ON INFORMATION AND COMPLEXITY

Robert Rosen

April 1985
CP-85-19

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
2361 Laxenburg, Austria

FOREWORD

This paper represents the written version of a lecture given at IIASA in September 1984 under the auspices of the Science & Technology and the Regional Issues projects. In its current form it will appear as a chapter in the forthcoming IIASA book, *Complexity, Language and Life: Mathematical Approaches*, J. Casti and A. Karlqvist, eds.

Boris Segerstahl
Leader
Science & Technology Program

ABSTRACT

Arguing by analogy with Aristotle's four distinct categories of causation (material, formal, efficient and final), this paper argues that there are correspondingly distinct categories of information, and that the same mathematical language cannot be used to describe each of them. This fact leads to the conclusion that our mathematical language is somehow deficient, and that it must be supplemented by new structures. These considerations lead to a formalization of the ideas of a complex system and anticipatory control.

AFFILIATION

Robert Rosen is a Professor in the Department of Physiology and Biophysics at Dalhousie University, Halifax, Nova Scotia, Canada.

On Information and Complexity

Robert Rosen

Introduction

We introduce the rather wide-ranging considerations which follow with a discussion of the concept of *information* and its role in scientific discourse. Ever since Shannon began to talk of information theory (by which he meant a probabilistic analysis of the deleterious effects of propagating signals through channels; cf. Shannon and Weaver, 1949), the concept has been relentlessly analyzed and reanalyzed. The time and effort expended on these analyses must surely rank as one of the most unprofitable investments in modern scientific history; not only has there been no profit, but also the currency itself has been debased to worthlessness. Yet, in biology, for example, the terminology of information intrudes itself insistently at every level; code, signal, computation, recognition. It may be that these informational terms are simply not scientific at all; that they are a temporary anthropomorphic expedient; a *facon de parler* which merely reflects the immaturity of biology as a science, to be replaced at the earliest opportunity by the more rigorous terminology of force, energy, and potential which are the province of more mature sciences (i.e. physics), in which information is never mentioned. Or, it may be that the informational terminology which seems to force itself upon us bespeaks something fundamental; something that is missing from physics as we now understand it. We take this latter viewpoint, and see where it leads us.

In human terms, information is easy to define; it is anything that is or can be the answer to a question. Therefore, we preface our more formal considerations with a brief discussion of the status of interrogatives, in logic and in science.

The amazing fact is that interrogation is not ever a part of formal logic, including mathematics. The symbol "?" is not a logical symbol, as, for instance, are "$\vee$", "$\wedge$", "$\exists$", or "$\forall$"; nor is it a mathematical symbol. It belongs entirely to informal discourse and, as far as I know, the purely logical or formal character of interrogation has not been investigated. Thus, if information is indeed connected in an intimate fashion with interrogation, it is not surprising that it has not been formally characterized in any real sense. There is simply no existing basis on which to do so.

I do not intend to go deeply here into the problem of extending formal logic (always including mathematics in this domain) so as to include interrogatives. What I want to suggest here is a relation between our informal notions of interrogation and the familiar logical operation "$\Longrightarrow$"; the conditional, or the implication, operation. Colloquially, this operation can be rendered in the form "If A, then B". My argument involves two steps. First, that ***every*** interrogative can be put into a kind of conditional form:

If A, then B ?

(where B can be an indefinite pronoun like who, what, etc., as well as a definite proposition); and second, and most important, that every interrogative can be expressed in a more special conditional form, which can be described as follows. Suppose I know that some proposition of the form

If A, then B

is true. Suppose I now change or vary A; that is, replace A by a new expression, δA. The result is an interrogative, which I can express as

If δA, then δB ?

Roughly, I am treating the true proposition "If A, then B", as a reference, and I am asking what happens to this proposition if I replace the reference expression A by the new expression δA. I could, of course, do the same thing with B in the reference proposition; replace it by a new proposition δB and ask what happens to A. I assert that every interrogative can be expressed this way, in what I call a ***variational form***.

The importance of these notions for us lies in their relation to the external world; most particularly in their relation to the concept of ***measurement***, and to the notions of causality to which they become connected when a formal or logical system is employed to represent what is happening in the external world; that is, to describe some physical or biological system or situation.

Before discussing this, I want to motivate the two assertions made above, regarding the expression of arbitrary interrogatives in a kind of conditional form. I do this by considering a few typical examples, and leave the rest to the reader for the moment.

Suppose I consider the question

"Did it rain yesterday?"

First, I write it as

"If (yesterday), then (rain)?"

which is the first kind of conditional form described above. To find the variational form, I presume I know that some proposition like

"If (today), then (sunny)"

is true. The general variational form of this proposition is

"If δ(today), then δ(sunny)?"

Then, if I put

δ(today) = (yesterday),

δ(sunny) = (rain)

I have, indeed, expressed my original question in the variational form. A little experimentation with interrogatives of various kinds taken from informal discourse (of great interest are questions of classification, including existence and universality) should serve to make manifest the generality of the relation between interrogation and the implicative forms described above; of course, this cannot be *proved* in any logical sense since, as noted above, interrogation remains outside logic.

It is clear that the notions of observation and experiment are closely related to the concept of interrogation. That is why the results of observation and experiment (i.e. data) are so generally regarded as being information. In a formal sense, simple observation can be regarded as a special case of experimentation; intuitively, an observer simply determines what *is*, while an experimenter systematically perturbs what is, and then observes the effects of his or her perturbation. In the conditional form, an observer is asking a question which can generally be expressed as

"If (initial conditions), then (meter readings)?"

In the variational form, this question may be formulated as follows: assuming the proposition

"If (initial conditions = 0), then (meter readings = 0)"

is true (this establishes the reference, and corresponds to calibrating the meters), we ask

"If δ(initial conditions = 0), then δ(meter readings = 0)?"

where, simply

δ(initial conditions = 0) = (initial conditions)

and

δ(meter readings = 0) = (meter readings).

The experimentalist, essentially, takes the results of observation as the reference and asks, in variational form, simply

"If δ(initial conditions), then δ(meter readings)?"

The theoretical scientist, on the other hand, deals with a different class of question; namely, those that arise from assuming a δB (which may be B itself) and asking for the corresponding δA. These are questions that an experimentalist cannot approach directly, not even in principle. It is the difference between the two kinds of questions which distinguishes between experiment and theory, as well as the difference between the explanatory and predictive roles of theory itself; clearly, if we give δA and ask for the consequent δB, we are predicting, whereas if we assume δB and ask for the antecedent δA, we are explaining.

It should be noted that exactly the same duality arises in mathematics and logic themselves; that is, in purely formal systems. Thus, a mathematician can ask (*informally*): If (I make certain assumptions), then (what follows)? Or, the mathematician can start with a conjecture, and ask: If (Fermat's Last Theorem is true), then (what initial conditions must I assume to construct explicitly a proof)? The former is analogous to prediction, the latter to explanation.

When formal systems (i.e. logic and mathematics) are used to construct images of what occurs in the world, then interrogations and implications become associated with ideas of causality. Indeed, the whole concept of natural law depends precisely on the idea that causal processes in natural systems can be made to correspond with implication in some appropriate, descriptive inferential system (e.g. Rosen, 1984, where this theme is developed at great length).

