SYSTEMS THEORY, KNOWLEDGE, AND THE SOCIAL SCIENCES

Richard N. Langlois

It is not too much of an overstatement, I think, to suggest that systems theory has not lived up to the great promise it has long been supposed to hold for the social sciences. Considering the magnitude of the claims, of course, disappointment was probably inevitable. Systems theory has been hailed by some as a way of unifying the methodology of all sciences. Others have intimated that this approach would finally endow the social sciences with that natural-science-like rigor and precision those disciplines allegedly lack. And others—more modestly—have looked to systems theory as a way of combatting the fragmentation and specialization of the sciences.

But systems theory has failed—at least so far—in its various attempts to bring all the disparate parts of inquiry under the sway of its organizing force. In fact, it is fairer to say that systems theory has itself succumbed to the diversity and complexity of modern scientific inquiry. Defining what we mean by systems theory or the systems approach is virtually impossible outside the context of a particular discipline, and we might almost say that there are as many versions of systems theory as there are would-be systems theorists.

In trying to be all-encompassing, systems theory has become unsystematic; and in trying to become systematic, it has become narrowly specialized. On the one hand, the attempt to describe at a broad level the system-theoretic approach to this or that inevitably ends up sounding like a Sears Roebuck catalog of vaguely connected concepts, models, and definitions. On the other hand, mathematical system theory[1]—the most rigorous and

This research was supported in part by the Division of Information Science and Technology of the National Science Foundation under grant IST-8110537. Note: Portions of this paper appeared under the title "Systems Theory and the Meaning of Information" in "Perspectives on Systems Methodology and Information Research," a supplement to the *Journal of the American Society for Information Science,* vol. 33 (November 1982), pp. 395–399.

[1]See Hassan Mortazavian's contribution to this volume.

tightly knit version of the systems approach—has coagulated out of the more general body of systems ideas to become a well-circumscribed mathematical specialty, even if one with certain imperialist pretensions.

I say this largely by way of apology—for I do intend in this paper to provide a perspective on systems theory. But because systems theory is, in fact, so multifarious, my portrayal of it will necessarily be selective, picking out those pieces that I find useful and letting the rest alone. This seems to me the only sensible compromise between the encyclopedic and the axiomatic.

My general concern will be with systems theory's relation to the social sciences, especially economics. More particularly, though, I will be concerned with the notions of knowledge and information and the ways in which systems theory can help illuminate those concepts. Indeed, I am hopeful that this essay might prove a useful preface to some research I am now beginning on the connection between the information sciences, broadly construed, and economic theory.

WHOLES AND PARTS

Perhaps the principal reason systems theory has failed to revolutionize scientific methodology is that at the broadest and most general level, it has nothing particularly revolutionary to offer.

At its most philosophical, systems theory is a confrontation with the age-old problem of *the whole and the parts*. Systems theorists discovered—or, rather, rediscovered—complexity . They saw the approach of the physical sciences, which is supposed to be analytic or reductionistic, as applicable only to phenomena of "organized simplicity." In dealing with the increasingly prevalent problems of "organized complexity," they argued, we cannot understand the parts in isolation from the whole. [Weaver, 1948.] Thus, we need to study the total system, to see the big picture. These are not new ideas, of course. In fact, it is striking the extent to which modern systems theorists—notably proponents of so-called general systems theory—have tended to retrace a lot of old footsteps in the long-standing debate over the doctrine of methodological holism. [D. C. Phillips, 1976, esp. pp. 46-67.]

Notwithstanding the holist rhetoric of most systems theorists, however, systems theory is not at all methodological holism in any strong sense. And at its best, systems theory is, in fact, a form of intelligent methodological individualism.

Since this is a somewhat unconventional thing to say, and since the individualism/holism debate is so frequently misunderstood, perhaps a brief digression is in order.

The conventional wisdom runs something like this. There is a doctrine in the social sciences called methodological individualism. It is a form of the analytic or reductionist method, and it therefore holds that knowing the properties of the parts—the individuals in society—is fully sufficient for grasping all there is to know about the whole—society. In other words, the properties of the whole can be deduced from the properties of the parts; or to put it in more familiar (and more naive) terms, the whole is just the aggregate or sum of the parts. This view is to be contrasted with methodological holism (or, sometimes, collectivism), which insists that wholes possess "emergent" properties that cannot be derived from the properties of constituent parts. To the holist, the whole is greater than the sum of its parts.

This is a familiar story. But, as is often the case with familiar stories, it is almost entirely wrong. There may well be some writers who espouse in principle this sort of naive methodological individualism. But no one can actually put such a view into practice. Consider the case of a neoclassical economist analyzing the effect of a tax suddenly placed on a certain commodity. In proper individualist fashion, he or she will instantly begin to model the choice problem faced by a representative economic agent, discovering, as always, that the individual will most likely consume less of the commodity than before. But this finding about individual demand tells the economist nothing about the whole—total demand—until he or she adds a global fact: that total demand is the sum of individual demand. In this (trivial) case, the whole *is* just the sum of the parts; but even here, the whole could not be deduced from the parts, since the relation among them-addition—is not logically contained in the individual-choice model itself. Moreover, there is nothing sacred about addition, and methodological individualists are quite willing to specify very different sorts of relations among the parts. In economics, for example, the aggregate result is very often exactly the opposite of what we would have expected from considering individual behavior alone.[2]

Far from denying the importance of emergent phenomena, the economist and philosopher Friedrich A. Hayek, widely (and correctly) cited as an archproponent of methodological individualism, reminds us that the entire *objective* of the social sciences is to explain how the behavior of individuals leads to orderly patterns and institutions that none had consciously planned—to explain, in other words, the emergent results of individual action. [Hayek, 1979, pp. 146-147.]

