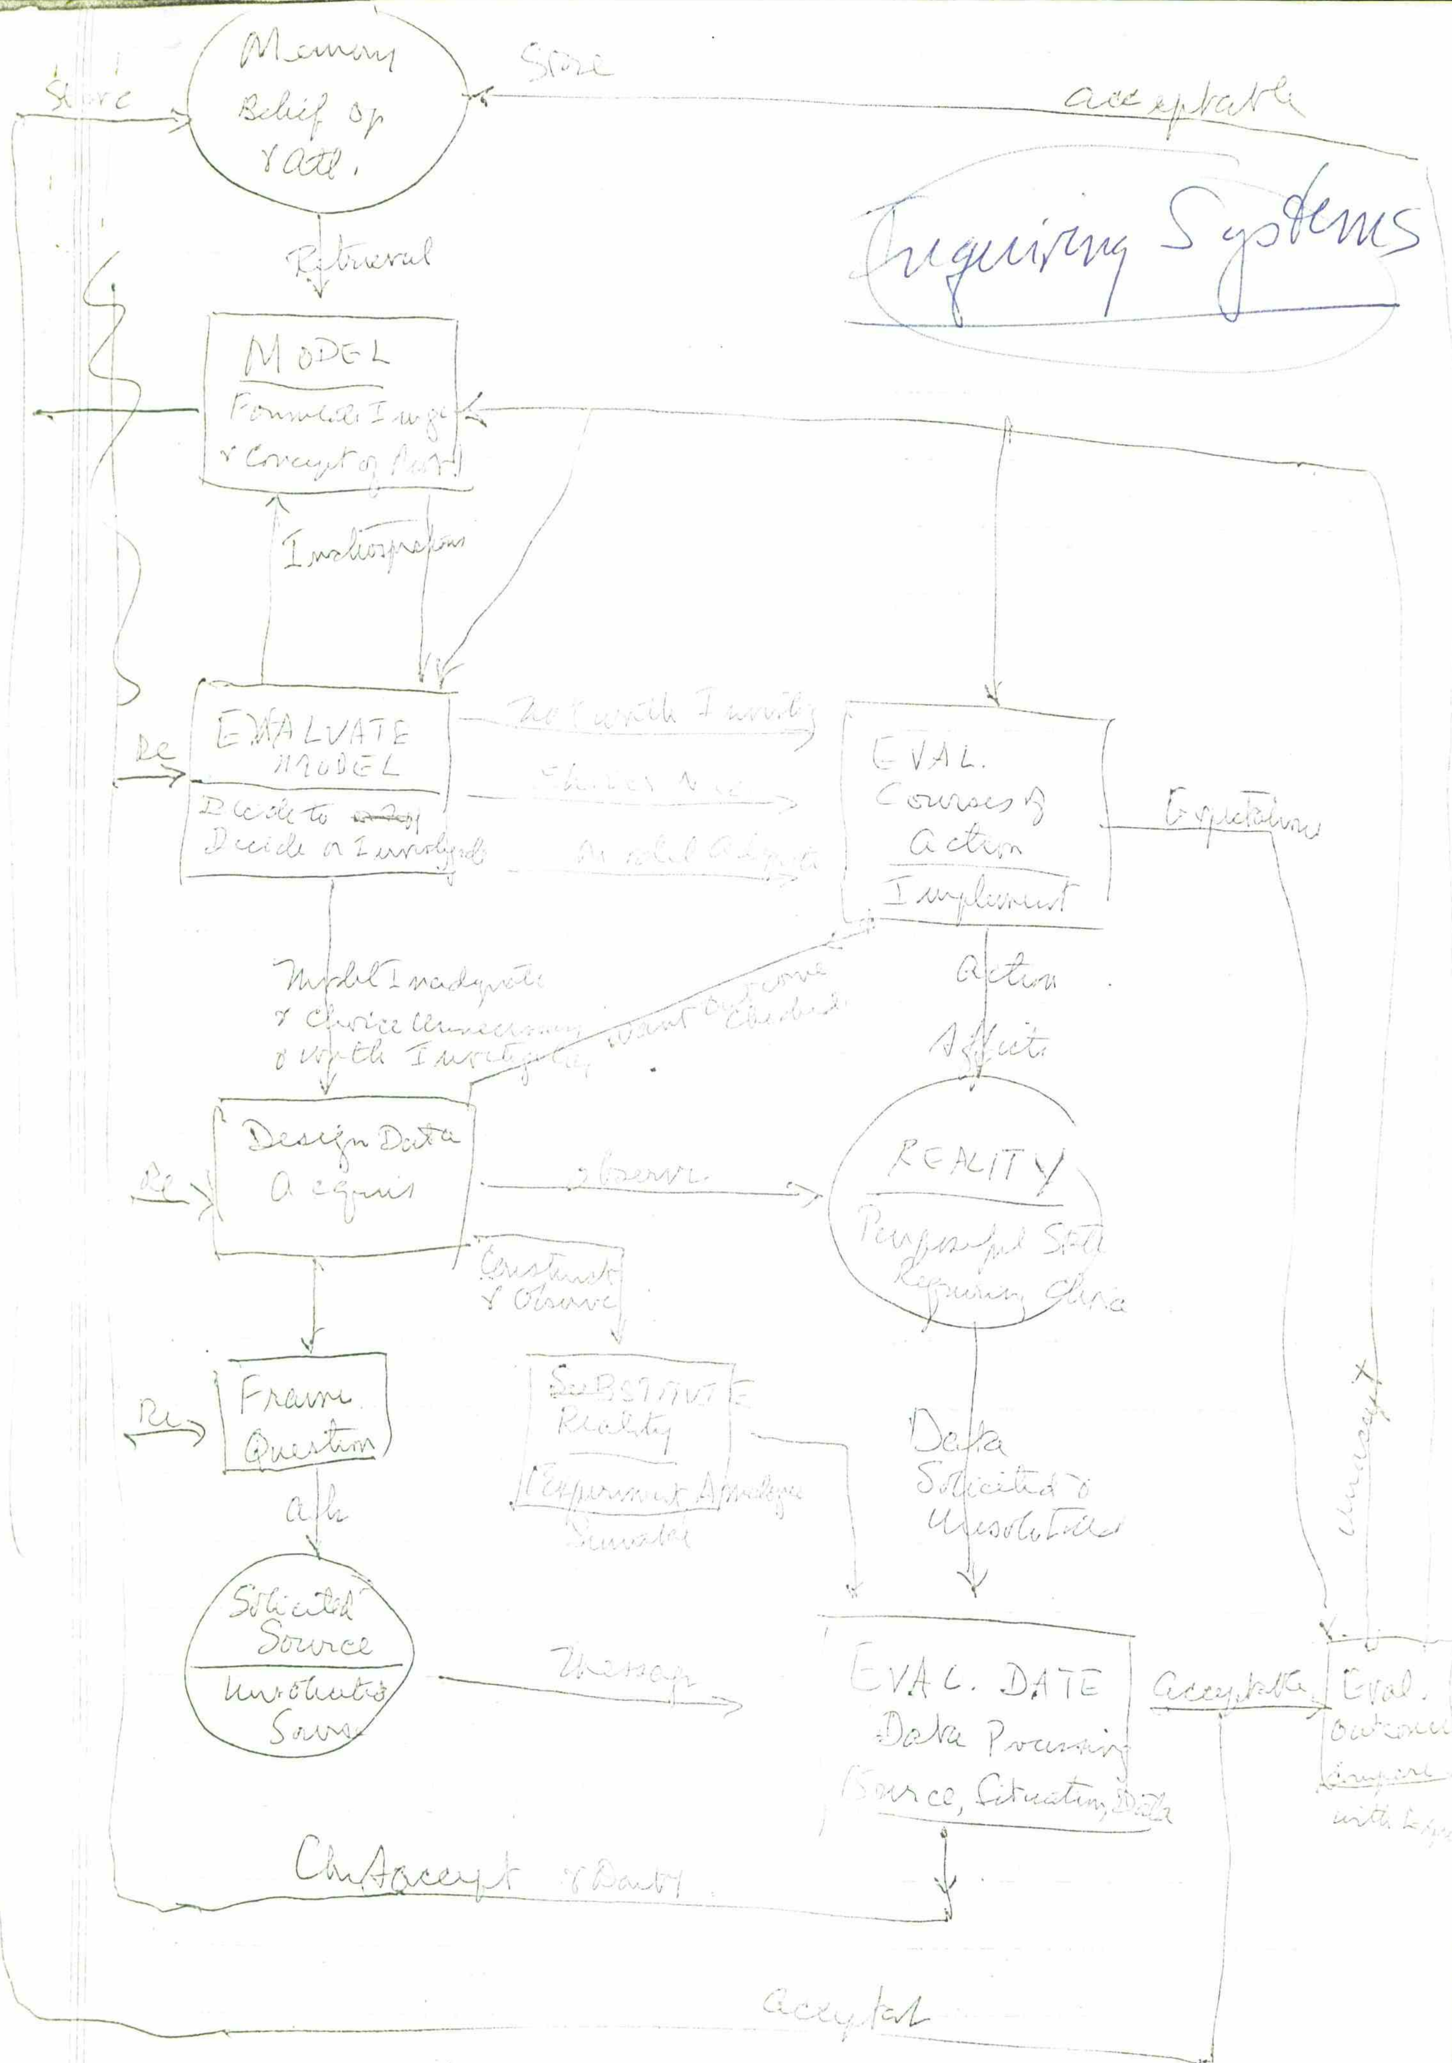


ACKOFF'S INQ. SYSTEM

Bert Raphael's Schema for Question-Answering Schemes

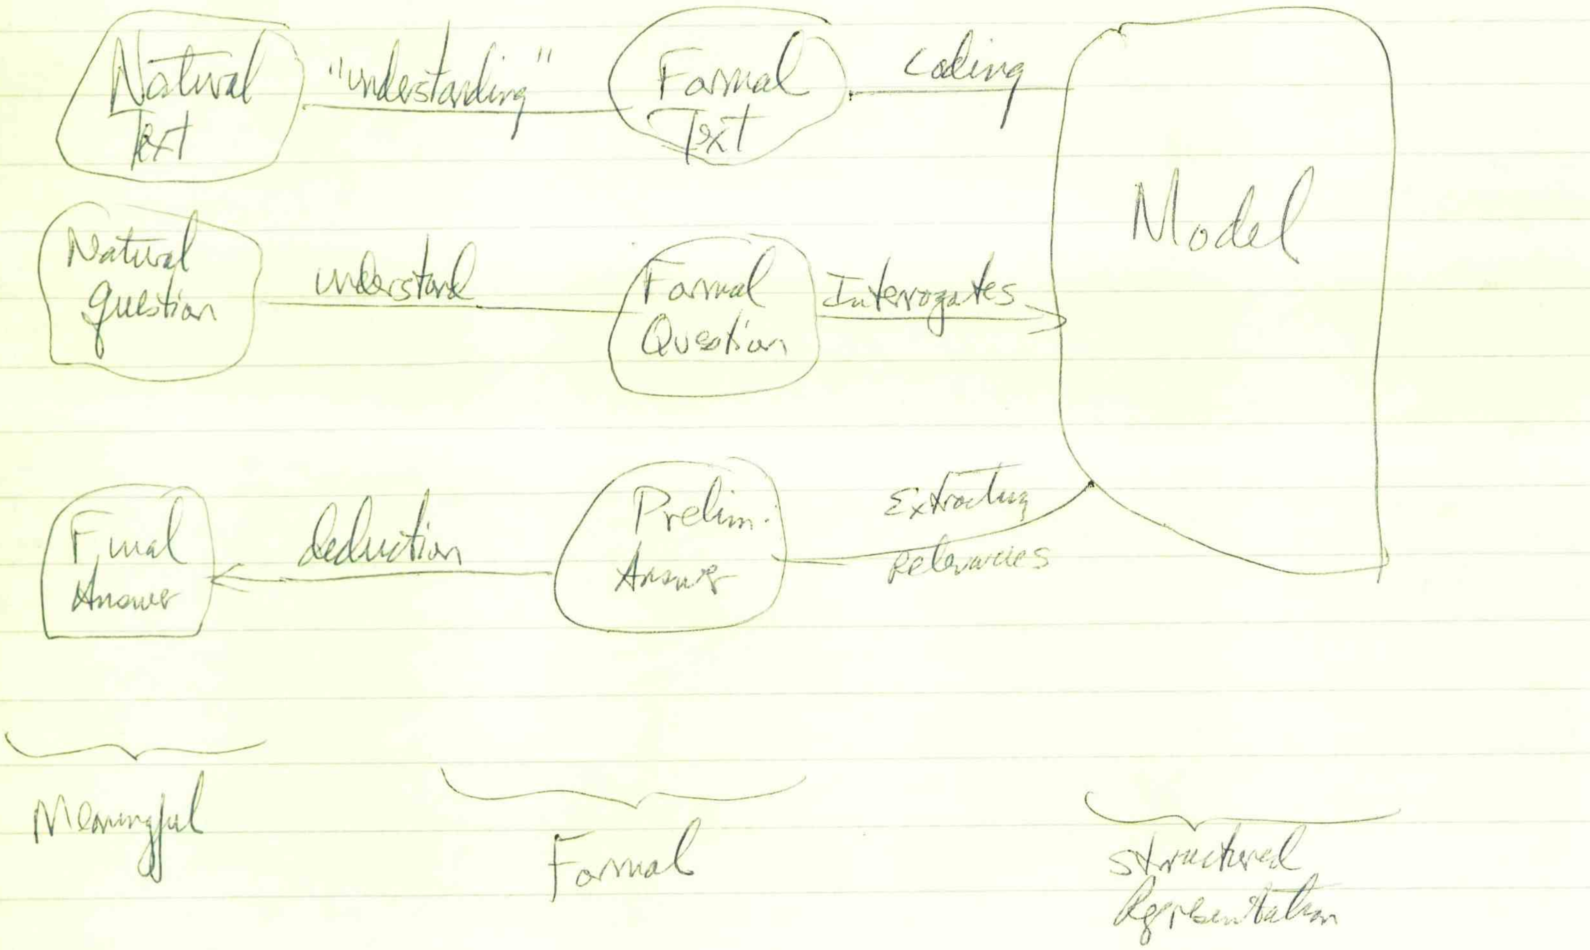

(SP Series)

SP-877

On Inquiring Systems

C. West Churchman

13 July 1962

SYSTEM DEVELOPMENT CORPORATION, SANTA MONICA, CALIFORNIA

A-1162

## FOREWORD

The first draft of this essay appeared as Working Paper #42 of the Center for Research in Management Science, University of California, Berkeley. The present draft incorporates some suggestions made by the members of the Center, and is specifically directed to those whose interests lie in system development. The aim is to ask whether it is possible to look at research as a system, and if so, what language is appropriate and what problems can be expressed in that language. I am aware that the language chosen in this paper is not the one that systems scientists most often use. It is a language more familiar to some philosophers. In this respect, the essay runs the risk of being misunderstood. But one must always express his muddles in the language that is easiest to use, with the hope that out of drafts will come a final statement in the most appropriate terms.

The general theme of the paper is this. Suppose one were to accept the following principle of system design: no system components can be properly designed until the designer has estimated the desirable properties of the whole system. If one accepts this principle, then what happens to the design of research, i.e., to the design of inquiring systems? Obviously a great deal happens, and this essay is only one tentative exploration of the consequences.

Two precautions should be given to the reader, in the light of comments on the first draft. First, if we propose to examine the design of research, we do not imply that individual freedom of the researcher is undesirable. Clearly, freedom of individual decision making may very well turn out to be a desirable property of inquiring systems. All we do imply is that the question of freedom or no freedom is a legitimate question to study. Second, this essay does not delve deeply into the value system of the designer, although this is an issue of some real concern. I believe that it is possible for inquiring systems to establish the proper value orientation of their designers, but the manner in which this can feasibly be done is not the topic of this essay[1].

C. West Churchman
Director of Research

---

[1] See E. A. Singer, Experience and Reflection (1960); and C. W. Churchman, Prediction and Optimal Decision (1961).

ON INQUIRING SYSTEMS

1. Man lives in a world of many systems. Most of these he has inherited from his predecessors, who themselves simply added to what was given them. Our systems have been designed in large part by the accidents of historical necessity. A reflective mind, therefore, finds much to think about when it poses the question whether reason could improve the design of a system if all the historical "givens" were wiped away. For example, how would we design our communication system if we could start from scratch with all our present knowledge? How would we design our transportation system, our marketing system, our system of international control?

Is there much point to this question that reflection can so easily ask? We know we can't throw away all equipment, and therefore a more practical approach would be to decide how to modify what we have, rather than design an abstraction. Yet, further thought shows that the question is not naive and useless, for if we succeed in designing an ideal system, we shall inevitably come to understand what kind of a system we're talking about, and we shall arrive at criteria of effectiveness that will help us to modify existing aspects of the system. More important, we shall be able to ask some very good questions, because our ideal will stand as a challenge to reality: why should matters be run as they are when there are far better ways to run them?

As we reflect on these questions we come to realize that the participants in the activity of systems rarely consider these issues. More likely than not, the evidence for a clear understanding of a system will fail to include the judgments of the humans who play very active roles within the system.