But the concept of causality is itself a complicated one; a fact largely overlooked in modern scientific discourse, to its cost. That causality is complicated has already been pointed out by Aristotle, for whom all science was animated by a specific interrogative: Why? He said explicitly that the business of science was to concern itself with "the why of things". In our language, these are just the questions of *theoretical* science: If (B), then (what A)? and hence we can say B *because* A. Or, in the variational form, δB *because* δA.

However, Aristotle argued that there were four distinct categories of causation; four ways of answering the question *why*. These categories, which he called *material cause*, *formal cause*, *efficient cause*, and *final cause*, are not interchangeable. If this is so (and I argue below that, indeed, it is), then there are correspondingly *different kinds of information*, associated with different causal categories. These different kinds of information have been confused, mainly because we are in the habit of using the same mathematical language to describe each of them; it is from these inherent confusions that much of the ambiguity and murkiness of the concept of information ultimately arises. Indeed, we can say more than this: the very fact that the same mathematical language does not (in fact, cannot) distinguish between essentially distinct categories of causation means that the mathematical language we have been using is, in itself, somehow fundamentally deficient, and that it must be extended by means of supplementary structures to eliminate those deficiencies.

The Paradigm of Mechanics

The appearance of Newton's *Principia* toward the end of the seventeenth century was surely an epochal event. Though nominally the theory of physical systems of mass points, it was much more. In practical terms, by showing how the mysteries of the heavens could be understood on the basis of a few simple, universal laws, it set the standards for explanation and prediction which have been accepted ever since. It unleashed a feeling of optimism almost unimaginable today; it was the culmination of the entire Renaissance. More than that: in addition to providing a universal explanation for specific physical events, it also provided a language and a way of thinking about systems which has persisted, essentially unchanged, to the present time; what has changed has only been the technical manifestation of the language and its interpretation. In this language, the word information does not appear in any formal, technical sense; we have only words like energy, force, potential, work, and the like.

It is important to recognize the twin roles played by Newtonian mechanics in science: as a reductionistic ultimate and as a paradigm for representing systems not yet reduced to arrangements of interacting particles. The essential feature of this paradigm is the employment of a mathematical language with an inherent duality, which we may express as the distinction between *internal states* and *dynamical laws*. In Newtonian mechanics, the internal states are represented by points in some appropriate manifold of phases, and the dynamical laws represent the internal or impressed forces. The resulting mathematical image is thus what is called nowadays a ***dynamical system.*** However, the dynamical systems arising in mechanics are mathematically rather special ones, because of the way phases are defined (they possess a symplectic structure). Through the work of people like Poincaré, Birkhoff, Lotka, and many others over the years, however, this dynamical system paradigm, or its numerous variants, has come to be regarded as the universal vehicle for the representation of systems which could not, technically, be described in terms of mechanics; systems of interacting chemicals, organisms, ecosystems, and many others. Even the most radical changes occurring within physics itself, like relativity and quantum theory, manifest this framework; in quantum theory, for instance, there was the most fundamental modification of what constitutes a *state*, and how it is connected to what we can observe and measure; but otherwise, the basic partition between states and dynamical laws is relentlessly maintained. Roughly, this partition embodies a distinction between what is inside or intrinsic (the states) and what is outside (the dynamical laws, which are formal generalizations of the mechanical concept of impressed force).

This, then, is our inherited ***mechanical paradigm***, which in its many technical variants or interpretations has been regarded as a universal language for describing systems and their effects. The variants take many forms; automata theory, control theory, and the like, but they all conform to the same basic framework first exhibited in the *Principia*.

Among other things, this framework is regarded as epitomizing the concept of causality. We examine this closely here, because it is important when we consider the concept of information within this framework.

Mathematically, a dynamical system can be regarded simply as a vector field on a manifold of states; to each state, there is an assigned velocity vector (in mechanics it is, in fact, an acceleration vector). A given state (representing what the system is intrinsically like at an instant) together with its associated tangent vector (which represents what the effect of the external world on the system is like at an instant) uniquely determine how the system will change state, or move in time. This translation of environmental effects into a unique tangent vector is already a causal statement, in some sense; it translates into a more perspicuous form through a process of *integration*, which amounts to solving the equations of motion. More precisely, if a dynamical system is expressed in the familiar form

$$dx_i / dt = f_i(x_1, \ldots, x_n) \quad i = 1, \ldots, n \tag{7.1}$$

in which time does not generally appear as an explicit variable (but only implicitly through its differential or derivation, dt), the process of integration manifests the explicit dependence of the state variables $x_i = x_i(t)$ on time,

$$x_i(t) = \int_{t_0}^{t} f_i[x_1(\tau), \ldots, x_n(\tau)] d\tau + x_i(t_0) \ . \tag{7.2}$$

This is a more traditional kind of causal statement, in which the state at time t is treated as an *effect*, and the right-hand side of equation (7.2) contains the *causes* on which this effect depends.

Before going further, let us take a look at the integrands in equation (7.2), which are the velocities or rates of change of the state variables. The mathematical character of the entire system is determined solely by the *form* of these functions. Hence, we can ask: What is it that expresses this form (i.e. what determines whether our functions are polynomials, or exponentials, or of some other form)? And given the general form (polynomial, say), what is it that picks out a specific function and distinguishes it from all others of that form?

The answer, in a nutshell, is *parameters*. As I have written the system (7.1) above, no such parameters are explicitly visible, but they are at least tacit in the very writing of the symbol f_i. Mathematically, these parameters serve as coordinates for function spaces; just as any other coordinate, they label or identify the individual members of such spaces. They thus play a very different role to the state variables, which constitute the arguments or domains of the functions that they identify.

Here we find the first blurring. For the parameters which specify the form of the functions f_i can, *mathematically*, be thrown in as arguments of the functions f_i themselves; thus, we could (and in fact always do) write

$$f_i = f_i(x_1, \ldots, x_n, a_1, \ldots, a_r) \tag{7.3}$$

where a_i are *parameters*. We could even extend the dynamical equations (7.1) by writing $da_i/dt = 0$ (if the a_i are indeed independent of time); thus, mathematically we can entirely eradicate any distinction between the parameters and the state variables.

There is still one further distinction to be made. We pointed out above that the parameters a_i represent the effects of the outside world on the intrinsic system states. These effects involve *both* the system *and* the outside world. Thus, some of the parameters must be interpreted as intrinsic too (the so-called *constitutive* parameters), while others describe the state of the outside world. These latter obey their own laws, not incorporated in equation (7.1), so they are, from the standpoint of equation (7.1), simply regarded as *functions of time* and must be posited independently. They constitute what are variously called *inputs*, *controls*, or *forcings*. Indeed, if we regard the states $[x_i(t)]$, or any mathematical functions of them, as corresponding *outputs* (that is, output as a function of input rather than just of time) we pass directly to the world of control theory.

So let us review our position. Dividing the world into state variables plus dynamical laws amounts to dividing the world into state variables plus parameters, where the role of the parameters is to determine the *form* of the functions, which in turn define the dynamical laws. The state variables are the arguments of these functions, while the parameters are coordinates in function spaces. Further, we must partition the parameters themselves into two classes; those which are *intrinsic* (the constitutive parameters) and those which are *extrinsic*; that is, which reflect the nature of the environment. The intrinsic parameters are intuitively closely connected with the *system identity*; that is, with the specific nature or character of the system itself. The values they assume might, for example, tell us whether we are dealing with oxygen, carbon dioxide, or any other chemical *species*, and, therefore, cannot change without our perceiving that a *change of species* has occurred. The environmental parameters, as well as the state variables, however, can change without affecting the species of the system.

