

SYSTEMS

1. *Systems in the News*

Dynamical-systems analysis is nowadays ubiquitous. From engineering (its point of origin and natural home) to physiology, and from psychology to ecology, it enjoys surprisingly wide application. Sometimes the analysis rings decisively false—as, for example, when adopted in certain treatments of historical narrative;¹ other times it is provocative and controversial, as when applied to the phenomena of mind and cognition.² Dynamical systems analysis (or “Systems” with a capital “S,” as I shall sometimes refer to it) is simply a tool of analysis. It mobilizes the language and mathematical technology of differential equations, and brings into play the distinctive concepts of *equilibrium* and *attractor*, as well as *gain*, *coupling* and *neighborhood*, that are not obviously proprietary property of any particular domain of objects or regime in the world.³ It is the ecumenical language of engineers, universal in scope.

Still, Systems, as a mode of analysis, itself stands in need of clarification. Once that clarity has been attained, one can then ask: are there limitations or bounds on proper application of Systems analysis, that are themselves premised upon considerations internal to the analysis itself? This is one of several questions to which the present essay is devoted. Before it can be attempted, however, we shall require some groundwork that clarifies the mode of analysis that is Systems—the family of analyses to which it belongs. This will begin to bring out (among other things) the precise difference of subject matter between biology and physics. And the ecumenicality of Systems analysis is bound to have its own distinctive commitments, as we shall see.

Practitioners attuned to the signal characteristics of Systems analysis—characteristics that set it apart—proclaim its many advantages

¹“Systems” by Mariam Thalos.

The Monist, vol. 92, no. 3, pp. 452–478. Copyright © 2009, THE MONIST, Peru, Illinois 61354.

over other styles of analysis. Indeed some Systems prophets, both early and late, proclaim that it resolves numerous philosophical difficulties, including the perennial mind-body problem and the late-comer "problem of consciousness."⁴ It would thus be of some considerable use to have a systematic (pun intended) treatment of these advantages, and why they fall to the Systems approach over against rivals. Indeed it would be of some considerable use simply to have a clear characterization of the contrast between the Systems approach and its rivals. That characterization shall be the centerpiece and focus of this study.

Norbert Wiener, apostle of the Systems approach to the natural sciences and engineering—and who, incidentally, preferred the term "cybernetics" (from the Greek term for steersmanship) for this style of treatment—contrasted it with Newtonian causal analysis. He maintained that the Newtonian analysis is "linear," but that the Systems analysis provides also for "circular" causation.⁵ The position was premised on analysis of the way *time* entered in each; according to Wiener, time in the Newtonian approach, is reversible; hence there is nothing new under the sun because (according to Poincaré's theorem) there is recurrence in systems governed by Newtonian mechanics: systems under Newtonian dynamics must return to their original condition (or, as close to it as we might like to insist) time and again. By contrast, Wiener pointed out, on the Systems approach time is unidirectional, and hence the novel is ubiquitous. What do these two diverse claims, the one pertaining to the "shape" of causal relations and the other pertaining to novelty and the directionality of time, amount to? Do they have anything to do with each other? Wiener did not say. (Perhaps he thought—mistakenly, of course—that their relationship is obvious.) I think there is something to what Wiener says. Something surprising. Something that Wiener himself might not have been prepared to accept. This something, as I shall be arguing, amounts to proclaiming that Systems analysis does not, strictly speaking, warrant the name of causal analysis at all, and moreover, that it may very well be incompatible with Newtonian causal analysis—not just different. Indeed, that Systems analysis is concerned with a topic—namely, control—that goes beyond mere causation. It is of course in the area of social science that systemic explanation has, notoriously, been contrasted with causal explanation—a matter to which we will return at the end of this essay.

2. Control, Not Causality

Causal analysis seeks to model the sequence of causes and effects, via causal laws. Here is an example with which we can work: water promotes plant growth. Such laws often must be construed as resistant to counterexample. For example, sometimes watering a plant too much will promote its decay rather than its healthy growth; and too little is no better; but neither fact should be sufficient to displace the generalization as a true causal law. And so causal laws are often construed as couched in (sometimes implicit) provisos that allow for a mean or average effect (and so we get the “exceptions that proves the rule”). The provisos spell out the details of a “background” against which the cause, when it appears, explains production of the effect. Sometimes these are referred to as “ordinary” or “standard” conditions. What are the standard conditions? They are very much case-specific. Causal analysis is thus straightforwardly statistical analysis of means, variations, and the conditions that make for variation from the mean.