2. Science is a system. It has come down to us from our predecessors, who themselves modified what they inherited. Historical necessity has often been responsible for the form of the system of science. Have we any reason to suspect that historical accidents have bequeathed us a system that is altogether satisfactory? Have we any reason to suspect that our science is a very poorly designed system?

In other words, if we could start from scratch to design a system especially suited to create knowledge, i.e., an inquiring system, how would we proceed?

The design of systems is a series of problems; in each phase one must display the alternatives and try to select the best from the resources available. The really difficult problem is to discern a reasonable set of alternatives, especially when men become so used to one possibility. The purpose of this paper is to pose some questions about the design of inquiring systems and not to arrive at a design itself or even general criteria for design.

But we can't permit ourselves to ask these questions until we have more assurance that they mean something. There are many different viewpoints about what "inquiry" means and about the meaning of knowledge. The term "system," though very popular these days, is just as troublesome. Each discoverer of "systems science" has his own definition of a system.

The framework we shall require in order to state the questions rests on a specific meaning of "systems," to which we turn first.

3. We say that systems are examples of teleological things, i.e., things some of whose properties are functional. This postulate excludes the "solar system" if we look at the planet and stars as modern science has taught us to do. It may even exclude "formal systems" like geometry in many cases. If these exclusions seem improper, and if "system" is to be reserved for more general usage, then our postulate requests that we concentrate on a subclass for present purposes.

What follows is somewhat tedious, so that a homely example may help to follow the logic. Consider an ordinary electric stove, with four knobs controlling the heat of the burners. The "world" to be observed is initially described in terms of the position of the knobs, the temperature of the burners and the state of the water in the pots (cool, warm, simmering, boiling). Any two or more knobs in the same position belong to the same morphological class. Similarly, any two burners with temperatures in the same range (e.g., within

5 degrees F of each other) belong to the same <u>morphological</u> class, as do any two pots of water in the same state. Now suppose at time $t_o$, one of the knobs is turned but all else remains fixed. We say that the knob has changed its morphology. At a later moment of time $t_1$, one of the burners changes its temperature. We say the knob's position at $t_o$ <u>produced</u> the burner's temperature at $t_1$, meaning that had this knob been in any other position, some other temperature of this burner would have occurred. Also, at some still later time $t_2$, the knob's position at $t_o$ and the temperature of the burner at $t_1$ produce boiling water in a pot. We also note that the boiling water could have been produced by other positions of the knob at $t_o$. In this case, we call the set of boiling-producing positions a <u>functional</u> class: the members of the class have a different morphology but a common product. Finally, introduce into the scene a cook, who can produce any turn of any knob. This means that he can produce functional entities. We call such a cook a purposive individual; the final end (the boiling water) is his purpose; and the set of things he can produce is a teleological class.

4. We now try to explicate the meaning of teleological classes more precisely, by turning to a more formal and technical exposition of the ideas contained in this example. Singer (1960) provided us with a deep insight into the meaning of teleology. Functional classes, he says, are made up of entities that are alike with respect to their production of a certain end result. More precisely, functional classes can only be defined in the framework of a cause-effect model in which aspects of any time slice can be individuated and identified. An entity A, in a specific region of any time slice $S_o$, is a producer of an entity B in another time slice $S_1$, if it satisfies three conditions: (1) $S_o$ precedes $S_1$; (2) A must occur in its region in $S_o$ if B occurs in its region in $S_1$; and (3) A and B are proper subsets of $S_o$ and $S_1$ respectively.