In sophisticated discourse, the question of emergent properties is in no way an issue between individualist and holist. *Both agree* that social phenomena must often be considered emergent wholes whose behavior cannot

[2] In his best-selling introductory textbook, Paul Samuelson (a methodological individualist in microeconomics, at least) finds it necessary to warn the student against fallacy-of-composition errors as early as page 11 of the text, where he presents a long list of such fallacies exploded by economic analysis. [Samuelson, 1980.] Indeed, Kenneth Arrow, Samuelson's fellow Nobel Laureate, has written that "the notion that through the workings of an entire system effects may be very different from, and even opposed to, intentions is surely the most important intellectual contribution that economic thought has made to the general understanding of social processes." (Arrow, 1968, p. 376.)

be entirely reduced to the behavior of individuals.[3] The issue is not whether we can deduce the nature of the whole from the properties of the parts; the issue is whether or not we should consult the parts *at all*. Holism is not the doctrine that we should study emergent phenomena; holism is the doctrine that we should somehow study wholes *directly* without considering the workings of the parts in a meaningful way.

It should be immediately clear that this position does not follow at all from a recognition of systemic interactions or emergent phenomena. The methodological individualist cannot deduce such phenomena from his or her knowledge of the parts; but—what is often overlooked—the holist is in no better position to understand such phenomena than is the individualist. The problem is simply a lack of knowledge about emergent properties, and calling oneself a holist does not instantly convey that knowledge. [D. C. Phillips, 1976, pp. 14-15.]

More to the point, adopting the holist stance arguably puts us in an epistemological position *inferior* to that of the individualist. The holist differs from the (sophisticated) individualist only in that the former insists on *throwing away* useful information. Here lies the real disagreement. The methodological individualist holds that the social scientist should always keep in the closest contact with the level of the parts (the individuals) and utilize fully whatever knowledge of the parts—however incomplete—he or she can bring to bear. [Machlup, 1969 and 1979b; Hayek, 1979, esp. chaps. 4, 6.]

The best way to appreciate this may be by analogy with literary criticism. No one would deny that a work of literature is more than the sum of the words and sentences it comprises. Yet the modern critic insists on studying these words and sentences carefully. Similarly, methodological individualism in social science is nothing more than an insistence on "sticking to the text," whereas holism is a license to engage in that most heinous of "lit-crit" solecisms, "reading in."

Although naive methodological individualism is probably impossible in practice, holism in this sense is not. Social science is rife with holistic formulations in which hypostatized concepts like *society, the capitalist class*, or *the public interest* take on operational significance in and of themselves—formulations in which, as Jacques Barzun said of Marx, the terms of reference become "entities stuffed with people but not composed of them." (Barzun, 1958, p. 182.) One result, ironically enough, is that holists are much more prone to fallacy-of-composition errors than are methodological individualists. For if we throw away information about the parts, we are inclined

[3] "The overall order of actions in a group is in two respects more than the totality of regularities observable in the actions of the individuals and cannot be wholly reduced to them. It is not so only in the trivial sense in which a whole is more than the mere *sum* of its parts but presupposes also that these elements are related to each other in a particular manner. It is more also because the existence of those relations which are essential for the existence of the whole cannot be accounted for wholly by the interaction of the parts but only in their interaction with an outside world both of the individual parts and the whole." (Hayek, 1967, pp. 70-71.)

directly to read in a logic of operation for the whole from some other source; and the logic of operation nearest at hand is that of the individual human. [Hayek, 1979, p. 101; Langlois, 1981, chap. 3.] The notion that, for example, a society permitting economic self-interest is therefore a greedy society is a bit of naive individualism characteristic of holists far more than of methodological individualists.

In order to make a plausible case for holism, we have to argue that less information is somehow better than more. This is no easy task. We might for instance invoke the Gestalt and assert that attention to the parts destroys our understanding of the whole. But this is to confuse perception and understanding. And, as my analogy with literary criticism was meant to suggest, the sophisticated methodological individualist has no compunction against stepping back to survey the Gestalt—so long as the (epistemologically more accessible) parts are also carefully analyzed. Historically, holists have taken a rather different (if not ultimately unrelated) line of attack. Hegel and his followers argued that the parts could not be studied in isolation because the parts acquire their very nature from their relation to the whole, which nature is necessarily altered if the parts are considered apart from that whole. But if taken at all seriously, this argument leads to logical absurdity, and it falls quickly apart when translated from the quasi-mysticism of essentialist rhetoric into a modern nominalist vocabulary. [D. C. Phillips, 1976, pp. 5-20.]

Proponents of general systems theory have unknowingly reinvented and invoked this Hegelian formulation in discussing the holism of systems theory. [Ibid., esp. pp. 45–50.] But most systems theorists—both mathematicians and practitioners alike—conceive of systems theory in a way antagonistic to this Hegelian view. They are certainly inclined to recite holist cant about "phenomena that depend on the conditions of the environment in which they exist, interact with this environment, and thus cannot be properly studied in isolation." (Mortazavian in this volume.) But all they mean by this is that the behavior of the whole cannot be understood without knowledge of the relations among the parts. The parts are conceived of as logically distinct elements of a mathematical set; those elements exist and are fully defined independently of any relations that might be specified. A system is just the set of parts plus a set of relations among the parts.[4] This is a formulation that would trouble a serious holist far more than it would a methodological individualist.

The philosopher Mario Bunge has recently transformed this set-theoretic definition of a system into a methodological position—systemism—that "combines the desirable features of individualism and holism." (Bunge, 1979a, p. 13.) To Bunge, a society (which he finds it necessary to call a) may be represented as the ordered pair (S, R), where S is the set of individuals in

[4] Cf. Mihajlo Mesarović's "implicit (syntactical) definition" of a system in Mesarović [1964a, p. 7].

the society, and ***R*** is the set of relations among them. That systemism does in many ways combine the best of both worlds is, I believe, an entirely unobjectionable assertion. What is ***not*** true is that systemism somehow represents a new methodological alternative: The basic ideas of what Bunge calls systemism are essentially identical to what sophisticated methodological individualists have believed all along.[5]

It is important to notice that while Bunge's definition of society qua system is consistent with the general mathematical definition of a system [Mesarović, 1964a, p. 7], not all applications of systems theory are in accordance with the tenets of systemism understood as intelligent methodological individualism. The "parts" in Bunge's formulation are individuals in society, but this is by no means always the case in social science. The elements in a typical systems model are aggregate variables of one sort or another—dollars, commodity levels, energy use—that are represented as impinging directly on one another with no reference whatever to the existence of human beings. A good many systems models in the social sciences must, therefore, be classified as instances of naive holism; and while such models may often prove interesting and illuminating, they cannot—as is often claimed—serve as a general foundation for economics and other social sciences. This is a point to which I will return.