These distinctions cannot be accommodated with the simple language of vector fields on manifolds; that language is too abstract. We can only recapture these distinctions by (a) superimposing an informal layer of *interpretation* on the formal language, as we have done above, or (b) changing the language itself, to render it less abstract. Let us examine how this can be done.

In order to have names for the various concepts involved, I call the constitutive parameters, which specify the *forms* of the dynamical laws, and hence the species of system with which we are dealing, the system *genome*; the remaining parameters, which reflect the nature of the external world, I call the system *environment*, and the state variables themselves I call *phenotypes*. This rather provocative terminology is chosen to deliberately reflect corresponding biological situations; in particular, I have argued (cf. Rosen, 1978) that, viewed in this light, the genotype–phenotype dualism which is regarded as so characteristically biological has actually a far more universal currency.

The mathematical structure appropriate to reflect the distinctions we have made is that of genome-parameterized mappings from a space of environments to a space of phenotypes; that is, mappings of the form

$$f_g : E \to P$$

specified in such a way that given any initial phenotype, environment-plus-genome determines a corresponding trajectory. Thus, we have no longer a simple manifold of states, but rather a fiber-space structure in which the basic distinctions between genome, environment, and phenotype are embodied from the beginning. Some of the consequences of this scenario are examined in Rosen (1978, 1983); we cannot pause to explore them here.

Now we are in a position to discuss the actual relation between the Newtonian paradigm and the categories of causation described earlier. In brief, if we regard the phenotype of the system at time t as *effect*, then

(1) Initial phenotype *is* material cause.
(2) Genome g *is* formal cause.
(3) $f_g(\alpha)$, as an operator on the initial phenotype, *is* efficient cause.

Thus, the distinctions we have made between genome, environment, and phenotype are directly related to the old Aristotelian categories of causation. As we shall soon discover, that is why these distinctions are so important.

Note that one of the Aristotelian categories is missing from the above; there is no *final cause*. Ultimately, this is the reason why final cause has been banished from science; the Newtonian paradigm simply has no room for it. Indeed, it is evident that any attempt to superimpose a category of final causation upon the Newtonian world would effectively destroy the other categories within it.

In a deep sense, the Newtonian paradigm has led us to the notion that we may effectively *segregate the categories of causation* in our system descriptions. Indeed, the very concept of system state segregates the notion of material cause from other categories of causation, and tells us that it is correct to deal with all aspects of material causation independent of other categories: likewise with the concepts of genome and environment. I, in fact, claim that *this very segregation* *into independent categories of causation is the heart of the Newtonian paradigm.* When stated in this way, however, the universality of the paradigm perhaps no longer appears so self-evident.

Information

We said above that information is, or can be, the answer to a question, and that a question can generally be put in the variational form: If δA, then δB?. This serves as the connecting bridge between information and the Newtonian paradigm. In fact, it has played an essential role in the historical development of Newtonian mechanics and its variants, under the rubric of *virtual displacements*.

In mechanics, a virtual displacement is a small, imaginary change imposed on the *configuration* of a mechanical system, while the impressed forces are kept fixed. The animating question is: If such a virtural displacement is made under given circumstances, then what happens? The answer, in mechanics, is the well-known *Principle of Virtual Work*: if a mechanical system is in equilibrium, then the virtual work done by the impressed forces as a result of the virtual displacement must vanish. This is a static (equilibrium) principle, but it can readily be extended from statics to dynamics, where it is known as *D'Alembert's Principle*. In the dynamical case, it leads directly to the differential equations of motion of a mechanical system when the impressed forces are known. Details can be found in any text on classical mechanics.

In what follows, we explore the effect of such virtual displacements on the apparently more general class of dynamical systems of the form

$$dx_i / dt = f_i(x_1, \ldots, x_n) \quad i = 1, \ldots, n \ . \tag{7.4}$$

There is, however, a close relationship between the general dynamical systems (7.4) and those of Newtonian mechanics; indeed, the former systems can be regarded as arising out of the latter by the imposition of a sufficient number of nonholonomic constraints.[1]

As we have already noted, the language of dynamical systems, like that of Newtonian mechanics, does not include the word information; the study of such systems revolves around the various concepts of *stability*. However, in one of his analyses of oscillations in chemical systems, Higgins (1967) drew attention to the quantities

$$u_{ij}(x_1, \ldots, x_n) = \partial / \partial x_j (dx_i / dt) \ .$$

These quantities, which he called cross-couplings if $i \neq j$ and self-couplings if $i = j$, arise fundamentally from the conditions which govern the existence of oscillatory solutions to equations (7.4). It turns out that it is not so much the magnitudes as the signs of these quantities that are important. In order to have a convenient expression for the signs of these quantities, he proposed that we call the jth state variable, x_j, an *activator* of the ith, in the state $(x_1^0, \ldots, x_n^0)$, whenever the quantity

$$u_{ij}(x_1^0, \ldots, x_n^0) = \frac{\partial}{\partial x_j} \left(\frac{dx_i}{dt} \right)_{(x_1^0, \ldots, x_n^0)} > 0$$

and an *inhibitor* whenever

$$u_{ij}(x_1^0,\ldots,x_n^0) < 0 \ .$$

Now, activation and inhibition are *informational* terms. Thus, Higgins' terminology provides an initial hint as to how dynamical language might be related to informational language, through the Rosetta stone of stability.

Now let us examine what Higgins' terminology implies. If x_j activates x_i in a particular state, then a (virtual) increase in x_j increases the *rate of change* of x_i or, alternatively, a (virtual) decrease of x_j decreases the rate of change of x_i. It is, intuitively, eminently reasonable that this is the role of an activator. Conversely, if x_j inhibits x_i, it means that an increase in x_j decreases the rate of change of x_i, etc.

Thus, the n^2 functions, $u_{ij}(x_1,\ldots,x_n)$; i, $j = 1,\ldots,n$, constitute a form of informational description for the dynamical system (7.4), which I have elsewhere (Rosen, 1979) called an *activation–inhibition pattern*. As we have noted, such a pattern concisely represents the answers to the variational questions: If we make a virtual change in x_j, what happens to the rate of production of x_i?.

There is no reason to consider only the quantities u_{ij}. We can, for instance, go one step further, and consider the quantities

$$u_{ijk}(x_1,\ldots,x_n) = \partial / \partial x_k [\partial / \partial x_j (\mathrm{d}x_i / \mathrm{d}t)] \ .$$

Intuitively, these quantities measure the effect of a (virtual) change in x_k on the *extent* to which x_j activates or inhibits x_i. If such a quantity is positive in any particular state, it is reasonable to call x_k an *agonist* of x_j with respect to x_i; if negative, an *antagonist*. That is, if u_{ijk} is positive, a (virtual) increase in x_k increases or facilitates the activation of x_i by x_j, etc. The quantities u_{ijk} thus define another layer of informational interaction, which we may call an *agonist–antagonist pattern*.

We can iterate this process, in fact to infinity, to produce at each state r a family of n^r functions, $u_{ij\ldots r}(x_1,\ldots,x_n)$. Each layer in this increasing sequence describes how a (virtual) change of a variable at that level modulates the properties of the preceding level.