The background conditions are traditionally referred to as “all-things-being-equal” conditions. They can fall into one of two classes: (1) irrelevant conditions—conditions that make no causal contributions of any kind to production of the effect; and (2) conditions whose contribution (whatever it might be) to production of the effect is *independent* of the contribution being made according to the particular law we are now considering, so that with an appropriate randomization of the population under scrutiny with respect to all these factors, the effects of the causally relevant factor will be distributed *normally* (on a bell curve). So, for example, watering causes plant growth. But so also do certain soil nutrients and nitrogen. When the contributions of each of these things to soil growth is independent of the others, then plant growth responds to *each* factor independently of how it responds to the others. When these independence conditions in fact obtain, a suitably chosen population will show a random distribution in plant height, and treatment with water will shift the mean of that population significantly. Similarly, a treatment with soil nutrients will shift the mean significantly as well, even when no change in water treatment is applied.

Suppose, however, that this independence does not obtain—perhaps (to stipulate the facts of a hypothetical situation) nitrogen levels boost plant growth only in the presence of sufficient levels of moisture, whereas moisture always boosts growth, with or without the presence of nitrogen.

In this case, we might say that moisture *modulates*, *mediates* or *moderates* the effects of nitrogen. It is no small feat to make sense of this very simply describable situation within a causal paradigm. Here is why.

Cutting-edge causal modeling methodology—causal path analysis, utilizing analysis of variance—relies upon a certain paradigm for diagnosing the causal bearing (“relevance”) of one factor upon another. The result of such an analysis is to divide up the variance, and attribute portions of it (of potentially different sizes) to independent players in the causal drama.⁶ According to the paradigm, causal bearing is revealed in, as Christopher Hitchcock refers to it, “test situations”: “A test situation is a conjunction of factors. When such a conjunction of factors is conditioned on, those factors are said to be ‘held fixed’. To specify what the test situations will be, then, we must specify what factors are to be held fixed.”⁷ (And this is why standard conditions are referred to as “all-things-being-equal” conditions.) In our hypothetical case, nitrogen (N), water (W) and plant growth (G) are the causal factors in question. To diagnose the causal bearing of N on G, we have to hold W “fixed,” but of course since water modulates the effects of nitrogen, holding water fixed can mislead us as to the bearing of N: if we hold it fixed too low, we will not see a proper shift in the mean. We will thus not be able to diagnose a straightforwardly positive causal bearing of N on G. (Effectively, what happens is that, rather than shifting a randomly collected sample population from one bell curve of growth distributions to another bell curve, manipulations of nitrogen will produce either no shift at all, or a non-bell curve, indicating a more complex relation among the factors than causal-bearing analysis can effectively handle.)

To handle this sort of example, causal modeling adopts an “interaction” paradigm, on which the factors relevant to plant growth (in this case two) are multiplied, so that not only do we have water and nitrogen playing independent parts in the causal drama, but also their “interaction” too—a proportion of their “cross product” (the cross product of their vector representations in the mathematical model). This fictitious entity is modeled as an independent causal player. How satisfying is this? One way to try to answer this question is to examine the founding principles of cutting-edge causal modeling.

Causal modeling is built upon an important assumption. This is the so-called “Markov Condition,” which is a generalization of Hans Reichenbach’s Principle of Common Cause. Very roughly, the Markov Condition

stipulates that the "causes" of X must *screen X off from* all other variables except for its own effects, so that correlations between X and other variables disappear once conditioned upon its causes. (Less roughly, if Q_i and Q_j are correlated and Q_i is not a cause of Q_j , and Q_j is not a cause of Q_i , then there are common causes of Q_i and Q_j in the set $\{Q_1, \dots, Q_n\}$ such that Q_i and Q_j are independent conditional upon these common causes.) And this condition stipulates that factors that have not (yet) interacted must be uncorrelated. Even if this apparently innocuous assumption is true simply as a matter of fact (rather than a matter of necessity), what place does such a principle play in causal analysis? Let's turn to the example again.