Perhaps I should apologize for the length of this digression. But, in a sense, perhaps this has not been a digression at all. After all, philosophical issues lie nearer to the surface in systems theory than they do in most intellectual endeavors, even if they are not for that reason more often perceived. More importantly, though, there is a sense in which it is methodological individualism that lies behind the systems-referential view of knowledge and information that I now wish to present.

SYSTEMS AND THE MEANING OF INFORMATION

There is a strain of thought on matters informational that I find somewhat disturbing. It is what I think of as the "oil-flow'' model of information, with its attendant "oil-tank'' model of knowledge. According to this (implicit) view, information is some sort of undifferentiated fluid that will course through the computers and telecommunications devices of the coming age much as oil now flows through a network of pipes; and the measure of our

[5] Bunge does admit that Hayek and other methodological individualists recognize "the reality of social relations" (p. 17), but he continues to paint them as naive individualists. "The individualist might not wish to dispute the systemist's thesis," says Bunge, "but, if he is consistent, he must insist that the structure of ***R*** is somehow 'contained' in, or deducible from, the properties of the individual members of society" **(p. 19).** But, as we saw (cf. again footnote **3),** Hayek for one does ***not*** so insist, and I cannot see why he is therefore inconsistent in any way. Bunge is simply mistaken in his characterization of the individualist position. [Bunge, 1979a.]

knowledge in this world will be the amount of "info-fluid" we have managed to store up.[6]

This model of knowledge and information has some mutually reinforcing affinities with information theory in the well-known Shannon-Weaver sense, which developed for purposes of communications engineering a quantitative measure of something called information. [Shannon and Weaver, 1949.] Communications theorists are themselves quick to deny the connection between their concept of information and the term's everyday meaning, distinguishing not only between semantic and nonsemantic information but also between the concept of information itself and the notion of *amount of information.* The logic of systems can help illuminate these distinctions, I believe, and may even prove able to offer a conceptual model for information and knowledge alternative to the pipe and tank.

It is tempting to think that the distinction between the communications theorist's *information* and the more general sense of the term is at base a matter of mechanism versus antimechanism (the quantification of information is something appropriate to machines — to computers and switchboards, wires and transmitters), and that the very different character of information and knowledge in everyday experience arises from the distinctly nonmechanistic nature of the human system. There is more than a grain of truth in this, and I will later suggest that such a dichotomy between mechanistic and nonmechanistic systems is indeed in order. But, in the end, this view gives us only part of the picture. To a large extent, I believe, the real issues of information and knowledge actually cut across the mechanism/antimechanism dichotomy.[7]

To see what this means, let us look at how even the most mechanistic sorts of systems models use a concept of information.

Systems theorists often distinguish between terminal or causal systems and goal-seeking systems. [Mesarović, 1964a, p. 21; 1962, p. 13.] The former react to their environment according, as it were, to the logic of proximate cause: The environment affects the system's inputs, which, in turn, affect the behavior of the system in a strictly preprogrammed fashion. By contrast, goal-seeking systems operate according to something nearer the final cause. In this case, there are "certain invariable aspects of the system which reflect its goal" (Mesarović, 1962, p. 13); a teleological subsystem receives the system's inputs and guides its behavior in light of the goal. The distinction

[6]This oil-flow model is not without implications. More than one writer has suggested that information policy be predicated on the inevitable dependency of future society on this "info-substance" just as energy policy is supposed to deal with our dependence on oil. As a consequence, we should worry about the availability of "info-fluid" to disadvantaged groups like Chicanos much as we recently used to fret about the availability of heating oil to poor New Englanders.

[7]This is precisely the theme pursued by the British physicist Donald MacKay. [MacKay, 1969.] The next section of this paper will draw heavily on his ideas.

can perhaps best be illustrated using the familiar stimulus/response (**S/R**) model from behaviorist psychology. In the causal approach, we posit some direct mapping of the stimulus to the response; each stimulus directly causes a particular response. In the goal-seeking approach, we assume a slightly more sophisticated version of **S/R** in which an intermediate processing stage comes between stimulus and response.[8] The stimulus, of course, is the information with which we are concerned.

As a kind of objective correlative, we could think of a stimulus as involving a (weak) form of energy—a small electric current, a pattern of light, a sound wave—which elicits as response the release of another form of energy (often stronger or different in character).

In both causal and goal-seeking models, we can talk about the *meaning* of a piece of information. In the former case, the meaning of a signal is the response it elicits. In the latter, response is also the ultimate criterion of meaning, even if we cannot necessarily understand the meaning of a signal without first knowing the goal that the system is pursuing.

Now, it is important to notice that not all signals will be equally meaningful. A signal of the wrong form—the wrong voltage, frequency, or code, for example—cannot be understood by the system. At the same time, not all meaningful signals will have the same meaning. A small voltage may elicit a smaller response than—and thus have a meaning different from—that of a larger voltage. A message coded ABCD may have implications for achieving a system's goal that are very different from those of a message coded DCBA. And this provides a clue to the distinction between semantic information and the nonsemantic information of communications theory. The latter is concerned only with the extent to which a message is within the set of meaningful messages; it is not at all concerned with the message's implications for the system. For example, the messages ABCD and DCBA—which contain the same characters—would typically have identical information content in the communications-theory sense, even if they had very different meanings from the system's point of view.

To illustrate this, let me draw on an example from the (normative) theory of economic decision-making. Here we are dealing with a goal-directed system, one in which the "teleological subsystem" takes the form of a decisionmaker who maximizes profit (or, as we shall wish to introduce stochastic elements, who maximizes so-called expected profit, which is, in effect, profit weighted by probability).