It is clear in many cause-effect models that entities of quite different kind can produce the same kind of an end product. For example, many computers, once they are running, are virtually cause-effect systems. Some items in the memory may be essential for some later result, so that Singer's conditions are all satisfied (e.g., if the item does not occur in its special region at time $t_o$, the desired result will fail to occur at $t_1$). The item is therefore a producer of the result. But so are many other items in the computer's memory. Now if one concentrates on the producer-product relationship and forgets about the differences in structure of the items, then one thinks of the items as functional classes. That is, all items in $S_o$ that are essential for an output in $S_1$ belong to the same functional class.

Singer shows how one may effectively relax the conditions just specified for the definition of functional classes and thereby speak of entities having a common potential product. This is done by confirming within the model that some items of a given structure have produced a specific product in a given type of environment, and some have not. In this case, all items of the class are said to be potential producers. Further, one can introduce a metric in the classes and describe the probability of production. Also, one may want to define functional classes (e.g., a machine's output) in terms of things that are common products of one producer.

5. The extension of the definition of function to teleology is fairly straightforward. Suppose that some of the elements of a functional class that could occur in an environment are the output of one individual. If so, we call the potential outputs "means" and the common potential product the "end." The elements of the functional class are then alternatives relative to an end product. In this case we can call the class a teleological class, and the end product the purpose. Since all members of a teleological class are potential products of one individual, they are functional in two ways: as common potential products of the same producer and as common potential producers of the same product. Next, we note that we need not restrict our attention to one potential end product. If we consider several, then the teleological class is defined in terms of several ends. But we also want to maintain the same sort of a metric that was developed for a single end. In other words, we want all teleological classes to be ordered in terms of a

"more effective" relationship. Perhaps the simplest way to accomplish this is to weight the end products and develop a metric from the product of the weights and probabilities.

The important point in weighting the objectives, however, is that the weights be functions of some property of the individual who can produce the alternative means. Specifically, we shall want to say that the weights correspond to the individual's "intentions" or "utilities" or "values." In this case, we can speak of the individual as a purposive entity.

In sum, a teleological class is a set whose members (1) have a common producer; (2) have a common potential class of products; and (3) can be ordered by a relation "is more effective than," which depends on the properties of the common producer as well as the probabilities of production.

Thus, an entity is teleological by virtue of the fact that other entities could (or do) exist that might produce the same results and that might be produced by the same designer.

Systems, in the sense of this discussion, are teleological entities. This means that if something is to be called a system, there must be alternative systems, and there must be a designer whose intentions are expressed in terms of the common potential produces of the set of systems.

6. But not all teleological entities are systems. The differentiating feature of systems is that they can be separated into parts, and that the parts work together for the sake of the whole. Hence we postulate that a system consists of at least two teleological parts, and that the effectiveness measures of all parts, if maximized, yields a maximum effectiveness measure of the whole system.

More precisely, a system can be divided into parts in such a way that if the state of all the parts but one is held fixed, then the effectiveness of the system will increase as the part is made more effective. It should be emphasized that the postulated relationship between part effectiveness and system effectiveness is a relatively weak one. It merely asserts that, in system design, one can always improve the system by improving one of its parts. But the pathway to the optimal system may not consist of improving each part, one by one, until it is perfect, because even though one part is perfected it may lose its top position when another part is changed. More generally, the rank order for one part may completely change if another part is changed: the rank orders may be functions of the state of the parts. This consideration is central in what follows, because we shall be interested in two kinds of systems: those in which the rank orders remain fixed independent of the other parts, and those in which the rank orders depend on the state of the other parts.

We now define a "system" in more formal terms.

Consider a set of entities, $S_1, S_2, \ldots$, all of which are members of the same teleological class S.

Let $M_1, M_2, \ldots$, be the effective measures associated with $S_1, S_2, \ldots$, respectively. Let M* be the $\max\{M_i\}$, and S* the corresponding member of S. Let the relation P stand for "is a part of" (P is binary, asymmetric, transitive, and its domain is the set of all entities).

S is a class of systems and each $S_i$ in S is a system if and only if for every $S_i$ there exists a set of at least two entities $\{s_i^j\}$, such that (a) each $s_i^j$ is a member of some teleological class $\bar{s}_i^j$; (b) $s_i^j$ is a part of $S_i$ for every j; (c) there exists a variable y which takes as its values the measures of effectiveness of the members of S, and a variable $x_j$ which takes as its values the measure of effectiveness of $\bar{s}_i^j$, such that y is a monotonic nondecreasing function of each $x_j$, and for some values of $x_j$, y is a monotonic increasing function of $x_j$. Hence: (d) a system S* has a maximum effectiveness M* in S if and only if every part has maximum effectiveness.

In other words, every alternative system $S_i$ is made up of teleological parts [conditions (a) and (b)]; one system is more effective than another if its parts are more effective [condition (c)]. Hence, the most effective system has the most effective parts [condition (d)].

The designer of systems seeks to find measures of effectiveness of the whole system and of its components so that he can use these measures to guide him in the design. He would like the component effectiveness measures to be related to the system effectiveness measures at least in the sense described above: namely, that an increase in component effectiveness implies an increase in system effectiveness.

The crucial problem in system design is to determine what the parts are and to create the appropriate measures of effectiveness. In other words, we are not considering a design problem in which the parts are given. Rather, the designer must decide what a part is, and how it should be used, i.e., how its effectiveness should be gaged.

Now it may seem reasonable to hope that in some cases a part of a system can be designed without considering the design of other parts. Whether or not this is possible is a matter of fundamental concern to the designer, and will be the principal concern of this discussion of the design of inquiring systems.

We say that a class of systems S is <u>separable</u> with respect to some part if the measure of effectiveness of the part is independent of the states of the other parts.

More precisely, a part of a system is separated in a weak sense if the optimal state of the part does not depend on the states of the other parts; a part is separated in a strong sense if the ordering of all its possible states remains invariant with respect to the states of all the other parts.

We say that a class of systems is <u>weakly</u> separable if it can be so described that at least one of the parts is weakly separated; it is <u>strongly</u> separable if it can be so described that at least one of its parts is strongly separated.

A system is <u>completely</u> separable (in either a weak or strong sense) if all parts can be (weakly or strongly) separated.

Production systems afford excellent illustrations of nonseparability, since they are among the most carefully studied systems we know about. Consider a manufacturing system. One typical way of partitioning the system is to break it into these subsystems: procurement, order processing, production scheduling, production control, labor force, inspection, packaging and distribution. In all production systems with which I am familiar, none of these parts is separated in even a weak sense. The optimal procurement policy always depends on the way in which items are scheduled; the optimal production scheduling depends on how labor is deployed; all parts depend on the extent and timing of control; and so forth. Indeed, most students of production question whether manufacturing is ever a separated part of the whole organization (e.g., the optimal manufacturing policies depend on the state of the subsystem that controls investments).

Specifically, one may try to redesign a production scheduling system by more or less ignoring the other aspects of the firm. As a result one may move to a system that works perfectly as long as the rest of the policies remain fixed. But when the deployment of labor changes, or the marketing system is modified, the production system may become worse than before it was redesigned.

On the other hand, two men assigned to dig a trench might be considered as a separable system: what work each accomplishes simply adds to the total. Some systems that collect and store items may also be construed as separable. Also, a system designed to solve problems in a formal framework may be separable: the optimal method of solution of one problem may not depend on how the other problems are solved.

It is not difficult to describe the concept of separability in mathematical terms. Suppose there is a teleological class of systems such that for every $x,y \geq 0$ there exists a corresponding system with effectiveness $T(x,y)$. Suppose $T(x,y)$ is a continuous function in $x,y$ with partial derivatives $T_x$ and $T_y$. Finally, suppose that $T_x$ is a function of $x$ only and $T_y(x,y)$ is a function of $y$ only (i.e., $T(x,y)$ can be expressed as $A(x) + B(y)$, where $A(x)$ is a function of $x$ only, and $B(y)$ is a function of $y$ only). Then the class of systems is separable into two parts: the one controlling the variable $x$, the other the variable $y$. In this case, both parts are strongly separated.

But our interest in this paper is not in the formal properties of separability, but in the evidence for separability. In other words, if it is clear that some function of a set of variables can be used to rank the effectiveness of a class of systems, then it may be relatively easy to determine whether the systems are separable. The problem is to prove that a specific mathematical function is appropriate.

In this regard, the lesson to be learned from the study of production systems is an important one relative to the empirical determination of the effectiveness of systems, and therefore should be explained in more detail. When we examine a procurement policy, we normally attempt to measure some of the relevant properties of the system: the way in which demands come into the system, the way in which prices vary, the costs in placing orders, the delays in receiving orders, and the costs of storing items. These measures are then combined in a model to form a measure of the procurement effectiveness. In attempting to measure these properties, we must examine how other parts of the whole system are operating. The demand, for example, is the consequence of policies of the production department or sales department. The cost of placing orders is the consequence of policies of the order department. The cost of holding stored items is in part the consequence of policies of investment of the firm. Now these policies of other parts of the system can be varied; in fact they determine the effectiveness of the part with which they are associated. As a consequence, any so-called optimal policy of the procurement system is asserted to be optimal only because one assumes that the other relevant policies of the other parts are optimal. In other words, the procurement system's effectiveness is measured in terms of the effectiveness of the other parts. In this case, we infer that the procurement activity is a nonseparable part.

This example can be generalized into the following principle of nonseparability: a part is nonseparable if its effectiveness is measured in terms of the teleological properties of other parts.

There is a strong tendency for the system designer to create separable parts. We can therefore summarize this discussion of the nature of systems by noting two fundamental problems of the designer: to identify the parts (i.e., to define teleological entities such that the system effectiveness is a monotonic function of the part effectiveness), and where possible to design the part so that it is separable.

We should note that a system may be designed as though one or more parts were separable, and operated in this manner until the separation is no longer feasible, at which point the system is redesigned. This, indeed, is the natural way to regard the separation of parts, by means of "management by exception." Later in the paper, we shall return to this type of design and try to make its meaning more explicit.

7. Supposing that "system" has been well enough defined, we turn to the meaning of "inquiring" in the phrase "inquiring systems." In general terms, inquiring systems have as their purpose the generation of knowledge. But these general terms are not very helpful in discussing the systematic properties of inquiry. We can approach a little nearer to the precise definition of the goal of inquiry if we say that it is always to increase man's understanding of his world. When a man understands, he knows why an event occurred. This suggests that the understanding man is never surprised. This in turn suggests that the event that is understood is in some sense redundant for the understanding person. And finally this suggests that the aim of inquiring systems is to maximize the redundancy of events. These suggestions are all tantalizing, and require further exploration. For the moment, though, we shall leave the definition of inquiry in this suggestive form, in order to discuss the nature of inquiring systems. Later, we shall return to the definition itself.