Applying the Markov condition to our nitrogen example, we have to say that the causes of plant growth—nitrogen, water and their interaction (water**nitrogen*)—must be "initially" uncorrelated. In other words, that the water**nitrogen* factor is, not only an independent causal factor, but also uncorrelated with water (as such) and nitrogen (as such). Another way to say this is that we have to *orthogonalize* the contributions of each of these players. We therefore use the Markov condition as a regulative ideal rather than as a simple assumption. And we do this in order to maintain the fiction that it is possible to attribute proper credit for a portion of variance to each of independent causal players.

Now, even if it makes some sort of sense to say that the interaction between nitrogen and water (the element water**nitrogen*) is itself independent of both nitrogen (as such) and water (as such)—and indeed it does make sense because we simply construct the water**nitrogen* factor (the cross product) so as to ensure it obeys our definition of independence—how can we construe the proposal that this interaction factor is a *third* factor? Presumably, all we are trying to get at, on this causal modeling proposal, is that water**nitrogen* is a third thing that one needs to keep track of in order to make sense of what's going on. But if we accept the proposal, have we really made sense of the idea that water *modulates* the effect of nitrogen? Or have we simply wished it away? Has the modulation simply been masked over for expediency's sake? I submit it has. I submit that the best this causal modeling can do for our special case is that it has identified some sort of dependence or link between the effects of water and nitrogen. But among other things, this analysis has missed out

on the asymmetricality in that dependence relation. The simple description “W modulates the effects of N” speaks clearly of an asymmetricality.

Some thinkers have suggested that the notion that causal factors must be thought of as independent keeps us from making sense of the notion of modulation.⁸ My own view is that causal analysis, involving the diagnosis of causal bearing as a categorical (on-off), is too coarse a tool to handle the sorts of modulations we find in this hypothetical nitrogen example. What is happening in this case is that an environmental or contextual factor is *coupling* with the variable whose magnitudes we are manipulating, in a way that obviates analysis in terms of “independent contributions” of different causal pathways. Thus the idea that the effect is under the control of the so-called “causal” variable whose bearing we are querying is inapt. The very notion of *control* itself requires treatment. To handle modulation, I submit we require a different sort of diagnosis altogether—one that allows for the context sensitivities we should like to put on display. Here is another way to put this point.

Causal analysis, resting as it does on diagnosis of causal bearing, involves the foregrounding of particularly salient contributory factors, against backgrounded conditions that are treated as largely negligible. It casts attention upon certain changing factors as causes (independent contributions, contributory or detracting to production of an outcome), at the cost of casting other stationary conditions (to wit, conditions that are not changing in the present context) in the role of background—and by implication noncontributory—factors. Where the assumption of “negligible” falls short, causal analysis papers over this fact by introducing the “cross product” to recover the negligibility, without illuminating how the cross product has done this. This sort of analysis has its place, of course. It gives us information that we might not otherwise be able to obtain. It gives us some very rough clues as to how to get things accomplished—for instance, how to ensure that adding nitrogen will have a desired effect. But where conditions favor it, there is room for another sort of analysis as well—one that retains in its crosshairs the “modulation” properties of the so-called *background*.

This essay will thus argue that, for the purposes of accounting for important features of large-scale behavior, the factors typically “backgrounded” by causal analysis—and in particular, those comprising training

histories, which make for the “precorrelations” that causal analysis attempts to define away—are really much more significant than normally given credit for in causal analysis. Indeed, that large-scale behavior is sharply attuned to the modulations and resonances (stored in “memories”) between entities and their surroundings. This is the province of systemic forms of analysis, and where causal modeling misses out—by its own design—on important features of the (preconfigured) landscape. Because some systems exercise a capacity for memory, their interactions with their environments require different conceptual foundations. And moving to a science of Systems acknowledges that a System’s context requires more attention than causal modeling can provide. It requires attention to the very notion of *control*.

To control is *not* simply to stand in a causal relationship (where by that we mean some kind of probability-raising relationship, as presumed in path analysis) to some event, or even to a network of such events. Control is a matter of *authorities*, a matter of which directives have a trumping effect (winning out when there is conflict); it is not essentially a local matter of where a process originates, nor is it a matter of which features of the world ultimately explain what the products of the process ultimately look like. Control is a global, all-over affair, a matter of *rank* and *protocol*, and not well correlated with the outcome or goal of a process or procedure.⁹ So it is very difficult to diagnose when encountered on the ground.