Let us suppose that the decisionmaker is a farmer and that he has certain "decision variables" under his control, for example, the number of acres of wheat he can choose to plant. There are also variables—say, the weather-that are not within his grasp. The profit experienced by the decisionmaker depends on both the weather and the number of acres planted. If the de-

[8] Mesarović holds—correctly, I believe—that any mathematical system can be represented either way. A causal system can be modeled as if it were pursuing a goal; and a goal-directed system can be reduced to a direct mapping of inputs to outputs. [Mesarović, 1964a, p. 22.]

cisionmaker knows the weather—if he knows how much rainfall there will be—he can easily determine the acreage appropriate for maximum profit. If he is uncertain about the level of rainfall, he can articulate a probability distribution for that variable, which allows him to select the acreage that maximizes expected profit. But the decisionmaker who can anticipate the level of rainfall perfectly will very likely elect to plant an amount of acreage different from that selected by the farmer who is uncertain about the weather; more to the point, the better informed decisionmaker will realize a higher profit than his ill-informed counterpart.

If the decisionmaker who is uncertain can obtain some kind of information about the weather—divine revelation, perhaps, or the weather bureau's probability distribution for rainfall—then he can improve his acreage decision. But if this information is costly, the decision whether to acquire it can be represented as yet another decision-system of the same sort as the original acreage decision. In general, our profit-maximizing farmer should wish to acquire the information if its expected value exceeds its cost. [See, for example, Howard, 1966.]

This *expected value of information* is the interesting quantity. In order to examine it more closely, let us simplify our example even further. Suppose there are only two possibilities, heavy rain and light rain. The farmer must optimize his acreage planted in light of whatever information he might have about which of these two possibilities will occur; for example, he might consider the amount of rain experienced in past years or various freely available predictions by the Weather Bureau or the Department of Agriculture. On the basis of this imperfect information, the farmer estimates a probability of, say, 0.7 for heavy rain and 0.3 for light rain. Now, suppose that a fully reliable clairvoyant stands willing (for a stiff fee) to disclose which of the two alternatives nature actually has in store. How much is this revelation worth to our profit-maximizing decisionmaker? We have to consider how the new information would affect the farmer's acreage choice and, thus, his profits. If the clairvoyant says "heavy rain," he can optimize the acreage in a way that increases profits over the expected level. If the clairvoyant sends the message "light rain," he can also optimize acreage—in a different direction—to increase profit over the expected profit. In deciding whether to buy the clairvoyant's information, the farmer must decide if the increase in profit will be enough to justify the clairvoyant's fee—and he must decide *before he knows which message the clairvoyant will send*. So, the decision whether or not to buy the information must be based on the farmer's prior probability assessment of what the clairvoyant will say. Since the clairvoyant is merely revealing what nature will do, the farmer's weather forecast or other prior information is exactly as relevant to predicting the clairvoyant's message as to predicting nature's response—since the former is nothing but the latter moved forward in time.

Thus, the farmer must use his original probability assessment (0.7 chance of heavy rain; 0.3 chance light rain) in deciding whether to buy the information. And the expected value of perfect information is thus the sum of two magnitudes: the first is the increase in profits from adjusting to heavy rain if there will, in fact, be heavy rain, multiplied by 0.7, the probability that there will be (and, therefore, that the clairvoyant will say) heavy rain; the second is the increase in profits from optimizing for light rain given that there will, in fact, be light rain, times the probability that there will be (and, therefore, that the clairvoyant will say) light rain.[9] In other words, the expected value of perfect information (EVPI) is the sum of the value to the system of a set of possible messages weighted by the probability of occurrence of each message.

By contrast, the information measure of communications theory is oblivious to the value a message holds for the system that receives it: In analogy with the entropy measure of thermodynamics, the information content of the message ***heavy rain*** is, to the communications theorist, proportional only to the logarithm of the probability of that message being received.[10] Although the value measure and the entropy measure can move in the same direction, there is no general reason why this should be so.[11]

[9] The length of this sentence suggests that there are times when mathematics has its expository (or at least space-saving) advantages. Let $\Pi(a|w)$ be the (realized) profit from planting a acres under weather conditions w. The farmer's original problem is

$$\max\,[p\Pi(a|w = H) + (1 - p)\Pi(a|w = L)],$$

where H indicates heavy rain, L light rain, and p is the farmer's prior probability assessment on heavy rain. (The probability of light rain has to be $(1 - p)$ since there are only two possibilities.) Suppose that a^* is the value of a—the number of acres—that maximizes the quantity in brackets. Then,

$$\text{EVPI} = p[\Pi(a^*_H|H) - \Pi(a^*|H)] + (1 - p)[\Pi(a^*_L|L) - \Pi(a^*|L)],$$

where p is the farmer's assessed probability that the clairvoyant will call for heavy rain (which is necessarily identical to his original assessment of the probability of heavy rain), a^*_H is the optimal choice of acreage under conditions of heavy rain, and a^*_L is the optimal acreage under light rain conditions.

[10] In the notation of footnote 9, the (nonsemantic) information content of the message *heavy rain* is proportional to $-\log_2 p$. The information content or entropy of the light-raidheavy-rain information system is the information content of each possible message weighted by its probability,

$$H = -[p\log_2 p + (1 - p)\log_2(1 - p)].$$

[11] In our now familiar example, it happens that the two measures are closely related: Information content is high when value is high and vice versa. (The reason is that a high-entropy [low "info-content"] message is one with a high probability. If a farmer anticipated heavy rain with a high probability, then a^* is likely to be already near a^*_H, and $[\Pi(a^*_H|H) - \Pi(a^*|H)]$ is low. Similarly, the more unexpected message would have both a higher information content and a higher value, since a^* would be less close to a^*_L. I believe this monotonicity can be shown rigorously to hold if the profit function is concave.) But such a connection is merely fortuitous. Consider an industrial research laboratory, for example, in which one experiment has a very high "info-content" (e.g., the experiment has a 50-50 chance of resulting in the message *success)* but low profit implications for the company, while a second experiment has a low "info-content" (e.g., a 90 per cent chance of success) but big profit implications. The research manager who used entropy as a decision-making criterion would be sorely misguided.