So far we have considered only the effects of virtual changes in state variables, x_j, on the velocities, $\mathrm{d}x_i / \mathrm{d}t$, at various informational levels. We could similarly consider the effects of virtual displacements at these various levels on the second derivatives, $\mathrm{d}^2 x_i / \mathrm{d}t^2$ (i.e. on the *accelerations* of x_i), the third derivatives $\mathrm{d}^3 x_i / \mathrm{d}t^3$, and so on. Thus, we have a doubly infinite web of informational interactions, defined by the functions

$$u^m_{ijk\ldots r}(x_1,\ldots,x_n) = \frac{\partial}{\partial x_r}\left(\cdots \frac{\partial}{\partial x_j}\left(\frac{\mathrm{d}x_i^m}{\mathrm{d}t^m}\right) \cdots \right)$$

If we start from the dynamical equations (7.4), then nothing new is learned from these circumlocutions beyond, perhaps, a deeper insight into the relations between dynamical and informational ideas. Indeed, given any layer of informational structure, we can proceed to succeeding layers by mere differentiation, and to antecedent layers by mere integration. Thus, knowledge of any layer in this infinite array of layers determines all of them and, in particular, the dynamical equations themselves. If we know, for instance, the activation–inhibition pattern

$u_{ij}(x_1,\dots,x_n)$, we can reconstruct the dynamical equations (7.4) through the relationship

$$df_i = \sum_{j=1}^{n} u_{ij}\,dx_j \tag{7.5}$$

(note in particular that the differential form on the right-hand side resembles a generalized *work*), and then set the function $f_i(x_1,\dots,x_n)$ so determined equal to the rate of change, dx_i/dt, of the ith state variable.

However, *our ability to do all this depends fundamentally on the exactness of the differential forms which arise at every level of our web of informational interaction*, and which relate each level to its neighbors. If the forms in equation (7.5) are not exact, there are no functions $f_i(x_1,\dots,x_n)$ whose differentials are given by it, and hence *no rate equations of the form (7.4)*. In such a situation, the simple relationship between the levels in our web breaks down completely; the levels become independent of each other, and must be posited separately. So two systems could have the same activation–inhibition patterns, but vastly different agonist–antagonist patterns, and hence manifest entirely different behaviors.

To establish firmly these ideas, let us examine what is implied by the requirement that the differential forms

$$\sum_{j=1}^{n} u_{ij}\,dx_j$$

defined by the activation–inhibition pattern be exact. The familiar, necessary conditions for exactness here take the form

$$\frac{\partial}{\partial x_k} u_{ij} = \frac{\partial}{\partial x_j} u_{ik}$$

for all $i, j, k = 1,\dots,n$. Intuitively, these conditions mean that *the relations of agonism and activation are entirely symmetrical* (commutative); that x_k as an agonist of the activator x_j is exactly the same as x_j as an agonist of the activator x_k; and similarly for all other levels.

Clearly, such situations are extremely degenerate in informational terms. They are so because the requirement of exactness is highly nongeneric for differential forms. Thus, these very simple considerations suggest a most radical conclusion: that *the Newtonian paradigm, with its emphasis on dynamical laws, restricts us from the outset to an extremely special class of systems, and that the most elementary informational considerations force us out of that class.* We explore some of the implications of this situation in the following section.

Meanwhile, let us consider some of the ramifications of these informational ideas that hold even within the confines of the Newtonian paradigm. These concern the distinctions made in the preceding section between environment, phenotype, and genome; the relations of these distinctions to different categories of causation; and the correspondingly different categories of information which these causal categories determine.

First, let us recall that according to the Newtonian paradigm, every relation between physical magnitudes (i.e. every equation of state) can be represented as a genome-parameterized family of mappings

$$f_g : E \to P$$

from environments to phenotypes. It is worth noting specifically that every dynamical law or equation of motion is of this form, as is shown by

$$\mathrm{d}\boldsymbol{x}/\mathrm{d}t = f_g(\boldsymbol{x}, \boldsymbol{\alpha}) \ . \tag{7.6}$$

Here, in traditional language, $\boldsymbol{x}$ is a vector of states, $\boldsymbol{\alpha}$ is a vector of external controls (which together with states constitutes *environment*), and the phenotype is the tangent vector $\mathrm{d}\boldsymbol{x}/\mathrm{d}t$ attached to the state $\boldsymbol{x}$.[2] In this case, then, the tangent vector or phenotype constitutes *effect*; the genome g is identified with formal cause, state $\boldsymbol{x}$ with material cause, and the operator $f_g(\ldots, \boldsymbol{\alpha})$ with efficient cause.

By analogy with the activation–inhibition networks and their associated informational structures, described above, we can consider formal quantities of the form

$$\frac{\partial}{\partial(\text{cause})}\left[\frac{\mathrm{d}}{\mathrm{d}t}(\text{effect})\right] \tag{7.7}$$

As always, such a formal quantity represents an answer to a question: If (cause is varied), then (what happens to effect)? This is the same question as we asked in connection with the definition of activation–inhibition networks and their correlates, but now set in the wider context to which our analysis of the Newtonian paradigm has led us. That is, we may now virtually displace *any* magnitude which affects the relation (7.6), whether it be a genomic magnitude, an environmental magnitude, or a state variable. In a precise sense, the effect of such a virtual displacement is measured by the quantity (7.7).

It follows that there are indeed different *kinds* of information. What kind of information we are dealing with depends on whether we apply the virtual displacement to a genomic magnitude (associated with formal cause), an environmental magnitude (efficient cause), or a state variable (material cause). Formally, we can now distinguish at least the following three cases:

(1) Genomic information,

$$\frac{\partial}{\partial(\text{genome})}\left[\frac{\mathrm{d}}{\mathrm{d}t}(\text{effect})\right] .$$

(2) Phenotypic information,

$$\frac{\partial}{\partial(\text{state})}\left[\frac{\mathrm{d}}{\mathrm{d}t}(\text{effect})\right] .$$

(3) Environmental information,

$$\frac{\partial}{\partial(\text{control})}\left[\frac{\mathrm{d}}{\mathrm{d}t}(\text{effect})\right] .$$

We confine ourselves herein to these three, which generalize only the activation–inhibition patterns described above.

We now examine an important idea; namely, *the three categories defined above are not equivalent*. Before justifying this assertion, we must briefly discuss what is meant by equivalent. In general, the mathematical assessment of the effects of perturbations (i.e. of real or virtual displacements) is the province of *stability*. For example, the effect on subsequent dynamical behavior of modifying or perturbing a system state is the province of Lyapunov stability of dynamical systems; that of perturbing a control is part of control theory; and that of perturbing a genome relates to structural stability. To establish this firmly, let us consider genomic perturbations, or *mutations*. A virtual displacement applied to a genome g replaces the initial mapping f_g determined by g with a new mapping $f_{g'}$. Mathematically, we say that the two mappings, f_g and $f_{g'}$, are equivalent, or similar, or conjugate, if there exist appropriate transformations

$$\alpha : E \rightarrow E \ ,$$

$$\beta : P \rightarrow P \ ,$$

such that the diagram

commutes; that is, if

$$\beta[f_g(e)] = f_{g'}[\alpha(e)]$$

for every e in E. Intuitively, this means that a mutation $g \mapsto g'$ can be counterbalanced, or nullified, by imposing suitable *coordinate transformations* on the environments and phenotypes. Stated yet another way, a virtual displacement of genome can always be counteracted by corresponding displacements of environment and phenotype so that the resultant variation on effect vanishes.