Control is best illustrated by examples where by hypothesis we know what the *protocol* is, or simply where we stipulate the control structure. Military organizations or militias are perhaps the best illustrations of what we may call *hierarchical* control, and other examples of control are contrasts to it. Protocol is less a matter of what forces, factors, or influences are actively exercised (changing or being changed) during the relevant real-time process, and more a matter of global features of the organization involved in that process. Control over my car, for example, is what I, as driver, lose when the brakes of my car fail at the top of the hill, even if I have occasion to call upon their service (and then entirely without success when I do) only at the bottom. During the period of time when I have no occasion to call on their service, matters will proceed exactly as they might well have done, had I indeed enjoyed control over them; still, I do not enjoy that control. Elsewhere I have adopted the

following rough and provisional analysis of *authority* or *rank*, preliminary to a mature account:

Axiom A: An entity, unit, or function A *ranks over* a second entity, unit, or function B, when the edicts or instructions or processes of A win out over those of B, in cases where the edicts or instructions or processes of the two are in conflict, or cannot go forward simultaneously.

Now, this axiom is palpably rough, and its roughness is instructive. Nothing in it takes into account any of three important qualifying ideas that will ultimately have to be handled in a mature conceptual treatment of control: (1) (*Partiality*) that ranking can be complete or partial; in other words, that it may leave entities unranked as to rank; (2) (*Circumscription*) that it may be circumscribed rather than global, applying to a very clearly delimited range of operations; and (3) (*Tempering*) that these two ways that a rank can be qualified or circumscribed might intersect or overlap in such ways as quickly confounds diagnosis of the precise balance of partiality and circumscription on the ground. For now all we need is the simple conception of ranking, as a piece of the relevant metaphysics for our examination of Systems analysis: it allows us to conceive of control structures as structures that overlay lower-level causal structures, and govern when and where lower-level structures interact. They choreograph. They direct. They are the stuff of upper management.

One concept that has been common in the disciplines of psychology and neurophysiology since the 1950s is that of *inhibition*.¹⁰ This concept is indeed a genuine control concept, much more than it is a causal one. The concept of inhibition is (roughly) the concept of veto or brake: the controlling entity or process prevents an action, even one that is already going forward, from going to completion. Indeed, this notion presupposes a notion of completion that has no place in path analysis. Inhibition is roughly one half the concept of ranking. The full concept also includes the controlling entity being, at least potentially, in a position to "arouse" or "excite" or "trigger" as well as inhibit. (This idea too has been in circulation in physiology and psychology, in the theories of memory, attention, motivation, and learning.) Put together, arousal and inhibition would seem to be two necessary, and perhaps also jointly sufficient, pieces of the

concept of rank. Perhaps it will turn out that to achieve more conceptual traction with the idea of ranking one will need to break it down further into these two component ideas. (Arguably, the two conceptions of feedback, which are discussed below, are the two corresponding pieces.) And it seems at least initially plausible that the two halves of the control function can be separated—that one entity or process can have inhibitory control over a second, but not excitatory control. It would serve us to investigate these sorts of questions further, at least conceptually. And when once a full complement of the necessary control concepts have been articulated, it will be possible to give a detailed account of homeostasis and similar physiological and systems notions in purely metaphysical terms.¹¹ For our purposes here, Axiom A will have to suffice.

When does a collection of objects under a scheme of control also qualify as a System in its own right? And how can we explain this process by which a System gets built out of components? This is the engineer's fundamental query.

3. *Simon's Criterion*

When one performs an analysis in terms of control, one must draw prominent attention to certain important features of the control landscape. The notion of control we have set down refers to global ("all-over") features of a system. A control analysis describes a system's actual state in relation to "nearby" potential states. It functions to track a system's *relative stability* in time, rather than track the specific means, in terms of forces and other features, that may (or may not) be holding it together. For certain purposes, this is only important information to keep track of—it is the kind of information never too far from the concerns of an engineer. Keeping track of such information about a system is very different from keeping track of its material characteristics as such. Keeping track of characteristics as bear on a system's stability in time involves keeping track of its large scale (global) relations to other things in its context, actual and possible.