MEANING AND STRUCTURE

Stimulus/response systems—what we could call cybernetic systems in a broad sense[12]—are mechanistic systems in that action alone matters. It is not so much that action—response—has taken the place of meaning; rather, meaning itself is ***defined*** solely in terms of the action released or controlled by the action in question.[13] In the case of our agricultural decision-system, for example, the meaning of a message about the weather is the acreage the decisionmaker chooses to plant. Now, because this **is** a goal-directed system, we can interpret the relation between message and action in light of the goal, and since the goal is a simple quantitative one, we can even measure the value of the message for the goal's achievement.[14] Nonetheless, all meaning is ultimately reflected in the system's behavior, in its output.

Yet, if we look at the question of meaning in a slightly different way, we may be able to generalize beyond simple stimulus/response systems. Another way to say that the meaning of a message within a cybernetic system arises from the action that message brings about is to say that the meaning is defined ***by the system itself.*** What makes the message ***rain will be heavy*** comprehensible to our decision-system—and the message ***tea tonight at eight*** meaningless to it — is that, by virtue in this case of its very structure, the system is ready to respond to the one and not the other. Furthermore, notions of information and meaning seem as applicable to complex and arguably nonmechanistic systems (such as brains) as to simple cybernetic systems.

In general, then, we might follow Donald MacKay in speaking of a system's structure as defining "conditional states of readiness." It is the overall configuration—not any particular response in isolation—that determines the meaning of a message.

> It isn't until we consider the range of other states of readiness, that ***might have been considered but weren't,*** that the notion of meaning comes into its own. A

[12] By a cybernetic system in a broad sense, I mean any system that is concerned with information and control. In a narrower sense, cybernetics is concerned specifically with feedback systems, where a monitoring signal is sent from output back to input in order to control that output and bring the system into an equilibrium condition called "homeostasis." The locus classicus here is Norbert Wiener's ***Cybernetics.*** [Wiener, 1948 and 1961.]

[13] "Depuis longtemps, le pragmatisme et le behaviourisme ont appris aux psychologues à mettre l'accent sur l'action plutôt que sur la conscience. La cybernétique adopte rigoureusement ce point de vue: le sens, la conscience dans l'information, n'a rien d'essentiel; ou plus exactement, le sens d'une information n'est rien d'autre que l'ensemble des actions qu'elle declenche et contrôle.'' (Ruyer, 1954.)

[14] Here, I think, we need to make the distinction between the (1) meaning, (2) meaningfulness, and (3) value of a message. A message to the decision maker reading *keep doing what you're already doing* is fully meaningful even though it does not entail a different action or necessarily result in a level of goal achievement higher than would have occurred in its absence. (Of course, a strict behaviorist would be unable to distinguish a meaningless message from a meaningful message to maintain the status quo, but this is not a problem in a goal-directed model.)

> change in meaning implies a different selection from the range of states of readiness. A meaningless message is one that makes no selection from the range. An ambiguous message is one that could make more than one selection. (MacKay, 1969, p. 24; emphasis original.)

MacKay offers the metaphor of a railroad switching yard in which the configuration of tracks and switches stands ready to direct the trains passing through it. By sending the right electronic signal—or, in older yards, by inserting the correct key in a switch box—we can rearrange the configuration of tracks. The meaningfulness of a message thus depends on its form-the shape of the key. And that meaning consists of the change the message effects in the arrangement of the yard, the selection it makes from the set of all possible configurations. (Notice that this example is somewhat less behavioristic than that of the decision-making farmer, in that the reception of a meaningful message does not, in this case, imply action. Although the selection operation implies ***potential*** action—the shunting of a train in one direction instead of another—it is meaningful even in the absence of action; it is a property of the system itself.) This view of information and meaning seems to me exceedingly suggestive and points to some implications we can generalize without, I think, much embarrassment.

The first implication is that meaning must always be defined in terms of the system—or person—receiving the signal. Meaninglessness, MacKay notes, "is a relative concept, and a precise definition of meaning would be useless unless it automatically reminded us of this." (MacKay, 1969, p. 86.) The meaning of a message to a cybernetic system—a robot, a spacecraft, an economic decision-system—depends on the system's structure and the message's form; and, while the human system is of a very different order, it remains that meaning cannot be defined independent of that system—the apprehending mind. "Meaning is always meaning *to* ***someone.***" (MacKay, 1969, p. 36; emphasis original.)[15] This may sound reasonable, but it is not an entirely noncontroversial view, for it stands in opposition to numerous attempts to define meaning entirely in extrapersonal empirical and logical terms.[16]

A correlative implication is that—perhaps surprisingly—this view tends to blur rather than sharpen the distinction between semantic and nonsemantic information. If we speak of information as involving a selection operation on the states of readiness of a system, then we can speak of the selective information-content of a signal as somehow measuring the extent or mag-

[15] And here, in an important sense, is where methodological individualism comes back into the picture.

[16] This is true of Wittgenstein in the *Tractatus logico-philosophicus* and also of the later logical empiricists. [Wittgenstein, 1922; Carnap, 1950; Bar-Hillel and Carnap, 1953b; Bar-Hillel, 1964.] I should also note that it was this attempt to eliminate the personal and (contra Mortazavian in this volume) *not* the use of subjective probability that was the ultimate problem with the positivist approach. It was a problem of too *little* subjectivism, not too much.

nitude of the selection operation performed. The expected value of information developed in the agricultural example is precisely this sort of measure. How does it differ from the entropy measure of communications theory?

> The communications engineer measures the selective information-content of . . . signals, not in terms of the selective operation performed by the symbol on the ensemble of states of readiness of the . . . receiver, but in terms of the selective operation performed by the signal on the ensemble *of signals.* The symbols are represented in this ensemble in the proportions in which they normally occur, the most frequently used occupying the largest space and being most easily selected. (MacKay, 1969, p. 75; emphasis added.)