8. The question of this paper is: to what extent can inquiring systems be designed with separable parts? We begin by considering some traditional answers to the question. Two schools of thought that characterize Western science are (1) the rationalist: that it is possible to think correctly, and (2) the empiricist: that it is possible to observe objectively. But these traditions are also phrased in much stronger terms by some scientists, namely, that thinking and/or observation are separable parts of any system of inquiry.

This stronger position, it seems to me, is the essence of positivism. The crucial point is not that thinking or observation may fall into error; everyone recognizes this possibility. The positivist position is that it is possible to do the best one can with some acts of thinking or observing without having to be concerned about the uses to which one's thoughts and observations are put or the way in which they are communicated to others. Again, operationism is a philosophy that seems to argue for separability of inquiry; it believes that a concept can be defined with maximum effectiveness by specifying a set of operations independent of the use to which the results may be put or the methods by which the results are communicated.

If one believes in the separability of observation, then one can argue that although the nonscientific politician may say to the scientist: "It is too bad you made this observation today, because as a result you have disturbed the political scene and I must try to prevent the recurrence of your act," the scientist can sensibly reply that the observation was still "correct." It was correct because, on the basis of the separability of observation, the act of observing can be perfected in itself, and its value does not depend on the manner in which the observational report is received.

Now, it seems clear that inquiring systems can be designed in such a way that the choice of one part does not depend on the choice of the others. The question is whether an inquiring system so designed is ever optimal. If one insists on intuitive grounds alone that observing and the use of observations be separated, then one has pre-empted the response to this question.

9. In order to discuss this problem, we can present one partitioning of a system of inquiry, without thereby claiming its superiority to other schemes. This one has the advantage that it seems to follow a logical sequence of steps, but this advantage may be illusory, of course.

Part 1 provides the resources of manpower and facilities.

Part 2 determines what specific objective an inquiring system shall pursue: i.e., defines the problem area.

Part 3 specifies the problem, i.e., creates a model within which the problem can be defined (the model often being expressed in a formal system).

Part 4 determines the logical consequences (theorems) of the model.

Part 5 specifies what data are required, in what form, to what degree of precision, in what amount.

Part 6 specifies how the data are to be collected.

Part 7 collects the data according to the requirements.

Part 8 transmits the data to a central point.

Part 9 analyzes the data.

Part 10 produces a set of results.

Part 11 stores the results and transmits them when needed.

Part 12 determines when stored results are needed and how they are to be used.

The rationalist thesis is that inquiring systems can be so designed that Part 4 is separable, and the empiricist thesis is that they can be so designed that Part 7 is separable.

10. Before examining these theses in detail, we can satisfy ourselves rather quickly that some of the other parts are not separable, and indeed that their proper design is still an unsolved problem.

Consider, for example, Part 11, the part of inquiry that stores and transmits information. In earlier times, men thought of communication systems in relatively simple terms, perhaps because the total amount of real information was small. The alternative designs consisted of mixtures of talking and writing and personal memory. Today, we all realize that the problem of storing and retrieving information has become serious and that the number of alternative designs is very large.

If we were to start from the beginning, would we build libraries with books? Would we publish journals? Would we hold meetings with papers? Or would each scientist's study be equipped with pushbutton panels that would call up what he needed to know on a television screen? What form would such information take? And how would a researcher in such an environment know when to request information?

At our present stage of technological development, we have some valid reasons for suspecting that science is poorly designed in its communications, so poorly that the resulting inefficiency may be colossal.

In any event, it seems clear that information retrieval systems are not separable. For example, the optimal information system depends on how the results are to be used (Part 12), as well as on the form of the results (Part 11) and their structure (Parts 3 and 9). This assertion seems clear if we conceive of inquiring systems as very large. For the time being, we shall want to reserve judgment on whether certain types of inquiry can be separated from the rest of inquiry, i.e., whether the most general inquiring system can be separated into inquiring parts.

11. Consider also Part 1 which could be called the part concerned with the political economy of science. Science has a politics of its own, and its political activities often come in conflict with the politics of business and of government. We sometimes tend to think that the politics of science is good, because, I suppose, everyone wants more knowledge and therefore the political activities of those who acquire knowledge for man must be good. But a little more honesty leads us to suspect that the manner in which some men become leaders of science may have been quite costly in terms of ruined careers and dollars, and that the kind of research that receives large grants may not be the best research for the development of science. Here again there is a prevalent and odd viewpoint that wants to distinguish between science and its politics; it sometimes takes the incredible form of arguing that science is a body of knowledge, and politics is people, and therefore the two must be separated. The point is that there is a design problem here. The problem is whether one can optimize the system of acquiring knowledge without considering the political problems that such a system generates. That is, can the optimal political strategy of science be determined independently of the rest of inquiry? For example, can we maximize our effectiveness in educating potential scientists without regard for the structure of inquiry itself?

More seriously, we have to ask ourselves what the proper size of Part I should be, in order to be able to gage its real effectiveness in the system of inquiry. If science competes with other systems in enlistment, then shouldn't the whole higher education system be regarded as a part of science?

12. All these questions can be made more difficult and their importance clearer if we realize that science is an institution devoted to survival. Its programs are designed for long-range objectives that cannot by their nature be accomplished in a generation. Hence, one strategic problem of a scientific project is the optimal bequeathing of the project to others, just as one overall strategy of science is bequeathing its work to the next generations. These strategic problems exist unless we accept the viewpoint that some inquiring systems can be designed with separable parts, each part being an inquiring system. This position would amount to asserting that "facts," once obtained, constitute knowledge even though no one besides their discoverer ever sees them. This, in turn assumes that some "facts" cannot be improved, a point which we shall want to examine with some care.

If we admit the bequeathing problem to be an essential part of the problem of scientific method, then we must include within science the strategies for the survival of science. The consequence is that the system of science will become very large in eras when politicians and industrial managers seek to use science for ends that conflict with the aim of scientific survival. In other words, the system of science must include estimates of optimal international policy, if international policy could threaten science's existence.

13. At this point we can afford to pause and ask whether the conclusion just reached makes a difference. Of course there are very many scientists who do not believe that world politics is a scientific issue. Their disbelief is based on their implicit assumption that "science" is simply the output of what I have called the system, i.e., the body of knowledge that the system creates. But most of these scientists would include the logic of scientific inference in the corpus of science. It follows that they are willing to include in the "body of knowledge" some assertions about how the system ought

to work. If this much is included, then why not the rest? On what basis are we to draw the borderline between the internal strategies of science and its external strategies?

More to the point, if science is taken to be an inquiring system, then for any part of science one must ask for the display of alternative designs. One way to design the politics of science is to ignore all considerations of the activities of international politicians. Another is to take some of these activities into account, and design a system of inquiry in which international policy making is a part. The question is: which alternative provides a better design?

One might reply that we don't know about international politics. This may be true enough, but our inability to solve a problem does not typically exclude the problem from science.

Still, even if those who want to keep science pure are illogical, does the issue make a difference? Most scientists are concerned about world politics and would like to help. What difference does it make whether they regard the issues to be scientific or nonscientific?

There are at least two ways in which it does make a difference. The first has to do with attitudes. If communication and world politics are regarded to be scientific issues, then the status of those who work on these issues is improved in the scientific community. Much more important, scientists themselves will realize the necessity of rigorous analysis and controlled fact finding in the study of these problems, qualities that seem to be lacking in many of the current discussions. Finally, if these issues are taken to be scientific, then our best minds may want to weigh the scientific value of working on them rather than on matters of energy and space.

The second difference is reflective. I said that most scientists would place the logic of inquiry in the corpus of science. But I think that very few of them have ever tried to check this logic in the same sense in which they check discoveries in their own fields. The validation of the logic of inquiry consists in showing that certain activities carried on within the scientist's laboratory or study are optimal ways of reaching results. In the remainder of this paper, I want to cast doubt on the operational effectiveness of some of the published accounts of the logic of inquiry; the doubt will be based on speculation about the system of science. Specifically, I want to examine the separability of Part 7.

14. Part 7 is the activity of collecting data. The history of reflection about modern science has centered about the "givens" of the scientific enterprise. But we know that these so-called givens are in fact "taken" in rather specific and elaborate ways. Our interest here is in the design of observation, i.e., in the alternative methods of observing the world.

At the outset we need to define "observation." The concepts of production and purpose which have been defined earlier will be useful in this regard. But before entering into the technical details, we ought to make a distinction between sensing and observing. The distinction is based on the notion that sensation is simply a reaction of a purposive individual to something that is occurring in his environment, whereas observation entails other properties as well. Specifically, observation entails the recording of the sensation, which seems to mean the transformation of the sensation into some entity which is relatively stable over time. Thus, we say that "X senses Y in environment N" is true if (a) X is a teleological entity in N, (b) Y is an entity in N, (c) Y produces a change in the behavior of X. Normally, we should also want to specify that X's reaction is purposeful, i.e., can be construed as a member of a functional class.

It seems quite clear that X can sense Y without "knowing" that he senses it, although an observer could know X's sensation. Awareness is not a necessary condition for sensation.

We can readily see that sensation is a fairly weak concept for our purposes. Most of us are sensing the world around us in many ways at most moments of time. Observation is the stronger concept we require because it involves the additional idea of a record. Although a record must be related to the sensation in some way, it need not be a perfect copy of it. A laboratory technician reacts to litmus paper and writes "red" or "acid." The record is not red at all, nor is it like his sensation in any obvious sense. Nor can we say that the record is merely something produced by the sensation. The sensations of the technician produce many things in his nervous system that we would not want to regard as observations.

The concept of a record seems to require other entities and properties for its meaning. The clue seems to reside in the idea that if a record is kept, then someone may go and examine it, and from his examination infer that such-and-such an observation was made in such-and-such an environment. Hence "X observes Y in N" is true if

(a) X senses Y in N;

(b) X's sensation produces in N an entity Z which exists in some later environment N';

(c) Z in N' potentially produces in some W a belief that a sensation of some type has occurred in N.