Herbert Simon, decades ago, sought to articulate a similar conception.¹² He thought that certain features of certain systems could not be captured correctly by analysis of their parts taken together with the specific and heterogeneously conceived interactions among them. His idea derives from his concern with building systems. The idea is that to

build a complex system (as nature does), one proceeds in stages, with the result that at the end of each stage, what is constructed must possess a stable structure, so as to "hold still" whilst the next phase of operation is launched. Without these intervals or layers of stability, complexity (according to Simon) is unsustainable. This makes complex systems typically: (a) modular; (b) intersubstitutive in their parts; (c) qualitatively similar with a change to their parts or their number; and (d) stable under reaggregations of parts.¹³ (Imagine trying to articulate these ideas within a path analysis!) Simon is thus reaching for the idea that a System under control has to be governed by high-order structures of stability that are relatively independent of the sorts of interactions (physiological, chemical, mechanical or what-have-you) that govern their parts more locally. A true System (with a capital S) is something special indeed, and subject to high-order laws of dynamics and control among its constituents. A system that does not obey such laws is one that very soon falls apart. As I will explain, this amounts to saying that a System, with a capital S, is one in which the aggregation of the parts has undergone a *reduction in degrees of freedom*.

This concern for layers of stability in complex systems comports very well with the notion of control we have here been at pains to articulate. Indeed, I would insist that Simon had his finger on just this notion of control when he talked about system construction stages. Having explored the structural features of the process by which Systems get built, we must now step back to look at the product, and therefrom extract a formal characterization of the key elements of it. The defining characteristic of the System that emerges from such a process is simply that it is built by the process—and is consequently subject to being added onto in the same way.

How do sites of systemhood coalesce? They do so along a variety of different dimensions. The first and most obvious is that control can be secured over a certain array of independent resources; when this occurs so, we speak of a *pool* of resources or capacities. I will refer to this as the *zeroeth* dimension of merger. And I will offer a roughly *Lego model* of Systems formation. The zeroeth dimension is the emergence of a unit Lego block—the individual System slice. In nature, this is the aspect of physiology.

Second, and also very familiar, sites of control can combine along different lifetime slices and developmental stages of the same entity or System; when they do, we speak of an *individual system*. This I will refer to as the *vertical* dimension of merger, conceiving of time as advancing

upwards. Individuals will be represented by columns or towers consisting of unit blocks.

Finally, systemic entities can coalesce (across various types of boundaries, living and nonliving); when they do, we speak of *coalitions*. I will refer to this as the *horizontal* dimension, conceiving of coalitions as overlaps or unions in the horizontal plane. In the Lego model, a coalition is a multi-unit block that joins more than one tower. Enduring coalitions will consist of multi-unit towers rising vertically from the point of merger.¹⁴

We have thus resisted the temptation to treat any particular sort of entity as an inalienably *independent* atom—a decision that amounts to a commitment to diversity about atoms (as well as about independence). And concomitant with this rejected commitment is another—a commitment to treating systematic correlations or coordinations amongst these independent systems as a matter of incidental, “external,” or “ecological” interactions between them, rather than as systemic and inalienable properties of their confederation—their Systemhood. We are therefore laying the groundwork for contentions to the effect that entities in community or communion, in addition to retaining a certain amount of independence (at least in certain of their dealings), coalesce (predictably) at important conjunctions of circumstance, forming molecular systems that interact in complex ways with each other, as well as with and upon the atoms themselves. We are thus urging that societies of entities can consist of systems within systems; that they are rich complexes of intricately nested and overlapping systems. And that therefore higher-scale behavior, in all its forms, is best handled within a Systems approach. On this view, bonding is the fundamental characteristic of an entity, on a Systems analysis, because an entity is, from the Systems perspective, fundamentally a bond. (Another way to put this is to say that Systems are fundamentally units of aggregation.) On any other formula, systemhood, as a phenomenon, is bound to disappear.