Thus, the entropy measure of communication theory is not so much a measure of nonsemantic information content as it is a measure of the semantic (i.e., selective) information content of operations on one particular system, albeit a system different from the one actually receiving the signal; it is a value-of-information measure for a system in which value and probability (or, rather, improbability) are identical.[17]

Another, more significant, aspect of this view—which we might call the system-referential view of information—is, it seems to me, that it argues strongly against the oil-flow model. Information is not homogeneous; meaning is a matter of form not of amount; and the value or significance of a message depends as much on the preexisting form of the receiver as on the message itself. More to the point, this view suggests that information is stored as knowledge in a system not as oil is stored in a tank, but by virtue of the change that information makes in the very organization of the system itself. In a fundamental sense, knowledge and organization are identical.

I will postpone some of the more resonant implications of this last, rather provocative, statement. For the moment, let me try to illustrate what such a connection between knowledge and organization might mean.

We are all accustomed to thinking of memory as a function that takes place in some isolated part of a system—a computer memory bank or a human brain, for instance. But information is not stored in those places like relics in an attic; information submitted to a system's memory changes the organization of that subsystem: The arrangement of magnetic elements in a

[17]We can see this clearly by comparing the equations in footnotes 9 and 10. For the value-of-information measure, we had

$$\text{EVPI} = p\,\Delta\Pi(w = H) + (1 - p)\Delta\Pi(w = L),$$

where $\Delta\Pi$ is shorthand for the expression in brackets in footnote 9. For the entropy measure, we had

$$H = -[p\log_2 p + (1 - p)\log_2(1 - p)].$$

In the first case, the probabilities weight terms that measure the effect of a message on the system; in the second case, the probabilities weight terms referring only to the selection of a message from the set of messages.

core memory, for example, is changed when data are stored. Furthermore, the memory bank is not the only locus of memory in the system. In an important sense, the entire organization of the system—hardware plus software, mind and body—contains functional knowledge to guide behavior. I do not think this is a new or particularly controversial way of putting things.

At the risk of a charge of idiosyncrasy, let me go beyond this to suggest a distinction between structural information and parametric information. The former is information that operates on—that changes—the basic structure of the system; the latter is information that operates on parameters of the system—on elements that adjust or calibrate the workings of the system within the dictates of, but without altering, its underlying structure.

I apologize if this all sounds a little vague. But the nature of form and structure is a problem that has animated philosophy since the Presocratics, and I make no pretense of trying to solve it. At an intuitive level, the notion of a system's structure is fairly clear: A system has various fixed (or relatively slow changing) attributes that define its form, that set the scope, ground rules, and boundary conditions for the system's more variable aspects. (This may have a physical correlative in hardware as distinct from software, but it need not; even within the software of a computer program, the mental representations of a human mind, or indeed any system in the abstract, we can speak of an underlying structure.) If we restrict ourselves to the mathematical realm, we can make the notions of structural and parametric information more precise, if in the end perhaps no less metaphoric.[18]

The agricultural decision-system again provides a clear example. Here, the model's structure lies in the form of the profit function; that function specifies a parameter, weather conditions, that can be altered by reception of an appropriate message. Our decisionmaker is able to obtain only parametric information and thus to gain only parametric knowledge from a signal. As the problem is formulated, no signal can change his profit function or any of the basic givens of the situation he faces.

In the modern mathematical "economics of information" [Hirshleifer and Riley, 19791, the focus is exclusively on parametric information of this sort. For a long time, the mainstream of economics had concentrated on "perfect-information" models, where the decisionmakers were portrayed as having full knowledge of all aspects of the decision-situations they faced. In newer

[18] Mesarović offers a definition of systems structure that, as best I understand it, is consistent with what I have in mind. The organization of a system consists in the systems relation that maps one set of parts into another. This relation $\boldsymbol{R}$ "can be considered as defined by an abstract relation and the specific values of the unspecified constituents, the so-called relational constituents; $\boldsymbol{R} = \{T, \zeta\}$, where T = systems structure, xi = set of relational constituents." (Mesarović, 1964a, p. 10.) As Mesarović further suggests (in Equation **9)**, these relational constituents can take the form of parameters of the system. Thus, my distinction between structural and parametric seems in accord with his view. Also see Mortazavian's paper in this volume.

models, the decisionmakers retain full structural knowledge of the problems they face; but now there are certain key parameters—such as weather conditions—obscured from their vision. As we might guess, the myopic decisionmakers invariably choose less efficiently than their better informed counterparts, a situation that has proven a boon to the legions of modern economists who take pleasure in identifying causes of what they term market failure much as nineteenth-century paleontologists delighted in unearthing new species in fossil. Meanwhile, the importance of imperfect ***structural*** knowledge—with its very different economic implications—has escaped widespread attention.

KNOWLEDGE, STRUCTURE, AND ECONOMICS

It is fair to conclude, I think, that systems theory and the theory of knowledge and information (broadly defined) must ultimately be related in a fundamental way. Both are concerned with form and organization. And it is little wonder that communications theory, with its entropylike measure of information content, has held a singular fascination for systems theorists. An organized entity is a nonrandom entity, one whose organization is unexpected in some sense, and it is unexpectedness—negative entropy—that communications theory measures.

Although, as we saw, the derivation of a scalar measure of organization is possible only in the very special case with which communications theory is concerned, it nonetheless remains that the knowledge content of a system is closely bound up with that system's organization—with its structure.

We have already seen two logics of organization identified by systems theorists: causal systems and goal-directed systems. I would now like to suggest that these two are not the whole story.

A causal system is an instance of what we might call a "mechanical" structure. The movements of the parts are causally related to one another within the dictates of a fixed structure. The centuries-old example is the mechanical clockwork. Each gear moves by virtue of the force impressed on it by a previous gear; and each carries out its function within the pattern ordained by the designer. Once cut loose from its creator, the mechanical system cannot increase its level of organization. Indeed, any change in structure (other than those effected by the ministrations of the designer) must lower the level of system organization, thus increasing entropy in some sense.