Even this definition requires further consideration to make it satisfactorily precise, but for our purposes the idea is clear enough and can be summarized less technically as follows: an observation entails both a sensation, and the transformation of the sensation into an entity which survives over time, and which may be used to infer the occurrence of the sensation. We may note in passing that observation, so defined, is not restricted to living beings.

Finally, we need to discuss the effectiveness of sensation and observation. Usually, we employ the term "accuracy" in this regard. A sensation is inaccurate if it does not "correspond" with reality. But the test of this correspondence lies in the purposes of the tester. In other words, X's sensation is accurate if his response serves some purpose very well, and is inaccurate if it does not. In the same manner, an observation is inaccurate if the belief it involves about an earlier sensation is not an effective belief; "belief" is itself a teleological action, and hence its effectiveness can be ascertained.

Thus, the two stages of observation (sensation and recording) may have varying degrees of effectiveness. The problem of the design of observation is to design the two stages so as to maximize their effectiveness. In this regard it is worthwhile mentioning that elusive concept called "fact." Facts are maximally effective observations. The concept is elusive because no one has a very clear idea about the proper design of observing systems. Often, the records themselves are taken as facts, without any regard for their origin in sensation or for their effectiveness relative to potential users.

15. We now examine some proposals for the design of observation by turning back some pages of history. We begin with John Locke, whose Essay Concerning Human Understanding can be read as a proposal for the design of an observer. The main idea here is to introduce the concept of "objectivity" into the observing process, and the central point about objectivity is that the observer does not "add" anything of himself to the sensation. For example, the observer does not add his own preconceptions or values. An observer who is convinced that something is true, or an observer who wants something to be true, may read his convictions or his wishes into his records. Both observers are "non-objective." The problem of designing an observing system is to capture this notion in a suitably general form. Now any sensing entity will put something of itself into its sensations; this follows from the meaning of sensation. What it must not put in are "unwarranted" responses. One might try to solve the problem by thinking of a whole class of observers of Y in N, and by saying that the "objective" component of their responses is the common class product of the responses. That is, what is objective about my sensation is what my sensation has in common with all other sensations of other potential sensers of Y in N. This seems to be the idea underlying the definition of objectivity in terms of "intersubjective agreement." But such a notion does require the determination of a suitable class of observers, which surely must have more than one member. Indeed, the number of members in the class seems irrelevant: it is the quality of the class that counts.

Locke's solution avoids the difficulties of defining the class of competent observers by introducing the very ingenious idea of the simplicity of a sensation.

Suppose that the inquiring system can establish the relationship "is simpler than" between any two sensations, and suppose in its application this relationship obeys the principles of simple ordering. Suppose further that there is a set of simplest sensations. Finally, suppose that any observation can be expressed as a combination of simple sensations. Then we could call the simplest sensations "objective," as well as any observation that can be reduced to a suitable combination of simplest sensations.[1]

Locke gives some hints as to how an inquiring system might accomplish these results. It must have built into it a power of reflection that enables it to answer certain questions about its responses. Thus it can always ask for the breakdown of a sensation into its parts, and it always receives reliable replies to its question. For example, if it sees a yellow square, it knows that the sensation is made up of the sensation of yellow and of square. It knows that the first part (yellow) cannot be broken down any further, and is a simplest sensation; it knows that the second part (square) can be broken down into a combination of a sensation of extension and of comparison.

The reflective powers of Locke's observer are extremely important, and without them the system would be quite inadequate to serve as a generator of observations. The responses the system receives to its queries about itself can also be ranked by the relation of "is simpler than." Also Locke's observer will be able to perceive, retain, discern, combine, etc. In general, the observer will not be able to observe <u>how</u> he does these things. Perhaps one of the most important things the observer is able to do is to form his observation into a sentence in some language that has a well defined syntax (including a logic of propositional functions).

---

[1] The relevant passages are Ch. 2, Bk II, <u>An Essay Concerning Human Understanding</u>. As to the objectivity of simplest sensations, see Chapter IV, Book IV.

One has no difficulty in discerning how very difficult it would be to design a Lockian observing system. The chief difficulty, as later philosophy pointed out, is the meagerness of basic structure that Locke permits. This meagerness is captured in Locke's own characterization of his system as a "wax tablet." In a way, this is an absurd analogy of a system that can rank inputs, and discern, retain, combine, create sentences, etc. Even so, the system at times is quite passive in its role of observation. It can ask itself whether it has had a specific sensation and receive a highly reliable response. It can assert that the simplest sensations it has actually had are "objective." But it cannot make any inferences beyond the capacity of its structure. Thus, although the system can anticipate, Locke fails to show how the anticipation can be made objective. Hence the "future" for such a system must be largely a mystery. Furthermore, the system is so weakly reflective that it cannot operate on its own powers; since it cannot understand how it discerns, it has no choice in discerning. Although it can willfully act to retain (record) a sensation, it has little ability to compare the effectiveness of different recording strategies.

Other difficulties of Locke's design are apparent: e.g., if the system can ask questions about its own actions can it also ask questions about its ability to query? But these issues need not concern us here, because it seems more important to consider a different kind of design in which the underlying structure is enriched, and to use Kant's first Critique for this purpose.

16. Kant's inquiring system can determine why it is able to discern. It has built into it a kinematics, which enables it to differentiate objects in terms of space and time. In other words, its observational sentences become part of a kinematical theory, and the observed objects therefore obey certain regular laws. Since the general kinematical structure resides in the system itself, it is "a priori." The observational sentences supply the details of this general structure. It should be noted that Kant argues that such a structure is essential or else the categories of "one" and "many," which are essential for any observational report, are meaningless. Hence it

is incorrect to say that in Kant's system "objects" are "subsumed" under a priori laws. Rather an object is created when, in the observational process, an observational report is generated from what is "given" and from the a priori structure. The great mystery of Kant's system is what is given. The system cannot sensibly ask for further detail on this question. Indeed, it cannot even pose the question, because the pure given (the input) cannot be identified or individuated. Of the Kantian system one can say that it observes and that the system's structure is a necessary but not a sufficient condition for the existence of an observational report.

There is one distinction between Kant and Locke that is especially important here. Locke's system can discern, but it cannot know how it does this. Kant's system has the potentiality of answering this question. This is why Kant felt that all objective science must be based on an a priori set of principles. A system such as Locke's, that cannot ask how it accomplishes what it does must learn inductively about its own structure. Thus it learns that objects obey certain geometrical, kinematical and mechanical laws. But it has no way of determining whether these regularities are a part of the scheme which feeds its inputs, or are a part of its own makeup. Hence Lockian systems are forever beset by the problem of regularity. What Kant was hinting was that a well designed system of inquiry can always ask for the nature of its own design. That is, if an inquiring system is well designed, it must be able to sweep in the designer in its series of inquiries. But Kant himself did not go this far, because he would not allow his system to query the origin of the inputs, or even the nature of the inputs. This is Hegel's criticism of Kant, and the exploration he conducts in his Phenomenologie des Geistes is a labored attempt to see what an inquiring system is like if it has the power of understanding its own design. Hegel's well known point is that the system's procedure of inquiry must consist of setting up a series of theses in such a form that the contradictory of each thesis is supported by the thesis itself.

17. Perhaps the most radical of all proposals for an observing system is that of E. A. Singer. We have seen that in Locke's system there is

parsimony with respect to the structure that the system itself imposes on the observations. This parsimony was carried to the extreme in Hume, and later on in this century in logical positivism: only enough structure is put in to enable the system to create observational reports as responses to stimuli. The task of such systems is to try to explain a lot on the basis of very little a priori structure. The process of inquiry of such systems consists, as in Locke, in building knowledge from elementary, direct observations. The direct observations are taken to be nondecomposable. In modern positivism, these basic building blocks are the maximally agreed upon reports of the system.

In Singer's system, inquiry begins by taking an observational question and decomposing it until disagreement is reached. The process can be illustrated quite simply. The system responds to a stimulus by the report "A is green." It then queries whether this same report occurs every time the stimulus is presented to a system like itself in this kind of an environment. If the answer is "yes," then the system must decompose the question. In Singer's formulation, the decomposition consists of finding a set of propositions such that (a) each proposition implies the original report, (b) each is a contrary of the others, (c) the original report implies that at least one of the propositions is true, and (d) the entire set can be ranked by some relation. Thus, let p be the original report. A decomposition of p is a set of propositions $p_1, p_2, p_3 \ldots p_n$ such that for every $i \leq n$, $p_i$ implies p, "$p_i$ and $p_j$" is logically false, "p implies $p_1$ or $p_2$ or $p_3$... or $p_n$" is logically true, and for which a relation R can be found satisfying the axioms of ordering over the $p_i$'s. In our example, "A is green" might be decomposed into "A is light green," "A is plain green," "A is dark green." The system now tries to make a new observational report with the same stimulus and environment, in which it is constrained to the decomposition. It again queries for agreement, and if agreement still obtains for one member of the decomposition set, it sets up a new decomposition. This process is continued until disagreement arises. The system proceeds with these disagreements as follows. It sets up a scheme for adjusting all observational reports to some standard set of conditions; that is, it accounts for some of the disagreements in terms of time-dependent laws. If the adjusted set of observations all agree, a new decomposition follows. If they disagree, the system applies an analysis of variance; that is, it accounts for the disagreements in terms of random variation. If the analysis of variance shows "no significant difference," the system will assume that any a priori structure it has used to frame its observational reports is satisfactory. It will continue to do this until either (a) the analysis of variance shows statistical differences too large to tolerate, or (b) some alternative structure which can operate equally well is proposed. If the latter situation develops, the system seeks for new decompositions and tests so as to resolve the question as to which a priori structure is correct. It may also neglect a proposed structure because of the structure's computational awkwardness or for some other reason. Whenever no structure can be found to accommodate all the observational differences, the system operates with several contradictory ones until a new structure is found.