We have elsewhere (Rosen, 1978) shown at great length that this commutativity may not always obtain; that is, that there may exist genomes which are bifurcation points. In any neighborhood of a bifurcating genome g, there exist genomes g' for which f_g and $f_{g'}$ fail to be conjugate.

With this background, we return to the question of whether the three kinds of information (genomic, phenotypic, and environmental) defined above are equivalent. Intuitively, equivalence would mean that the effect of a virtual displacement δg of genome, supposing all else is fixed, could equally well be produced by a virtual displacement of environment, δa, or of phenotype, δp. Or stated another way, the effect of a virtual displacement δg of genome can be nullified by virtual displacements $-\delta a$ and $-\delta p$ of environment and phenotype, respectively. This is simply a restatement of the definition of conjugacy or similarity of mappings.

If all forms of information are equivalent, it follows that there could be no bifurcating genomes. We note in passing that the assumption of equivalence of the three kinds of information defined above thus creates terrible ambiguities when it comes to *explanation* of particular effects. We do not consider that aspect here, except to say that it is perhaps very fortunate that, as we have seen, they are not equivalent.

Let us examine one immediate consequence of the nonequivalence of genomic, environmental, and phenotypic information, and of the considerations which culminate in that conclusion. Long ago (cf. von Neumann, 1951; Burks, 1966) von Neumann proposed an influential model for a self-reproducing automaton, and subsequently, for automata which grow and develop. This model was based on the famous theorem of Turing (1936), which established the existence of a universal computer (universal Turing machine). From the existence of such a universal computer, von Neumann asserted that there must also exist a universal constructor. Basically, he argued that computation (i.e. following a program) and construction (following a blueprint) are both algorithmic processes, and that anything holding for one class of algorithmic processes necessarily holds for any other class. This universal constructor formed the central ingredient of the self-reproducing automaton.

Now, a computer acts, in the language we have developed above, through the manipulation of efficient cause. A constructor, if the term is to bear any resemblance to its intuitive meaning, must essentially manipulate material cause. The inequivalence of the two categories of causality, in particular manifested by the nonequivalence of environmental and phenotypic information, means that we cannot blithely extrapolate from results pertaining to efficient causation into the realm of material causation. Indeed, in addition to invalidating von Neumann's specific argument, we learn that great care must be exercised in general when arguing from purely logical models (i.e. from models pertaining to efficient cause) to any kind of physical realization, such as developmental or evolutionary biology (which pertain to material cause).

Thus, we realize how significant are the impacts of informational ideas, even within the confines of the Newtonian paradigm, in which the categories of causation are essentially segregated into separate packages. We now consider what happens when we vacate the comforting confines of the Newtonian paradigm.

An Introduction to Complex Systems

Herein, I call any natural system for which the Newtonian paradigm is completely valid a ***simple system***, or ***mechanism***. Accordingly, a ***complex system*** is one which, for one reason or another, resides outside this paradigm. We have already seen a hint of such systems in the preceding section; for example, systems whose activation–inhibition patterns u_{ij} do not give rise to exact differentials $\sum u_{ij}\,dx_j$. However, some further words of motivation must precede a conclusion that such systems are truly complex (i.e. reside fundamentally outside the Newtonian paradigm). We must also justify our very usage of the term complex in this context.

What I have been calling the Newtonian paradigm ultimately devolves upon ***the class of distinct mathematical descriptions*** which a system can have, and the relations which exist between these descriptions. As noted earlier, the basis of system description arising in this paradigm is the fundamental dualism between states and dynamical laws. Thus, the mathematical objects which can describe natural systems comprise a category which may be called general dynamical systems. In a formal sense, it appears that any mathematical object resides in this category, because the Newtonian partition between states and dynamical laws exactly parallels the partition between propositions and production rules (rules of inference) which presently characterize all logical systems and logical theories. However, we argue that, although this category of general dynamical systems is large, it is not everything, and, indeed, it is far from large enough.

The Newtonian paradigm asserts much more than simply that every image of a natural system must belong to a given category. It asserts certain relationships between such images. In particular (and this is the reductionistic content of the paradigm), it asserts that among these images there is the universal one, which effectively maps on all the others. Intuitively, this is the master description or ultimate description, in which every shred of physical reality has an exact mathematical counterpart; in category–theoretic terms, it is much like a free object (a generalization of the concept of free semigroup, free group, etc.).[3]

There is still more. The ingredients of this ultimate description, by their very nature, are themselves devoid of internal structure; their only changeable aspects are their relative positions and velocities. Given the forces acting between them, as Laplace noted long ago, everything that happens in the external world is in principle predictable and understandable. From this perspective, everything is determined; there are no mysteries, no surprises, no errors, no questions, and no information. This is as much true for quantum theory as for classical; only the nature of state description has changed. And it applies to everything, from atoms to organisms to galaxies.

How does this universal picture manifest itself in biology? First, from the standpoint of the physicist, biology is concerned with a rather small class of extremely special (indeed, inordinately special) systems. In the theoretical physicist's quest for general and universal laws, there is thus not much contact with organisms. As far as he or she is concerned, what makes organisms special is not that they transcend the physicist's paradigms, but rather that their specification within the paradigm requires a plethora of special constraints and conditions, which must be superimposed on the universal canons of system description and reduction. The determination of these special conditions is an empirical task; essentially someone else's business. But it is not doubted that the relationship between physics and biology is the relationship between the general and the particular.

The modern biologist, in general, avidly embraces this perspective.[4] Historically, biology has only recently caught up with the Newtonian revolution which swept the rest of natural philosophy in the seventeenth century. The three-century lag arose because biology has no analog of the solar system; no way to make immediate and meaningful contact with the Newtonian paradigm. Not until physics and chemistry had elaborated the technical means to probe microscopic properties of matter (including organic matter) was the idea of molecular biology even thinkable. And this did not happen until the 1930s.

At present, there is still no single inferential chain which links any important effect in physics to any important effect in biology. This is a fact; a datum; a piece of information. How are we to understand it? There are various possibilities. Kant, long ago, argued that organisms could only be properly understood in terms of final causes or intentionality; hence, from the outset he suggested that organisms fall completely outside the canons of Newtonian science, which are applicable to everything else. Indeed, the essential telic nature of organisms precluded even the possibility that a "Newton of the grassblade" would come along, and do for biology what Newton did for physics. Another possibility is the one we have already mentioned; we have simply not yet characterized all those special conditions which are necessary to bring biology fully within the scope of universal physical principles. Yet a third possibility has developed within biology itself, as a consequence of theories of evolution; it is that much of biology is the result of *accidents* which are *in principle* unpredictable and hence governed by no laws at all.[5] In this view biology is as much a branch of history as of science. At present, this last hypothesis lies in a sort of doublethink relation with reductionism; the two are quite inconsistent, but do allow modern biologists to enjoy the benefits of vitalism and mechanism together.

Yet a fourth view was expressed by Albert Einstein, who wrote in a letter to Leo Szilard: "One can best appreciate, from a study of living things, how primitive physics still is".

So, the present prevailing view in biology is that the Newtonian canons are indeed universal, and we are lacking only knowledge of the special conditions and constraints which distinguish organisms from other natural systems within those canons. One way of describing this with a single word is to assert that organisms are *complex*. This word is not well defined, but it does connote several things. One of these is that complexity is a system property, no different from any other property. Another is that the *degree* to which a system is complex can be specified by a number, or set of numbers. These numbers may be interpreted variously as the dimensionality of a state space, or the length of an algorithm, or as a cost in time or energy incurred in solving system equations.