Information and control are closely related concepts in systems theory. In a strictly causal system, the only way to change behavior is by reprogramming the system; often, this can be accomplished by adjusting various control variables to modify the system's structure—much as we manipulate the steering wheel and foot pedals to alter the behavior of an automobile. This is called open-loop control, in that the controlling information comes entirely from outside the boundaries of the system.

The mechanical-control model has not been absent from economics. The various models of stabilization policy, now happily somewhat out of fashion, have long portrayed the economy as a system of aggregate variables (national income, consumption, investment, etc.), which the government-viewed as an entirely exogenous entity—could regulate by suitable manipulation of its control variables, for example, government expenditures.

As we have seen, goal-directed systems (or cybernetic systems, in the narrow sense of the term) work somewhat differently. These systems also have a fixed structure, but now there are certain manipulable variables that can be altered by information generated within the system itself. The system possesses a goal; and, by means of an information-feedback loop, it is able to compare its situation with that goal and make appropriate adjustments toward it. Among control theorists, this is called closed-loop control.[19]

I have already suggested that there is a definite congruence between causal systems and goal-directed ones, and while perhaps less horological than the former, the latter are not necessarily less mechanical in the everyday sense of the term. (An ordinary mechanical thermostat is a cybernetic device.) But there remains a tendency to see cybernetic systems as somehow more organic than causal systems. This is, I believe, largely because many bodily systems also operate in a cybernetic fashion.[20]

Rather than settling at a mechanical equilibrium point—like a clock running down or gas molecules coming to terms with the surrounding temperature—a cybernetic system achieves a condition of dynamic balance called homeostasis. Like a mechanical system, though, a cybernetic system in homeostasis is struggling to maintain a fixed level of organization. Any change in the system's state is, therefore, potentially dangerous, and prolonged changes represent worrisome imbalances that threaten to destroy the system. Indeed, many who take cybernetics seriously as a general metaphysics of order soon develop an attitude I tend to call thermodynamic Manichaeism—the implicit belief that every trend, every economic or social

[19] I should also mention here the branch of systems theory called optimal-control theory. The optimal-control theorist seeks the best way to control a system in order to achieve some goal. For example, we might wish to calculate that trajectory of a rocket between two points that minimizes flight time, fuel consumption, or some other objective. (And, as a matter of fact, it was precisely such aerospace problems that formed much of the early subject matter of this theory.) Mathematically, the optimal path is found using the calculus of variations (and something called the Pontryagin Maximum Principle) or the (closely related) technique of dynamic programming. Once found, this optimal trajectory can be imposed directly by open-loop control or implemented through a feedback system to create goal direction and closed-loop control. [See generally Bryson and Ho, 1968.] Optimal-control theory has found its way into economics, especially in economic-growth models (not unrelated to the simple stabilization policy models already mentioned) [Burmeister and Dobell, 1970, chapter 11] and in the theory of resource extraction [Sweeney, 1977]. See also Kamien and Schwartz [1981]] and Aoki [1976].

[20] Human biology was, in fact, an early inspiration for cybernetics. [Cannon, 1939.]

innovation is a potential victory for the forces of entropy and disorder over the forces of homeostasis. [Langlois, 1981, chap. 2, esp. pp. 73-82.] And it is thus probably no surprise that systems theorists [Mesarović and Pestel, 1974] and cybernetically oriented biologists [Hardin, 1977] are prominent among the new Malthusians who see dangerous social and ecological imbalances in the world's future.

In economics, there has been some effort afoot, especially by so-called post-Keynesians, to close the loop on the macroaggregate models of economic stabilization. These economists, as one observer correctly notes, ". . . might wish to replace the Newtonian-clockwork model by something they call a 'cybernetic' model, which may be an improvement (if it could ever be devised), but a shift from mechanical statics to sophisticated mechanical dynamics is no radical conceptual revolution." (Kristol, 1981, p. 212.)

What is the alternative to the mechanical causal and cybernetic models? The answer is best found by leaving the level of aggregate dynamics and returning to consideration of the parts—the human agents. A human being-at least in part or at times—is a goal-directed system. This is the basis of the ideal type—homo economicus—that underlies much of mainstream economic modeling. Yet everyone this side of B. F. Skinner recognizes that the flesh-and-blood human being does not operate with the stimulus-response compulsiveness of a simple goal-directed cybernetic system. There is something more, or perhaps different, at work.

One major school of thought holds, at least implicitly, that differences between the human mind and a mechanical cybernetic system are ones of degree rather than of kind. The mind is an immensely complex thing, and if we could only construct a cybernetic system—that is to say, a digital computer, the apotheosis of the mechanical cybernetic system—with sufficient complexity, we could largely replicate much of what the mind can do. Adherents to this view take heart from the old saying attributed to Marx that quantitative change allowed to go on long enough inevitably becomes qualitative change.

There are dissenters from this viewpoint, of course. Most notable among these is Hubert Dreyfus, whose analysis continues to stir controversy in the field of artificial intelligence. [Dreyfus, 1979.] In a significant sense, Dreyfus's thesis rests on the system-referential view of knowledge already articulated. Human knowledge, he argues, is very much tied to the biologically and culturally evolved structure of the human organism. What is meaningful to a human is meaningful only in reference to that structure, and meaning cannot be reduced to the system of explicit, extrapersonal, context-free statements that a computer, by virtue of its own structure, must employ. It is for this reason, Dreyfus argues, that no artificial intelligence program has yet been (or could be) written in which the important functions of discriminating meaning and significance are not preprogrammed by the human designer.

However this be resolved, it is clear that the human mind operates differently from *simple* cybernetic systems. The human mind is an example of a system in which information can result not merely in parametric but in structural knowledge. Unlike our simple cybernetic decisionmaker, the flesh-and-blood human is able, both spontaneously and as the result of signals, to change the problem formulation, rearrange the structure of his or her expectations, to alter his or her states of readiness. In Kantian terms, we might say that the human system is able to "evolve categories of description beyond those built into its design . . . ." (MacKay, 1969, p. 55.)