Singer's system is only mildly governed by the criterion of parsimony. That is, it may throw into the a priori structure a great deal that is not directly related to observational reports. It can challenge this affluence if it can propose some structure that will yield a satisfactory analysis of variance and that is in some sense more economical. At any stage it will accept any part of the structure it is using as "objective." That is, it makes no distinction between the objectivity of observation and of theory, simply because no observational report by itself ever means anything. Something is accepted as factually true by Singer's system only when it is a member of, or follows from, a set of observational sentences that differ and which are adjusted by a theoretical structure. Singer's system has no "direct" observations in it, so that it is not embarrassed by its ability to introduce into its structure whatever will make it operate. In psychologist's terms, it makes no distinction between intervening and non-intervening variables, since the whole structure that brings about consistency is as objective as any part. The system does, however, have some difficult decision problems to solve: e.g., what steps to take when two alternative theories both satisfy its criteria in the analysis of variance, and what steps to take when no theory satisfies the criteria. When two or more alternatives are satisfactory, it must decide whether to decompose the reports still further, or to test the adjustment rules in wider contexts, or to discard one or more alternatives on the basis of relative complexity. It must be admitted that in practice such systems are more often faced with the problem of proceeding without a satisfactory theory than in choosing among two or more satisfactory ones. That is, so-called inductive problems are less apt to occur than discovery problems.

18. We have described three alternative ways of designing observation, which clearly do not exhaust the possibilities. Which one should be chosen? It is obvious, I hope, that the traditional answer to this question won't do in the present discussion. This answer is that one should choose that system which best describes how men do ("in fact") learn. Undoubtedly, this was the intent of Locke and Kant, and, to some extent, of Singer. But our interest here is in the design of inquiring systems, so that the manner in which men have learned can be regarded as only one alternative design, and not necessarily the best.

The answer to the question seems to depend on what the inquiring system is supposed to do (Part 2). For example, compare a Lockian system with a Singerian. Locke's observing system puts an end to further inquiry about an object in an environment once it establishes objectivity. Singer's system takes agreement as an unsatisfactory state, i.e., as a problem, and sets to work to create a new problem about an object for which agreement does not obtain. What is the point of either procedure? In Locke's case, the answer must be that the aim of an inquiring system is to create a set of facts that are satisfactory to the employers of the system. Usually this is paraphrased by saying that the aim of inquiry is to satisfy intellectual curiosity. Hence a Lockian observer might work very well in a context in which the system effectiveness was very explicitly tied into states of anticipation and satisfaction of a set of employers of the system, especially if these states were directly and clearly recognizable. We are to imagine that a part of the inquiring system determines what is not known (by a judgment), that it determines that something not known would be wonderful to know (by a feeling), and that it knows directly whether a given process has satisfied its curiosity (by a judgment and a feeling).

Singer's system would seem to work best when the objective of inquiry is to create new problems, because the observer always seeks to put his inquiry to an ever more severe test. The satisfaction of intellectual curiosity for these purposes is irrelevant; the effect of the system is to create dissatisfaction of intellectual curiosity.

Needless to say, one could combine the two aims, or propose another (e.g., a "practical" objective of inquiry). The point is that however the question of objectives is answered, it now seems clear that the observing system is not separable: its excellence depends on Part 2, the part that determines the objectives of the inquiry. In other words, by the principle stated above, since the effectiveness of observation depends on the effectiveness of Part 2, observation is not a separable part of inquiry. We could pursue this point further by showing that the criteria of performance of observing systems also depend on the communication system (Part 11). For example, a communication system that simply transmits observational reports might be effective for a Lockian system, but it would be poor for a Singerian system which requires sufficient information to adjust for disagreement.

19. But it may be more useful to illustrate further the theme of radical alternatives in the design of inquiring systems. For this purpose, consider Part 6 (data collection) and Part 9 (analysis of data), i.e., the parts that are concerned with the design of experiments and the analysis of results. In the so-called well-designed experiment the following conditions must hold: it is agreed upon beforehand how to classify the objects of the experiment and what shall constitute the grounds for accepting or rejecting a hypothesis about these objects. The experiment consists of a series of operations, which result in certain observational reports that classify objects in the prespecified manner, and a series of operations applied to the reports, which lead to the

acceptance or rejection of the hypothesis. We are inclined to say that an experiment that fails to meet these two simple criteria is no experiment at all. That is, some aspects of correct experimental design can be specified independently of any other part of the inquiring system. But can they?

Consider an "experiment" which can be subdivided into a set of phases. At the end of each phase the inquiring system makes one of these decisions: (a) to keep the method of classification and the hypothesis-test procedure, and to gather more data; (b) to change the method of classification and/or the hypothesis-test procedure, and to gather more data; (c) to terminate the data collection and decide on a final classification and hypothesis-test procedure. Suppose we call this the "elusive" experiment. Is it an experiment at all? One is inclined to answer in the negative, for if the objective of the system were to prove something it wished to prove, it could not do better than to follow this design. But this would be contrary to the basic notion of objectivity: namely, that the system's own wishes do not influence its procedures.

Nevertheless, the elusive experiment does occur in a great deal of inquiry. As the doctor examines the patient, he keeps deciding and redeciding how to classify certain responses, as well as how to determine what hypothesis he should be testing. Consider, for example, the somewhat irrelevant challenges that are often leveled at psychoanalytic techniques, because they fail to "prove" that a certain method of treatment really helps or cures a patient. The critics argue that one should select a group of patients that has been treated and a similar group that has not; that one should be very clear at the outset what "treatment" means and what "cure" means; and finally that the test of the hypothesis that more treated than untreated patients are cured should be specified and logically defensible. But one might reply that the cure of any one patient is uniquely appropriate to him and to no other. Hence, only after one has examined the patient can one decide what a cure would be. But examination in this case <u>is</u> the treatment, and furthermore examination is also a unique procedure for each patient. In such circumstances it makes no sense to talk about a cure for an untreated patient, or to specify beforehand what a cure really is. Instead, as the process occurs, the analyst keeps changing his classification method and his concept of a cure.

The obvious reply is that all control has been lost in this procedure. But this need not be the case. For one thing, it might be possible to conduct an elusive experiment in which the decision rules were specified beforehand, so that one knows under what circumstances the inquiring system will make one of the three choices mentioned above. This could be done if the objectives of the inquiry were clearly enough stated so that some value could be placed on the outcomes of each of the choices, and if some estimate could be made of the way in which a choice is related to an outcome.

No one seems to have tried to formalize the elusive experiment, perhaps because present-day inquiring systems in science are largely unimaginative with respect to data collection, which is regarded as a rather dull task, better delegated to less creative minds. But it seems clear that the elusive experiment is a choice the inquiring system could make, and that experimental design is not a separable part of inquiry.

20. The lack of imagination in empirical methods has been partially compensated by the large amount of creative insight one finds in model construction (Part 3). A century or so ago, model construction was constrained by a certain kind of logic (two-valued), by a certain kind of arithmetic (without imaginary numbers), by a certain kind of geometry (Euclidian), by a certain kind of kinematics (absolute simultaneity), and so on. As our science grew, the part of inquiry that generates models shifted to more and more radical concepts of model construction. Perhaps one of our most sophisticated findings is that there are many model choices for a given kind of inquiry, and that model construction is not clearly separable from other phases of inquiry.

21. To finish this essay on the characteristics of inquiring systems, we must return to its central problem: namely, separability of functions.

So far, we have been considering the case of a fixed separation of a part or parts of a system. But this consideration may be too strong. Actually, all that is required is a temporary separation, so that one can proceed in the development of the system. One may want to proceed in this way because he feels that separability is a desirable characteristic of a system. It permits control relative to fairly specific objectives; it permits an adequate scanning of alternatives and a reasonable evaluation of each. Hopefully, then, if one can perform reasonably well in each segment, the whole will be reasonably satisfactory. In the case of inquiry, the dominating criterion of control is objectivity; one wants to be reasonably sure that the evidence for a state of affairs is not itself distorted by the feelings of one of the investigators, or some external but unknown influence. We believe that the larger the system, the greater the risk of non-objective evidence.

Why do we believe this statement about inquiring systems? Possibly because of our recent heritage in matters of inquiry. Modern mechanics began with the study of one body; when this system was understood, mechanics went on to the study of many bodies, of fields, and so on. So in modern game theory, we find students of conflict who believe that one starts with a simple constant-sum two-person game, and that the optimal behavior of people in such a situation can be fully understood regardless of any other conflict situation these people may be involved in, and regardless of any other characteristic of the world.

Thus, even though separability may never occur in a pure and permanent form, shouldn't we act as though it holds, as long as we can legitimately do so? The principle that motivates one to answer this question affirmatively may be stated as follows: so design the inquiring system that some of its parts are virtually separated. Proceed in the study of the separated parts and reconstitute the system only when separability is no longer feasible. This is surely the spirit of much of academic research and a large part of business and government research.

This paper is designed to cast doubt on the adequacy of the principle so stated. In order to make the conclusion more precise, it is necessary to phrase the principle in a more precise way, so that its contradictory becomes apparent.

22. Consider, as before, a system S with subparts $S_i$. Instead of considering the system as a fixed entity in time, consider a method of operating the system as though its parts were separable as long as the system behaves properly, and of changing the parts whenever the system fails to operate properly. We shall call the principle by which a part is changed a "transformation function." The crucial point in the design is whether one can recognize the unsatisfactory state of a part without having to study the entire system. This is equivalent to asking whether the transformation functions are, or are not, functions of the prior states of the parts alone. The separability principle in the preceding paragraph states that the transformation functions are functions of the prior states only. Specifically, it says that one can partition S into subparts in such a manner that

(a) decision rules can be constructed for each part;

(b) the application of the decision rules depends only on the prior states of the parts (i.e., the parts are separable);

(c) there exists a function of the prior state of a part only, which determines whether the part is stable (satisfactory) or unstable (unsatisfactory);

(d) there exists a decision rule applicable to the system that will transform an unstable part into a stable part;

(e) as soon as instability of a part occurs, the system is so transformed that the part becomes stable;

(f) the set of all possible decision rules governing each part, plus the transformation rules that send a part from an unstable state, according to (a) through (e), contains at least one member that is superior to any other rules for operating S, relative to S's objectives.

The reader may recognize that this principle underlies a good deal of present-day reflection about adaptive behavior. For example, the principle is inherent in statistical quality control procedures. Inspection, which plays the role of an inquiring system for production, partitions the production system into parts, identifies the properties of the parts, sets up standards of stability in terms of control charts, signals instability when it occurs solely on the basis of the data obtained from the part, reconstitutes its image of the system until the "assignable" cause is found. Similarly, the model of a "satisficing man" shows him to be one who breaks out reasonably sized problems of decision, pursues a problem to a solution that satisfices, recognizes satisfaction or dissatisfaction clearly, and takes these to be criteria of his control of the situation.