On a more empirical level, however, complexity is recognized differently, and characterized differently. If a system surprises us, or does something we have not predicted, or responds in a way we have not anticipated; if it makes errors; if it exhibits emergence of unexpected novelties of behavior, we also say that the system is complex. In short, complex systems are those which behave counter-intuitively.

Sometimes, of course, surprising behavior is simply the result of incomplete characterization; we can then hunt for what is missing, and incorporate it into our system description. In this way, the planet Neptune was located from unexplained deviations of Uranus from its expected trajectory. But sometimes this is not the case; in the apparently analogous case of the anomalies of the trajectory of the planet Mercury, for instance, no amount of fiddling within the classical scenario succeeded and only a massive readjustment of the paradigm itself (via general relativity) availed.

From these few words of introduction, we can conclude that the identification of complexity with situations where the Newtonian paradigm fails is in accord with the intuitive connotation of the term, and is an alternative to regarding as complex any situation which merely is technically difficult within the paradigm.

Now let us see where information fits into these considerations. We recall that information is the actual or potential response to an interrogative, and that every interrogative can be put into the variational form: If δA, then δB? The Newtonian paradigm asserts, among other things, that the answers to such interrogatives follow from dynamical laws superimposed on manifolds of states. In their turn, these dynamical laws are special cases of *equations of state*, which link or relate the values of system observables. Indeed, the concept of an observable was the point of departure for our entire treatment of system description and representation (cf. Rosen, 1978); it was the connecting link between the world of natural phenomena and the entirely different world of formal systems which we use to describe and explain.

However, the considerations we have developed above suggest that this world is not enough. We require also a world of variations, increments, and differentials of observables. It is true that every linkage between observables implies a corresponding linkage between differentials, but as we have seen, the converse is not true. We are thus drawn to the notions that a differential relation is a generalized linkage and that a differential form is a type of generalized observable. A differential form which is not the differential of an observable is thus an entity which assumes no definite numerical value (as an observable does), but which can be incremented.

If we do think of differential forms as generalized observables, then we must correspondingly generalize the notion of equation of state. A generalized equation of state thus becomes a linkage or relation between ordinary observables and differentials or generalized observables. Such generalized equations of state are the vehicles which answer questions of our variational form: If δA, then δB?

But as we have repeatedly noted, such generalized equations of state do not usually follow from systems of dynamical equations, as they do in the Newtonian paradigm. Thus, we must find some alternative way of characterizing a system of this kind. Here is where the informational language introduced above comes to the fore. Let us recall, for instance, how we defined the activation–inhibition network. We found a family of functions u_{ij} (i.e. of observables) which could be thought of in the dynamical context as modulating the effect of an increment dx_j on that of another increment df_i. That is, the values of each observable, u_{ij}, measure precisely the extent of activation or inhibition which x_j exerts on the rate at which x_i is changing. .

In this language, a system falling outside the Newtonian paradigm (i.e. a complex system) can have an activation–inhibition pattern, just as a dynamical (i.e. simple) system does. Such patterns are still families of functions (observables), u_{ij}, and the pattern itself is manifested by the differential forms

$$\omega_i = \sum u_{ij} dx_j$$

But in this case, there is no global velocity observable, f_i, that can be interpreted as the rate of change of x_i; there is only a velocity *increment*. It should be noted explicitly that u_{ij}, which define the activation–inhibition pattern, need not be functions of x_i alone, or even functions of them at all. Thus, the differential forms which arise in this context are different from those with which mathematicians generally deal, and which can always be regarded as cross sections of the cotangent bundle of a definite manifold of states.

The next level of information is the agonist–antagonist pattern, u_{ijk}. In the category of dynamical systems, this is completely determined by the activation–inhibition pattern, and can be obtained from the latter by differentiation:

$$u_{ijk} = \frac{\partial}{\partial x_k} u_{ij} .$$

In our world of generalized observables and linkages, u_{ijk} are independent of u_{ij}, and must be posited separately; in other words, complex (non-Newtonian) systems can have identical activation–inhibition patterns, but quite different agonist–antagonist patterns.

Exactly the same considerations can also be applied to every subsequent layer of the informational hierarchy; each is now independent of the others, and so must be posited separately. Hence a complex system requires an *infinite* mathematical object for its description.

We cannot examine herein the mathematical details of the considerations sketched so briefly above. Suffice it to say that a complex system, defined by a hierarchy of informational levels of the type described, is quite a different object to a dynamical system. For one, it is quite clear that there is no such thing as a set of *states*, assignable to such a system once and for all. From this alone, we might expect that the nature of causality in such systems is vastly different to what it is in the Newtonian paradigm; we come to this in a moment.

The totality of mathematical structures of the type we have defined above forms a category. In this category the class of general dynamical systems constitutes a very small subcategory. We are suggesting that the former provides a suitable framework for the mathematical imaging of complex systems, while the latter, by definition, can only image simple systems or mechanisms. If these considerations are valid (and I believe they are), then the entire epistemology of our approach to natural systems is radically altered, and it is the basic notions of information which provide the natural ingredients.

There is, however, a profound relationship between the category of general dynamical (i.e. Newtonian) systems, and the larger category in which it is embedded. This can only be indicated here, but it is important indeed. Namely, there is a precise sense in which an informational hierarchy can be *approximated*, locally and temporarily, by a general dynamical system. With this notion of approximation there is an associated notion of *limit*, and hence of topology. Using these ideas, it can be shown that what we call the category of complex systems is the completion, or limiting set, of the category of simple (i.e. dynamical) systems.

The fact that complex systems can be approximated (albeit locally and temporarily) by simple ones is crucial. It explains precisely why the Newtonian paradigm has been so successful, and why, to this day, it represents the only effective procedure for dealing with system behavior. But in general, it is apparent that it can usually supply *only* approximations, and in the universe of complex systems this amounts to replacing a *complex* system with a *simple subsystem*. Some of the profound consequences are considered in detail in Rosen (1978).

This relationship between complex systems and simple ones is, by its very nature, without a reductionistic counterpart. Indeed, what we presently understand as physics is seen in this light as *the science of simple systems*. The relation between physics and biology is thus not at all the relation of general to particular; in fact, quite the contrary. It is not biology, but physics, which is too special. We can see from this perspective that biology and physics (i.e. contemporary physics) develop as two divergent branches from a *theory of complex systems*, which as yet can be glimpsed only very imperfectly.

The category of simple systems is, however, still the only one that we know how to use. But to study complex systems by means of approximating simple systems resembles the position of early cartographers, who were attempting to map a sphere while armed only with pieces of planes. Locally, and temporarily, they could do very well, but globally, the effects of the topology of the sphere become progressively important. So it is with complexity; over short times and only a few informational levels, we can always make do with a simple (i.e. dynamical) picture. Otherwise, we cannot; we must continually replace our approximating dynamics with others as the old ones fail. Hence another characteristic feature of complex systems; they appear to possess a multitude of partial dynamical descriptions, which cannot be combined into one single complete description. Indeed, in earlier work (Rosen, 1977), we took this as the defining feature of complexity.