How can we take stock of the elementally systemic—the fundamentally confederate—dimension of the lives of Systems with their environments? How do we keep from making this dimension disappear—as we have seen that causal modeling makes it do? I am proposing that to ensure that the systemic does not disappear, we need to confer a degree of freedom upon aggregates or confederations of entities and variables, as such—in order to give them a role in explanation. And this constitutes a profound departure from causal approaches to explanation, which rest upon the assumption of

independent atoms, and so are incrementally reductionistic. (This remark illuminates the point that reductionism rests fundamentally on a search for independent atoms, and that adherence to the Markov condition is reductionistic in just this sense.) When we seek a Systems analysis, we are instead postulating—not only that many systems are governed by Systems notions—but also that precisely the features of the systems that make them Systems also subject these systems to higher-order laws governing scale and organization. In the language we will be adopting straightaway, this amounts to saying that Systems enjoy different *degrees of freedom* than might be ascribed to them by a strictly causal analysis in terms of their parts and matter—because there are degrees of freedom to be conferred upon molecular aggregates or bodies, *as such*, within a Systems framework. A suite of technical notions is required to render this analysis, to which we now turn. As we will see, it takes the framework of analytical mechanics to make these notions possible.

4. *Degrees of Freedom in Mechanics*

I outline roughly here the contours of two incompatible but nonetheless venerable formulations of mechanics: Newtonian and Analytical.¹⁵ The Newtonian formulation is causal in a straightforward sense; causes, on the Newtonian formula, are put one-to-one with forces. The analytical formulation is not similarly causal. This shall set the stage for subsequent argument (in sections following) to the effect that the Systems approach adopts the second—the analytical—conception of mechanics. But that Systems also goes beyond mechanics—adding to it. For the scope of Systems is rather narrower. The central concern of Systems analysis is not so much to furnish an account of *any* system undergoing change, as to furnish an analysis of *systems under control*. And the topic of control is itself beyond the reach and aspiration of mechanics proper. Toward its goals, Systems analysis articulates a notion of control, a notion that is nearly as foreign to analytical mechanics as it is foreign to Newtonian mechanics. This is in itself strong evidence for the thesis that the topic of control is not within the ambit of the theory of pure mechanics: any account of control, even when it treats systems all of whose characteristics are uncontroversially physical, must go beyond mechanics as such.

The central mission of mechanics, in the modern tradition, has been to secure an account that accurately describes how changes in quantities

take place over time. A natural description of motion introduces vectorial quantities, with vectorial variables to portray them. The Newtonian formulation of general mechanics (historically the first comprehensive formulation) takes simple vectorial description of such alterations one step further, for with each alteration in motion it postulates at least one entity—a vehicle or emissary, routinely interpreted as a cause—that brings about the alteration by communicating an influence in the appropriate direction or orientation. The distinguishing mark of a Newtonian treatment is its utilization of Newton's Second law to formulae that enable computations of changes in magnitudes of motion variables.¹³

Not every continuous course of magnitudes for a given quantity, beginning with the initial magnitude and ending with the final one, is of concern in a Newtonian analysis, where by contrast all of these alternative histories are much the object of scrutiny in the second venerable branch of physical mechanics.

Analytical mechanics, by contrast, originated with ideas of Leibniz, Euler and Lagrange, and culminated in Hamilton and Jacobi's equations of transformation. These equations enable us to write down differential equations describing covariations in certain observable quantities. Rather than conceiving of alterations in time as a blow-by-blow, push-me/pull-you narrative of a drama of forces-and-their-aftermaths, taking place amongst entities playing their individual parts on the physical stage, analytical mechanics conceives of alterations, instead, as a wave phenomenon—a kind of global choreography without the notion of perpetrators. These wave disturbances take place in the space of quantities known as the *phase space*. There is no counterpart to force-metaphysics, the vehicle of influence communication, in this formulation of mechanics. Simply, there are tides in the affairs of quantities, whereby they conspire to undertake alterations together, and in keeping with very general principles that make no mention of causes as such. The distinguishing mark of analytical mechanics is the *variational principle*. Famous examples are Hamilton's principle (to the effect that the integral, over a system's path, of the difference between kinetic and potential energies is always an extremum—either a maximum or a minimum) and Huygens's principle for optics which leads to Fermat's "principle of quickest arrival" (to the effect that the path of a light ray is distinguished by the property that if light travels from one given point M to another given point N along that

path, it will arrive in the smallest possible interval of time). Change in these quantities, on the analytical formulation, is thus treatable—and treated—without appeal to forces and their communication from one object, event, or state of affairs to another.