23. The negation of the principle might take one of several forms, depending on what aspects of its lengthy set of assertions one wants to hold fixed. The intent here is not to question the advisability of partitioning systems into parts or of attempting to control the parts. Rather in the examples set forth in the paper, it is clear that the proper performance of a part depends on some concept of the whole system. In other words, we have been saying that the criteria of the stability of the parts depends on a concept of the stability of the system S itself. Therefore, the alternative principle we have in mind is one that modifies rules (b) and (c) and hence (f). It asserts that one ought to determine instability by examining the whole system as well as each part.

This principle reads:

(a) (as before) decision rules can be constructed for each part;

(b') the application of the decision rules depends on the state of the whole system;

(c') there exists a function of the prior state of the part and of the whole system, which determines whether the part is stable or unstable;

(d) (as before) there exists a decision rule applicable to the system that will transform an unstable part into a stable one;

(e) (as before) as soon as the unstability of a part occurs, the system is so transformed that the part becomes stable.

(f') the set of possible decision rules of each part, plus the transformation rules that send a part from an unstable to a stable state, according to (a), (b'), (c'), (d) and (e), contains at least one member that is superior to any other rules for operating S relative to S's objectives.

I said earlier that our recent heritage leads us to be sympathetic with the first-stated principle of the design of systems and not with the second. But older traditions would reverse the preference. We don't need to be reminded of Plato's typical approach to system design: from the general idea of good to the specific details of goodness. But it was the seventeenth century that voiced the rationalism of system design most clearly. Despite the differences in metaphysical theories, Descartes, Spinoza and Leibniz all saw that an understanding of knowledge could occur only if one first came to an understanding of God, who is the supreme inquiring system. The tortuous path of Descartes' meditations led him ever and again to return to the most tormenting of all reflections: the deception of the rational mind. Convinced that deception is possible even for our most firmly established convictions, he could see no solution to be found in the individual inquirer himself. For Descartes there could be no point in a mind's asking itself whether it could conceive the opposite of what it holds to be so obvious. A negative answer could not constitute evidence for one's convictions. In this regard, he differed from his British colleagues, some of whom used and still use the introspective question as the fundamental evidence of their philosophies. But Descartes could construct a model of the universe in which the more convinced a man was, the more wrong he was; the more inconceivable a proposition, the more likely it would be to be correct. This was the universe of a deceiving God. In the language of this report, if God is that aspect of inquiring systems that controls the flow of evidence (data), then a deceiving God would create nothing but self-deceiving inquiring systems. Hence for Descartes the first task of the design of inquiry is to learn enough about God to show that self-deception is not possible. The concept of the whole system dictates, in part, the design of any component.

In Spinoza, the design of inquiry is made most explicit. His axioms are designed to establish the properties and the existence of the most general system, that system which cannot be a part of another. The most relevant assumptions are: "Id, quod per aliud non potest concipi, per se concipi debit" (Ax. II) and "Quicquid ut non existens potest concipi, ejus assentia non involvit existentiam" (Ax. VII) and "Per substantiam intelligo, id, quod in se est, et per se concipitur" (Def. III). From which it follows: "Deus, sine substantia constans infinitus attribis, quorum unumquodque aeternam et infinitam essentiam exprimit, necessario existit," (Prop. XI), and, finally, and most pertinent to this discourse: "Quiciquid est, in Deo est, et nihil sine Deo esse neque concipi potest" (Prop. XV).

Translated into modern systems language, these Spinoza postulate assertions tell us that one cannot establish the existence of contingent objects until one has first established the non-contingent, whole system.

Leibniz's "whole system" concept differs from Descartes' and Spinoza's in the model as well as in the purpose. But he saw much more clearly than did the others the central problem of optimal design. The Monadology is a bold and amazing attempt to delineate the essential aspects of a designed system. Its chief contribution, from the point of view of the present discussion, is that any one system must have all the aspects of any other system. In other words, one should not think of system design in terms of degrees of complexity, and it is incorrect to differentiate one system from another in terms of the number and type of their components. For Leibniz, the correct taxonomy of systems must be expressed in terms of the relative effectiveness of the operations of the components. Systems differ solely in the degree to which they operate effectively. This implies that the theory of systems design is meaningless unless it incorporates a definition of the most general system. It must be emphasized that Leibniz is always talking about systems in the sense of this paper, that is, teleological entities. He is not concerned with the mechanics of systems design.

Thus, "Chaque âme connait l'infini, connaît tout, mais confusiment."[1]

[1] Principes de la nature et de la grâce, par. 13.

The confusion of systems is the critical point of the Monadology and the central concept is the degree of distinct perception. Leibniz actually reserves the word "mind" for those systems having a degree of distinctiveness of perception above a certain minimum (Monadology, p. 19). God is the perfect system and provides the essential standard for the measure of the distinctness of perception (Monadology, p. 48). All systems for Leibniz have two primary functions, which we could translate as "goal seeking" and "perception," corresponding to our modern terminology of "output" and "input."

There are many statements that Leibniz makes to substantiate his claim that systems can only be understood in terms of total systems and that total systems can only be understood in terms of a perfect system. Perhaps the following will suffice to represent these:[1]

"Comme je n'aime pas de juger des gens en mauvaise part, je n'accuse pas nos nouveaux philosophes, qui prétendent de bannir les causes finales de la physique, mais je suis néanmoins obligé d'avouer que les suites de ce sentiment me paraissent dangereuses, surtout si je le joins à celui que j'ai réfuté au commencement de ce discours, qui semble aller à les ôter tout à fait, comme si Dieu ne se proposait aucune fin ni bien en agissant, ou comme si le bien n'était pas l'objet de sa volonté. Je tiens au contraire que c'est là où il faut chercher le principe de toutes les existences et des lois de la nature, parce que Dieu se propose toujours le meilleur et le plus parfait. Je veux bien avouer, que nous sommes sujets à nous abuser, quand nous voulons déterminer les fins ou conseils de Dieu, mais ce n'est que lorsque nous les voulons borner à quelque dessein particulier, croyant qu'il n'a eu en vue qu'une seule chose, au lieu qu'il a ne même temps égard à tout."

A translation into modern concepts (and a condensation) of this passage might read as follows:

"I don't want to prejudge people's intentions and therefore I don't morally criticize contemporary philosophers who wish to get rid of purpose in science and system design. Nevertheless, I must confess that the consequences

---

1 Discours de métaphysique, par. 19.

of their position seem to me to be quite dangerous. They believe that there is no such thing as an overall system or a most general system. On the contrary I hold that it is in the concept of an overall system and its performance that one will find the underlying principle of every component and the performance characteristics of the component, because the overall system is the only standard of good or excellent performance, I want especially to emphasize that one only ends in confusion when he tries to determine an optimal plan solely in terms of some particular design, as though the whole system had only this component to be concerned about, instead of its entire operations."

This part of the rationalist thesis was not new in the seventeenth century, of course, and even today one could easily find adherents to its basic viewpoint. The other aspect of rationalist doctrine is equally important. This is the objectivity of the whole system. To a rationalist mind as to an empiricist, nothing can be admitted to the fund of knowledge that has not passed the most carefully designed criteria of objective truth. This is what makes the "whole system" approach so difficult, and explains why theology and science today have no common meeting ground. A theology that above all must postulate a God regardless of its inability to satisfy the criteria of precise proof cannot expect to find acceptance as a branch of science in an age when the essential feature of science is its strict adherence to standards of precision.

The rationalist thesis in this regard was a very direct one. If God's existence is to be proved, it must be proved simply. This does not mean that it is a simple matter to find such a proof, as the tortuous passages in Descartes and Leibniz clearly show. But most who have tried mathematics have had that quite wonderful experience of finding, after hours or years of labor, a very simple way of proving something that was not obvious at the outset.

The failure of rationalism lay in its inability to find any such simple proof. It was Kant who finally exposed the fallacies of all the proposed simple proofs. The essence of the Kantian refutation was that the conceptual framework required by science to give meaning to experience was not logically strong enough to establish a God in the sense demanded by a Leibnizian theory of reality. In the post Kantian period, Hegel attempted to revise Kant's notion of this conceptual framework, and thereby to establish an Absolute Mind. His Absolute Mind plays exactly the role required of a "whole system," because it establishes the grounds for meaning in any aspect of reality. But Western science, at least, could not tolerate the ambiguities of Hegelian logic, which required contradiction as a necessary condition for proof. To-day contradiction still plays the same role as it always has in Western science; it is that which establishes the stopping point of formal inquiry.

Today we can look at the struggles of the last centuries in a somewhat more relaxed manner than did their philosophers. Do we have to establish the existence of a whole system in order to use the concept in designing systems? Furthermore, are we required to go the whole way, as Descartes, Spinoza and Leibniz thought, and establish an _ens realissimum_, i.e., a system that is perfect in all conceivable respects? Third, were the rationalists correct in asserting that the proof of the existence of the whole system can be established objectively? Since the consideration of these questions will serve to conclude this essay, we can best proceed by proposing some answers in the way of theses about systems design:

(a) The whole system must be real if it has any function at all in system design (realism);

(b) the whole system must be taken to be as perfect as our present estimates allow if it has any function at all in system design (monism);

(c) the proof of the existence of the whole system and its properties must meet the requirements of scientific evidence (rationalism).

Proposition (a) contradicts the philosophy of conventionalism. Few would question the _convenience_ of using constructs that enable us to integrate our empirical findings. For example, students of organization theory often act as though there really were a total organization "out there," just as political scientists sometimes seem to act as though there is such a thing as the federal government, and engineers as though a total generator plant existed. To a strict empiricist, however, these suppositions are merely convenient ways of tying together a series of observations. He would not permit the scientist to claim reality for his construct, since the construct is never observed.