I add a brief word about the status of causality in complex systems, and about the practical problem of determining the functions which specify their informational levels. Complex systems do not possess anything like a state set which is fixed once and for all. Also, the categories of causality become intertwined in a way which is not possible within the Newtonian paradigm. Intuitively, this follows from the independence of the infinite array of informational layers which constitutes the mathematical image of a complex system. Variation of any particular magnitude connected with such a system typically manifests itself independently in many of these layers, and thus reflects itself partly as material cause, partly as efficient cause, and even partly as formal cause in the resultant variation of other magnitudes. We feel that it is, at least for the most part, this involvement of magnitudes simultaneously in each of the causal categories which makes biological systems so refractory to the Newtonian paradigm.

Also, this intertwining of the categories of causation in complex systems makes the direct interpretation of experimental results of the form: If δA, then δB, extremely difficult. If we are correct so far, such an observational result as it stands is far too coarse to have any clear-cut meaning. In order to be meaningful, an experimental proposition of this form must isolate the effect of a variation δA on a single informational level, keeping the others clamped. As might be appreciated, this will in general not be an easy task. In other words, the experimental study of complex systems cannot be pursued with the same tools and ideas that are appropriate for simple systems.

One final conceptual remark is also needed. As mentioned earlier, the Newtonian paradigm has no room for the category of final causation. This category is closely linked to the notion of anticipation, which in turn is linked to the ability of systems to possess internal predictive models of themselves and their environments, which can be utilized for the control of present actions. We have argued at great length elsewhere (cf. Rosen, 1984) that anticipatory control is indeed a distinguishing feature of the organic world, and have described some of the unique features of such anticipatory systems. Herein we have shown that for a system to be anticipatory, it must be complex. Thus, our entire treatment of anticipatory systems becomes a corollary of complexity. In other words, complex systems can admit the category of final causation in a perfectly rigorous, scientifically acceptable way. Perhaps this alone is sufficient recompense for abandoning the comforting confines of the Newtonian paradigm, which has served so well over the centuries. It will continue to serve us well, provided we recognize its restrictions and limitations, as well as its strengths.

Notes

[1] Newton's original ***particle mechanics***, or ***vectorial mechanics***, is hard to apply to many practical problems, and was early on (through the work of people like Euler and Lagrange) transmuted into another form, generally called ***analytical mechanics***. This latter form is usually used to deal with extended matter (e.g. rigid bodies). In particle mechanics, the rigidity of a macroscopic body is a consequence of interparticle forces, which must be explicitly taken into account in describing the system. Thus, if there are N particles in the system (however large N may be) there is a phase space of $6N$ dimensions, and a set of dynamical equations which expresses for each particle the resultant of *all* forces experienced by that particle. In analytical mechanics, on the other hand, any rigid body can be completely described by giving only six configurational coordinates (e.g. the coordinates of the center of mass, and three angles of rotation about the center of mass), however many particles it contains. From the particulate approach the internal forces which generate rigidity are replaced by ***constraints***; supplementary conditions on the configuration space which must be identically satisfied. Thus, the passage from particle mechanics to analytical mechanics involves a partition of the forces in an extended system into two classes: (a) the internal or ***reactive*** forces, which hold the system together, and (b) the ***impressed*** forces, which push the system around. The former are represented in analytical mechanics by algebraic constraints, the latter by differential equations in the configuration variables (six for a rigid body).

A system in analytical mechanics may have additional constraints imposed upon it by specific circumstances; for example, a ball may roll on a table top. It was recognized long ago that these additional constraints (which, like all constraints, are regarded as expressing the operation of reactive forces) can be of two types, which were called by Hertz ***holonomic*** and ***nonholonomic***. Both kinds of constraints can be expressed locally, in infinitesimal form, as

$$\sum_{i=1}^{n} u_i(x_1,\dots,x_n)dx_i = 0$$

where $x_1,\dots,x_n$ are the configuration coordinates of the system. For a holonomic constraint, the above differential form is exact; that is, the differential of some global function $\varphi(x_1,\dots,x_n)$ is defined over the whole configuration space. Thus, the holonomic constraint translates into a global relation

$$\varphi(x_1,\dots,x_n) = \text{constant} \ .$$

This means that the configurational variables are no longer independent, and that one of them can be expressed as a function of the others. The constraint thus reduces the dimension of the configuration space by *one*, and therefore reduces the dimension of the phase space by *two*.

A nonholonomic constraint, on the other hand, does not allow us to eliminate a configurational variable in this fashion. However, since it represents a relation between the configuration variables and their differentials, it does allow us to eliminate a coordinate of ***velocity***, while leaving the dimension of the configuration space unaltered. That is, a nonholonomic constraint serves to eliminate one degree of freedom of the system. It thus also eliminates one dimension from the space of impressed forces which can be imposed on the system without violating the constraint.

Similarly, if we impose r independent nonholonomic constraints on our system, we (a) keep the original dimension of the configuration space; (b) eliminate r coordinates of velocity, and thus reduce the dimensionality of the phase space by r; and (c) similarly, reduce by r the dimensionality of the set of impressed forces which can be imposed on the system.

Let us express these facts mathematically. A nonholonomic constraint can be expressed locally in the general form

$$\varphi\left[x_1,\ldots,x_n,\frac{dx_1}{dt},\ldots,\frac{dx_n}{dt}\right]=0$$

which can (locally) be solved for one of the velocity coordinates (dx_1/dt, say). Thus, it can be written in the form

$$\frac{dx_1}{dt}=\psi\left[x_1,x_2,\ldots,x_n,\frac{dx_2}{dt},\ldots,\frac{dx_n}{dt}\right]$$
$$=\Psi(x_1,\mathbf{a})$$

where we have written $\boldsymbol{a}=(x_2,\ldots,dx_n/dt)$. [At this point the reader is invited to compare this relation with equation (7.6) in the main text.]

Likewise, if there are r nonholonomic constraints, these can be expressed locally by r equations

$$dx_i/dt=\psi_i(x_1,\ldots,x_r,\boldsymbol{a})\quad i=1,\ldots,r$$

where now $\boldsymbol{a}$ is the vector $(x_{r+1},\ldots,x_n,dx_{r+1}/dt,\ldots,dx_n/dt)$. These equations of constraint, which intuitively arise from the *reactive* forces holding the system together, now become more and more clearly the type of equations we always use to describe general dynamical or control systems.

Now what happens if $r=n$? In this case, the constraints leave us only *one degree of freedom; they determine a vector field on the configuration space*. There is in effect only one *impressed* force that can be imposed on such a system, and its only effect is to move the system; once moving, the motion is determined entirely by the *reactive* forces, and not by the *impressed* force. Mathematically, the situation is that of an autonomous dynamical system, whose manifold of states is the *configuration* space of the original mechanical system.

This relationship between dynamics and mechanics is quite different from the usual one, in which the manifold of states is thought of as generalizing the mechanical notion of *phase*, and the equations of motion as generalizing the *impressed* force. In the above interpretation, however, it is quite different; the manifold of states correspond now to mechanical *configurations*, and the equations of motion come from the *reactive* forces.

[2] The reader should be most careful not to confuse two kinds of propositions, which are equivalent mathematically but completely different epistemologically and causally. On the one hand, we have a statement like

$$dx/dt=f_g(\boldsymbol{x},\boldsymbol{a})\ .$$

This is a local proposition, linking a tangent vector or velocity $d\boldsymbol{x}/dt$ to a state $\boldsymbol{x}$, a genome g, and a control $\boldsymbol{a}$. Each of these quantities is derived from observables assuming definite numerical values at any instant of time, and it is *their values at a common instant* which are related by this proposition.