The lexicon of systems theory has a couple of words that sound as if they were intended to articulate a conception of the metamechanistic or the metacybernetic; these are *open system* and *self-organizing system*. The first is of thermodynamic origin and refers to a system that is not isolated and thus is able to exchange matter or information with an outside environment. A self-organizing system—by one definition, at least [Mesarovic, 1964a, pp. 11–131—is a system that can change its structure in response to the environment.[21] Clearly, a human system is both open and self-organizing. But I am not entirely persuaded that these terms by themselves fully capture our intuitive sense of the distinction between a mechanistic and nonmechanistic (or metamechanistic) system. Openness and self-organization are certainly necessary but perhaps not sufficient conditions to qualify a system for the latter category. Indeed, these distinctions do not seem to rule out various sorts of robots or things like "perceptrons" and "heuristic programs." [Ibid.][22] Perhaps the best distinction, despite its greater generality, is that between a "morphostatic" and a "morphogenetic" system. [Buckley, 1967, pp. 58–59.]

Models in the mainstream of economics—including, as I have suggested, those in the mathematical economics of information—portray the economic agent as a passive cybernetic reactor, a morphostatic system that responds to changes in data according to the logic of the economic problem programmed into it by the economist. As Fritz Machlup has long and patiently explained, this approach is not a statement about human psychology but a perfectly reasonable and justifiable technique of analysis—particularly for the sorts of economic problems to which the basic tools of partial-equilibrium comparative statics are normally applied. [Machlup, 1967.] But even granting all this, there remains a wide range of problems that demand a more active, morphogenetic ideal-type of the economic agent—one who can acquire structural knowledge and change the economic problem he or she faces.[23]

[21] Ludwig von Bertalanffy's definition of this term seems a bit different and comes nearer to the distinction I am looking for. He describes a self-organizing system as one capable of "evolving from a less to a more differentiated state." (Bertalanffy, 1968, p. 68.)

[22] The reason for this, I believe, is that Mesarovic is willing to classify systems as self-organizing if they change their structure within the dictates of a fixed higher level structure; that is, if they "change their structure by using a relation from the set Ω_R." (Mesarovic, 1964a, p. 11.)

[23] Morphogenetic man has not been entirely neglected in economics. Under the title of entrepre-

If economic agents are cybernetic reactors, they can adjust the economic mechanism in light of changed circumstances, but they cannot thereby increase the organization (decrease the entropy) of that mechanism. And any imperfection in the agent's knowledge (that is, any lack of correspondence between the economic problem perceived by the agents and the "true" economic problem) can lead to a bad equilibrium, a market failure. But if economic agents can alter the problem they face, if they can bring new ***structural*** knowledge into the system, then elaboration of, and increased differentiation in, the economic system becomes possible. Far from causing disorder or chaos, apparent departures from homeostatic equilibrium can actually result in an ***increase*** in system organization and a ***decrease*** in entropy. This is the phenomenon of "spontaneous order" [Hayek, 1967, p. 77] or, in technical jargon, "deviation-amplifying mutual causal effects." [Maruyama, 1963.][24]

Recognizing the existence and importance of morphogenetic processes has a number of far-reaching implications for both systems theory and economics, but this is not the place to explore them, I am afraid. Instead, let me close with some brief observations on relations between these two disciplines.

Systems theorists of my acquaintance are sometimes inclined to the opinion that economics lacks an adequate systems perspective and its scientific development would be rapidly enhanced by a complete subsumption of that discipline into systems theory. This viewpoint is not without its ironies. In an important sense, it was economists who were the first systems theorists. Moreover, the founders of economics were concerned with a fully morphogenetic version of systems theory. [Hayek, 1967.] Adam Smith's conception of economic growth based on the increasing division of labor is very much a theory about the evolution of economic structure from a less to a more differentiated state. [A. Smith, 1776 and 1936.] And, indeed, it is now widely recognized that the theory of evolution, known to us through the biological theory of Darwin, was articulated at least a century earlier in the social sciences.[25]

Of course, modern systems theory is also concerned with questions of morphogenesis, but mathematical systems theory—in the guise of non-equilibrium thermodynamics [Prigogine, 1971] and topology [Thom,

neur, he has been studied by Joseph Schumpeter and more recently, Israel Kirzner, among others. [Schumpeter, 1934; Kirzner, 1973.1

[24] Another name for this class of phenomena is autopoiesis, an area that is apparently attracting increasing interest among systems theorists. For a bibliographic introduction, see Zeleny [1981].

[25] David Hume, who numbered economics and social thought in general among his philosophical interests, adumbrated a genuinely Darwinian view of evolution in his ***Dialogues Concerning Natural Religion,*** first published in 1779 but written in 1759, a hundred years before ***The Origin of Species.*** [Hume, 1961, especially p. 478.]

1975][26]—is just beginning to grapple with these problems. In practice, the systems theorists who wish to bring economics under their wing are the control theorists and cyberneticians (in the narrow sense), the builders of naive-holist macromodels of the economy. In this sense, then, a subsumption of economics into systems theory is a step backwards—a step back in the direction of the French physiocrats, whose ideas Adam Smith transcended in founding the discipline of economics.[27] The problem is not a lack of high-powered systems mathematics in economics: Modern economists have not been slow in the least to adopt whatever new mathematical techniques emerge from systems theory. Indeed, if there is any need to make economics more system-theoretical, it is a need not to give it over to the cybernetic modelers but to return the discipline to the more sophisticated version of systems thought on which it was founded.

[26] This work by Thom, which presents the notion of catastrophe theory, has attracted a large following in a number of disciplines and remains something of a cause célèbre in mathematical systems theory. It promises mathematical interpretations of some of the most ancient and vexing problems of philosophy—the nature and origin of form and structure—and the text alternates enticing intimations of the profound with an utterly impenetrable formalism, thereby applying a time-tested formula for attracting a cult following.

[27] On the similarities between physiocrats and certain modern systems models of the economy, see Almarin Phillips [1955].