The empiricist has adopted one answer to the problem of the ontological status of sense impressions, an answer that he takes to be based on the principle of parsimony in inquiry: never accept any more than is strictly warranted by the evidence. The difficulty with his answer is demonstrated in the discussion of this paper. If he is telling us how to design inquiring systems, then we must ask whether the kind of parsimony he requires is desirable. The answer to this question depends on the manner in which the inquiring system gathers its evidence. To an empiricist, the inquiring system would say that the rose on the desk exists because it is observed. But it may question whether a rose as such is really there. It will assert that the appearance of a rose exists because it can construct this appearance out of direct observations. In Singer's system, no observation has any claim as evidence until it has been embedded in a sequence of differing observations. The only way in which this embedding can take place is for the system to construct a model of reality in which there are classes of ranked entities. Thus "x is white" is meaningful only if there is a class of entities that can be ranked along a color spectrum. One may still question whether the class itself really exists. One model of reality might assert that the objects of the world can only take on a certain finite set of values in a continuous range. The other values are logical possibilities that are never realized. A second model of reality would assert that the objects of the world can take on any value in the range: there exists one real entity having any given value in the range. The point is that in Singer's inquiring system either of these models is to be construed as a (partial) method of adjusting observations. The first model will take any reported reading and adjust it to one of the finite set of permissible values in the range while the second may adjust it to any value in the range. No observation exists unless some method of adjustment is available. At any stage of inquiry, the inquiring system will not be able to certify one correct model, and hence it will not be able to certify one correct adjusted observation. The real model and the real observation are unattainable limits of the system's activities. But the system requires the existence of these limits in order to operate at all. Hence, the whole system must be real (even though unknown) if it is to function at all in the design of the inquiring system, which is what proposition (a) asserts. It may be noted that neither model (the one that is constrained to adjust to a finite set or the one that is not) is necessarily more parsimonious than the other. Parsimony, in Singer's system, must be defined in terms of the costs of operating the system and not in terms of the simplicity of the entities of the system.

Thus, the validity of (a) depends on the settlement of an issue regarding the design of the system. I have given a very specific illustration of how this issue might be considered. Another illustration could be developed around the problem of information flow. Should the inquiring system regard an observation to be real independent of the method of its transmittal, or does an observation gain tentative acceptance only by virtue of the fact that it can be shown that it can be transmitted? If the latter answer seems best for the purpose of systems design, then one must accept the reality of a transmittal system if one accepts the reality of an observation. In more general terms, one might argue that most organizations "exist" because without organizations there is no such thing as information.

Proposition (a) makes no commitment about the meaning of "whole." Thus, one might argue that only small sections of a total possible reality must be assumed to exist, even though one designs inquiry along the lines that Singer's analysis suggests. But Proposition (b) asserts that the only adequate definition of "whole" is in terms of a perfect system.

In philosophical tradition, X is perfect if it is not limited in some respect. In other words, the general property "good" can be subdivided into a set of properties: intelligent, beautiful, knowing, powerful, and so on. Entities having these properties can be ranked, so that for example, "is more powerful than" orders the objects of the world. For each such property there is a maximal entity; e.g., an entity which is more powerful than any other entity. Finally, it is asserted that the maximal entity in all the properties of goodness is exactly the same. A most intelligent entity is also most powerful and most beautiful. The <u>ens realissimum</u> is that entity.

Thus Proposition (b) breaks down into a set of assertions, which can be more formally stated as follows:

(1) The world consists of a set of entities.

(2) The entities of the world can be strictly ordered by a set of relations.

(3) For each relation there exists an entity that is maximal.

(4) The maximal entity for all relations is the same.

(5) For all systems designs, the "whole system" must be the entity satisfying assertion (4).

Those who wish to restrict the whole system to what seems to be practically conceivable would deny (5). They would, in effect, argue that the theory of systems design does not have to commit itself concerning the properties of the most general system. In other words, they would not agree with Leibniz's stipulation that all systems are basically alike. One source of the disagreement might result from terminology; those advocating non-general whole systems may be thinking of systems in ateleological terms, or at least in terms that have no relevance to ultimate purposes or destinies.

But the real opposition to Proposition (b) is to be found in the disagreement with (4). Most systems designers go as far as they can in trying to conceive a system that will be best for some specific purpose. But they do not feel that these systems are best for all purposes. A missile system may be designed by the designer trying to conceptualize what an ideal missile should do, e.g., it is one that destroys an enemy stronghold perfectly. If so, he does imagine a perfect system within the limits of his imagination. But he would hardly say that the missile system was perfect in all respects. It is not very good at all for producing consumer goods, for example.

The Leibnizian answer is fairly obvious, of course. It simply says that for every missile-system designer there must be another systems designer who considers the missile system as a part of his system. Such a designer also tries to conceptualize the perfect system. For him, the ideal missile may not be one that destroys perfectly. It may, instead, be one that perfectly prevents destruction. In this case, the original missile designer made a mistake in his selection of the relations that rank entities. In short, only if one conceptualizes the most general system will one know what relations are appropriate in ranking entities.

Leibniz's "single system" concept can be formulated in another way. In our culture, we typically segregate the functions that men perform, in terms, say, of the professions of research, law, education, industry, government, and so on. The professions come in contact only on the periphery, so to speak, where a man of one profession consults a man of another. In the consultation, the one learns about the results of the other's deliberations, but does not take a hand in framing the results. This is essentially a partitioning of our social institution into presumably separable parts. Each profession can be understood by itself, by understanding the manner in which it works and the principles that guide its actions.

But suppose one were to deny all this segregation of the professions, and were to say, for example, that one cannot understand science unless one has understood it as a management profession, or a political activity, or a legal activity. For example, one might argue that science can manage an enterprise, or a part of it, and that operations research is just such a way of viewing science. One might further argue that there is some optimal way in which science can manage: an ideal of scientific management. Finally, one might argue that a necessary condition for understanding what science is, is the understanding of how it can and ought to manage.

It must be emphasized that all along we are discussing the design of systems. Hence, the question is not to understand how present-day science can manage, because this may be a very bad design question. The question is, how would science have to be designed in order for it to be a management? In other words, what is the design of a science which makes science an optimal management system?

By the same token, a necessary condition for a full understanding of management is to conceive of management as a science. This, indeed, is what is happening in research and development, where management is playing a stronger and stronger role in the planning of the phases of research, e.g., as in SCANS and PERT. There is an activity called the management of science. There is an optimal way in which management becomes a science, i.e., a generator of information. To understand management, one must understand it as a science. In other words, one way to understand management development is to determine in what way management can become a scientific system.

The same theme could be repeated in many contexts. To understand science, one must understand it as a legal profession, and to understand the law, one must understand it as a science. For example, T. A. Cowan argues that law is the system of controls for experimentation in the social sciences. It seems to me that he is trying to conceive of law as a science. I know of no one who has yet been bold enough to suggest how science becomes the law, except in the bad sense of a science that controls thought processes.

Even within the institutions themselves, the same principle could be applied. For example, one cannot understand psychology until one has understood in what way psychology is a physical science; i.e., understands how psychology must change so that it becomes a physics. The reverse is all too familiar to students of the history of science: the attempt to understand physics as a psychology.

Thus, the system's designer does not understand his system until he understands it in terms of all the basic functions. The designer of a missile system must understand how the missile system is a productive system, a communications system, an inquiring system, and so on, <u>if</u> he is to understand his system fully.

So put, Proposition (b) becomes intelligible, although not necessarily acceptable. It states that the understanding of systems demands a monism, a method of looking at all systems in one way.

Clearly, the Proposition demands some reasonable taxonomy of systems in order to be usable. If systems of type x must be understood as systems of type y, what ranges of concepts do x and y entail? Our present taxonomy, which has grown out of the tendency to separate functions, may be far too awkward to apply. Ackoff and I once suggested a four-way taxonomy: all systems must be conceived as <u>scientific</u> (discerning the proper means to ends), and <u>productive</u> (developing new means), and <u>cooperative</u> (coordinating teleological entities), and <u>changing</u> (creating interest in new ends).[1]

But such a taxonomy is obviously only a first step in the direction of a monism of systems.

We should note that Proposition (b) together with (a) implies that the perfect system "exists." Of course, we have not demonstrated such existence.

All we have done, in arguing that (a) is sensible, is to assert that "existence" does not mean "observed" or "observable." Instead, for something to be taken to exist, it must be assumed essential in the development of inquiry. One cannot separate out segments of inquiry and stamp "existence" or "reality" on these alone, because, according to the argument, these segments exist as segments only because of the rest of the system. We never know what really exists, but at any time we do the best we can to construct an image of the world in which our observations, our thinking, our feeling, our intuition will live as well together as possible. We <u>take</u> such an image to exist; but it is so taken only because we argue that there is a whole image of which ours is an approximation.

[1] Churchman, C. W., and Ackoff, R. L., <u>Psychologistics</u>, University of Pennsylvania, 1946, <u>Methods of Inquiry</u>, 1950. See also C. W. Churchman, <u>Prediction and Optimal Decision</u>, 1961.

So viewed, Propositions (b) and (a) state a hypothesis about reality: namely, that there is an *ens realissimum*. But far from the proof of the hypothesis being simple, it never will be proved. It is the most complicated hypothesis possible.

This brings us, finally, to Proposition (c), which states that the task of proof lies in the hands of inquiring systems. To summarize all that has been said here, inquiring systems can be partitioned, but an understanding of any part requires an understanding of how that part can function as another part: for example, how the observing part can be conceived as the part that develops models, or determines the goals of inquiry. Finally, then, Proposition (c), along with (a) and (b), implies that we cannot understand the nature of inquiry until we understand in what way inquiring systems are theological.

## Debts and References

I have found much value in the papers in General Systems, the yearbook of the Society for the Advancement of General Systems Theory. The reader's attention might be directed to the section "On Myths of Parts and Wholes," in Stafford Beer's "Below the Twilight Arch," General Systems, Vol. 5, 1960, pp. 15-17. Of special interest for the purposes of this paper is an article by A. D. Hall and R. E. Fagen, "Definition of System," which appeared in Vol. 1 of General Systems. The authors define systems in a manner quite similar to the definition that appeared in an earlier draft of this paper, where I defined a system as a set of entities, operators, and relations, with the usual formal rules governing the formation of new entities, of meaningful assertions, and valid assertions, with the further stipulation that the language be rich enough to enable one to speak of the whole system. My chief interest then was to define "real" systems. The definition was clearly inspired by recent writings on formal systems. This seems to be the case in the Hall-Fagen paper, where a system is defined to be "a set of objects together with relationships between the objects and between their attributes." The authors also introduce the concept of the "independence" of parts; a part is independent if a change of the part depends only on that part alone. In this paper I discarded the definition of systems in terms of formalized languages in favor of an emphasis on the effectiveness measures of the parts and the whole, because only thereby could I discuss the kind of independence or dependence that interested me, namely, the criteria by which systems are designed. It seemed to me that the formal-system approach takes the parts to be given, as well as the rules that govern the behavior of the parts, whereas I was interested in speculating about how one decides whether something is a part and how one evaluates the decision once made.

SDC

SYSTEM DEVELOPMENT CORPORATION · 2500 COLORADO AVENUE · SANTA MONICA, CALIFORNIA