On the other hand, the integrated form of these dynamical relations is

$$\boldsymbol{x}(t)=\int_{t_0}^{t}f_g[x,\boldsymbol{a}(\tau)]d\tau\ .$$

This relationship involves time *explicitly* and links the values of observables *at one instant* with values (assumed by these and other observables) *at other instants*.

Each of these epistemically different propositions has its own causal structure. In the first, we treat the tangent vector dx/dt as effect and define its causal antecedents as we have done. In the integrated form, on the other hand, we take $\boldsymbol{x}(t)$ as effect and find a correspondingly different causal structure. In general, the mathematical or logical equivalence of two expressions of linkage or relationship in physical systems does not at all connote that their causal structures are identical. This is merely a manifestation of what was discussed earlier, that the mathematical language we use to represent physical reality has abstracted away the very basis on which such causal discriminations can be made.

[3] It should be recognized that this reductionistic part of the Newtonian paradigm can fail for purely mathematical reasons. If it should happen that there is no way to effectively map the master description onto some partial description, then this is enough to defeat a reductionistic approach to those system behaviors with which the partial description deals. This is quite a different matter from the one we are considering here, in which no Newtonian master description *exists*, and the program fails for *epistemological* reasons, rather than mathematical ones.

[4] This statement is not simply my subjective assessment. In 1970 there appeared a volume entitled *Biology and the Future of Man*, edited by Philip Handler (1970), then President of the National Academy of Sciences of the USA. The book went to great lengths to assure the reader that it spoke for biology as a science; that in it biologists spoke with essentially one voice. At the outset, it emphasized that the volume was not prepared as a (mere) academic exercise, but for serious pragmatic purposes:

> Some years ago, the Committee on Science and Public Policy of the National Academy of Sciences embarked on a series of 'surveys' of the scientific disciplines. Each survey was to commence with an appraisal of the 'state of the art'.... In addition, the survey was to assess the nature and strength of our national apparatus for continuing attack on those major problems, e.g., the numbers and types of laboratories, the number of scientists in the field, the number of students, the funds available and their sources, and the major equipment being utilized. Finally, each survey was to undertake a projection of future needs for the national support of the discipline in question to assure that our national effort in this regard is optimally productive....

To address these serious matters, the Academy proceeded as follows:

> .... Panels of distinguished scientists were assigned subjects.... Each panel was given a general charge... as follows:
>
> The prime task of each Panel is to provide a pithy summary of the status of the specific sub-field of science which has been assigned. This should be a clear statement of the prime scientific problems and the major questions currently confronting investigators in the field. Included should be an indication of the manner in which these problems are being attacked and how these approaches may change within the foreseeable future. What trends can be visualized for tomorrow? What lines of investigation are likely to subside? Which may be expected to advance and assume greater importance?... Are the questions themselves... likely to change significantly?.... Having stated the major questions and problems, how close are we to the answers? The sum of these discussions, panel by panel, should constitute the equivalent of a complete overview of the highlights of current understanding of the Life Sciences.

There were twenty-one such Panels established, spanning the complete gamut of biological sciences and the biotechnologies. The recruitment for these Panels consisted of well over 100 eminent and influential biologists, mostly members of the Academy. How the panelists themselves were chosen is not indicated, but there is no doubt that they constituted an authoritative group.

In due course, the Panels presented their reports. How they were dealt with is described in colorful terms:

> .... In a gruelling one week session of the Survey Committee... each report was *mercilessly* exposed to the criticism of all the other members.... Each report was then rewritten and subjected to the *searching, sometimes scathing*, criticisms of the members of the parent Committee on Science and Public

Policy. The reports were again revised in the light of this exercise. Finally, the Chairman of the Survey Committee...devoted the summer of 1968 to the final editing and revising of the final work.

Thus we have good grounds for regarding the contents of this volume as constituting a truly authoritative consensus, at least, as of 1970. There are no minority reports; no demurrals; biology does indeed seem guaranteed here to speak with one voice.

What does that voice say? Here are a few characteristic excerpts:

> The theme of this presentation is that life can be understood in terms of the laws that govern and the phenomena that characterize the inanimate, physical universe and, indeed, that at its essence life can be understood *only* in the language of chemistry. [emphasis added]

A little further along, we find this:

> Until the laws of physics and chemistry had been elucidated, *it was not possible even to formulate* the important, penetrating questions concerning the nature of life... . The endeavors of thousands of life scientists...have gone far to document the thesis...(that) living phenomena are indeed intelligible in physical terms. And although much remains to be learned and understood, and the details of many processes remain elusive, those engaged in such studies hold *no doubt* that answers will be forthcoming in the reasonably near future. Indeed, *only two major questions* remain enshrouded in a cloak of *not quite* fathomable mystery: (1) the origin of life...and (2) the mind-body problem...yet (the extent to which biology is understood) even now constitutes a satisfying and exciting tale. [emphases added]

Still further along, we find things like this:

> While *glorying* in how far we have come, these chapters also reveal how large is the task that lies ahead... . If (molecular biology) is exploited with vigor and understanding...a shining, hopeful future lies ahead. [emphasis added]

And this:

> Molecular biology provides the closest insight man has yet obtained of the nature of life – and therefore, of himself.

And this:

> It will be evident that the huge intellectual triumph of the past decade will, in all likelihood, be surpassed tomorrow – and to the everlasting benefit of mankind.

It is clear from such rhapsodies that the consensus reported in this volume is not only or even mainly a scientific one; it is an emotional and aesthetic one. And indeed, anyone familiar with the writings of Newton's contemporaries and successors will recognize them.

The volume to which we have alluded was published in 1970. But it is most significant that nothing fundamental has changed since then.

[5] In the inimitable words of Jacques Monod (1971, pp 42–43):

> We can assert today that a universal theory, however completely successful in other domains, could never encompass the biosphere, its structure and its evolution as phenomena deducible from first principles... . The thesis I shall present...is that the biosphere does not contain a predictable class of objects or events but constitutes a particular occurrence, compatible with first principles but not deducible from these principles, and therefore *essentially unpredictable*. [emphasis added]

References

Burks, A. (1966) *Theory of Self-Reproducing Automata* (Urbana, IL: University of Illinois Press).

Handler, P. (Ed) (1970) *Biology and the Future of Man* (Oxford: Oxford University Press).

Higgins, J. (1967) Oscillating chemical reactions. *J. Ind. & Eng. Chem.* 59: 18–62.

Monod, J. (1971) *Chance and Necessity* (New York: Alfred A. Knopf).

von Neumann, J. (1951) The general and logical theory of automata, in L.A. Jeffress (Ed) *Cerebral Mechanisms in Behavior* (New York: John Wiley & Sons) pp 1–41.

Rosen, R. (1977) Complexity as a system property. *Int. J. General Systems* 3: 227–32.

Rosen, R. (1978) *Fundamentals of Measurement and Representation of Natural Systems* (New York: Elsevier).

Rosen, R. (1979) Some comments on activation and inhibition. *Bull. Math. Biophysics* 41: 427–45.

Rosen, R. (1983) The role of similarity principles in data extrapolation. *Am. J. Physiol.* 244: R591–9.

Rosen, R. (1985) *Anticipatory Systems* (London: Pergamon Press) in press.

Shannon, C. and Weaver, W. (1949) *The Mathematical Theory of Communication* (Urbana, IL: University of Illinois Press).

Turing, A.M. (1936) On computable numbers. *Proc. London Math. Soc. Ser.* 2, 42: 230–65.