Of considerable concern also in analytical mechanics are alternative courses of magnitudes for quantities under study. Variational principles operate upon these alternative courses of magnitude to select as distinguished (when there is a unique solution to the problem) the actual one. In analytical mechanics the set of possible trajectories must be in place, to make description of the *problem* complete, before application of a variational principle that selects among them a distinguished one as *solution* to that problem. For the dynamical laws—the variational principles—are in their essence contrastive: they select that trajectory which possesses an extremum of a particular characteristic, while presupposing that the set of alternative trajectories has an independent specification. (The actual trajectory emerges, according to Hamilton's Principle, as a result of the existence of a unique potential trajectory meeting all boundary conditions, and possessed of an extremum—typically a minimum—in the difference between potential and kinetic energies.) And so, without prior specification of the boundary conditions, not only is the system to be treated (the problem) underdescribed, but also—and consequently—the mechanical problem itself is ineligible of solution. For not only do the macroconditions and constraints help to identify the contrast set of potential trajectories, they also *define* that set. And in analytical mechanics they are treated as independent, both of actual magnitudes of microquantities and of the principles of contrast that constitute the dynamical laws. Analytical dynamics thus regards boundary conditions as (1) independent of dynamical law and (2) prior to it in ontology—since analytical dynamics must treat boundary conditions as a pre-existing problem subject to solution via variational principles. Analytical mechanics is thus a different beast, philosophically, from Newtonian mechanics.

Each of these two conceptions or formulations of dynamics, in traditional mechanics, requires embedding in a general theory of quantities and their potential dependence relations. For alone, dynamics does not suffice to produce a thoroughly general treatment of motion and change over time. Specifically, dynamics does not contain a systematic treatment of the nature of quantities, and how they relate to each other, both as time goes

on, and at any given instant in time. Once furnished, such an account will provide the means, at least in principle, whereby to assemble complete descriptions of systems comprising enormous numbers, conceivably nondenumerable, of interacting bodies, with many internal parts (and thus with proportional numbers of interacting quantities). But such a treatment will first have to answer questions like: How are macroquantities related to microquantities? Do microquantities conspire together to form macroquantities? How? Do they always do so in the same way? Can the macro play all the same roles in mechanics that the micro does?

Take this porcelain cup I now drink from. What shall we say is the relationship between its microscopic realities—the to-ings and fro-ings amongst molecules that constitute it—and the fact that the cup shatters upon impact with my kitchen floor? Is the relationship a one-way affair, in which the microscopic gets absolutely all the credit for what happens at the more macro level? If we answer yes, then we endorse a dogma to the effect that the microscopic is—always and everywhere—master, and that the macroscopic is always and everywhere slave. Tightly coiled in this master-slave doctrine is the conception that only micro quantities can serve as degrees of freedom—and this is a conception hostile to Systems analysis. But this very master-slave doctrine is nonetheless an implicit cornerstone of the Newtonian architecture of analysis: it is an implicit axiom, in the Newtonian framework, that macro causes bring about macro effects entirely in virtue of the fact that micro causes bring about micro effects.¹⁷ For example, the cup of our acquaintance shatters, according to this analysis, because certain intermolecular forces in it (in the porcelain that forms it), are in very fine and quantitative ways different from those in, say, a steel fork or a textile, and are such as to make the cup “fragile.”

By contrast, this master-slave doctrine is *not* enshrined in the Analytical formulation of mechanics. It is an open question, on the analytical formulation, whether a cup’s fragility is an additional, independent property that belongs to it, in addition to others it possesses at the micro level, which might be displaced as poolings, mergers, and alliances are forged in Systemic fashion. Analytical mechanics can accommodate a decision to model such properties as independent. Given its philosophical structure—which is far less rigid—analytical mechanics can thus allow a certain macroconditions, associated with a given System, to *swing free* so to speak from micro conditions, and (to use the language of diplomacy)

analytical mechanics permits us to *represent* the object so swung free to the rest of the world by a reduced number of (typically more macro) conditions. Analytical mechanics thus allows us to acknowledge that object's Systemhood as such. And this is what we require if we are to allow Systems, as such, to possess degrees of freedom as wholes.

The key to this achievement is that analytical mechanics neither acknowledges nor incurs an obligation to introduce quantities, such as forces, that mediate between the actual and the possible. Second, it views certain macroconditions of a system (conceived either as boundary conditions—for example, the condition of being fixed to a certain track—or as magnitudes of macroquantities) as capable of preceding in ontology magnitudes of microquantities with which that system, as a System, interacts. In other words, the laws acknowledged by analytical mechanics do not rule out (as do both Newtonian dynamics and the Markov condition), a pre-existing correlation between quantities in a System or in a System-situated. Analytical mechanics, as contrasted with Newtonian mechanics, allows for (even demands) the independence of conditions describing a “pre-existing” correlation. For example, it allows us to treat a train “clinging” to its track as a pre-existing condition—a kind of “designed” and therefore nonnegotiable feature. This shows decisively that the principles of analytical mechanics are quite different from those of Newtonian mechanics, and potentially incompatible with them, as they are potentially incompatible with cutting-edge causal analysis.

The analytical conception of a system in analytical mechanics is, as I will refer to it, *contextual*. By this I mean that analytical mechanics views it as inappropriate to request characterization of a physical system taken “in isolation,” if by that is meant “without specification of boundary and/or design conditions.” For, on the conception of system that emerges here, there is no such entity as *this system in splendid isolation*, if this system is not already in splendid isolation.¹⁸ Portrayal of a system, on the analytical approach, is always of a system situated, rather than of a system-by-itself. And this is just what we need for a Systems analysis: the room for ineliminable pre-existing conditions. What they amount to is something else entirely, as we will soon see.

Now, Systems analysis is itself an application of analytical mechanics. But Systems is a very special sort of application. In a Systems treatment, just as in an analytical one, there is no predesignation of certain

quantities as causes, and others as effects. Rather, there is a certain respectful silence regarding what may qualify as legitimate choice of dynamical variable. Moreover, Systems moots the question whether bodies are just assemblages of quantities. Systems, as I will now show, makes a beginning at an account of “body” or “unity,” by identifying a spatial boundary—a membrane or skin—between what shall count as “body” or “figure” and what shall count as “surround” or “ground.” Thereby Systems adds what we can with good reason refer to as “perspective.” Systems analysis thus makes a beginning at the job that mechanics (as such) declines: the job of counting the number of bodies, thought of as the number of unified objects, in a given region. And so Systems is an application of analytic mechanics, but one that takes an important philosophical step towards meeting biology and psychology half way.

5. *Systems, Feedback, and Control*

Analytical mechanics provides a hospitable framework for Systems analysis, providing the sufficient flexibility within which to confront directly the question of what makes a body a unity, and not just an unstructured assemblage of quantities. Systems analysis is thus a transformative rendering of mechanics, aiming to bring into focus within the discipline of dynamics concepts that are more unapologetically at home in physiology. How? By introducing an innovation: the notion of a *boundary* between that which is under the System’s control—its body—on the one side, and environment, on the other; it offers a nonarbitrary opposition between inside and outside, where mechanics alone will countenance only fields of force and their “effects” amongst quantities, or variational principles and their aftermaths. Thus Systems offers the making of a rudimentary conception of organ, and therewith function. Systems achieves this marvelous feat, as I will now show, by taking for granted the notion of equilibrium, which is to be sure a physical conception—indeed a thermodynamic one—but not a mechanical one. The condition of equilibrium obtains in the controlled (the “internal”) regime.

In thermodynamics, equilibrium is defined as that condition of a system in which a small number of thermodynamic quantities—distinguished among them, temperature—remain unaltered indefinitely into the system’s future.

Systems analysis of a given system begins with tabulations of certain quantities (I will refer to them as $\{\alpha_i\}$) of the target System S, together

with certain other quantities (I will refer to them as $\{\rho_j\}$) of that System's environment E. The whole business is assumed to remain throughout the time of analysis in a state of quasi-equilibrium, or (as I will refer to it hereafter) a state of *control*. (The term "homeostasis" is sometimes used here. I will not use it, as it suggests—falsely—the existence of a unique state or condition that control processes are targeting.) This means that, throughout the time of analysis, alterations to features of S+E off-list of tabulated quantities are for all intents and purposes irrelevant. This assumption is not trivial or innocuous. And it is not available for all systems in which someone might wish to take an interest. In fact, it is precisely the condition under which a Systems analysis is valid. The condition obtains only under circumstances in which control (as we will now discuss) is possible.

Thus Systems analysis supposes, first of all, *ab initio*, that control is possible. Then it proceeds with description of how that control is attained. Control is attained when the controller system S acts (literally) to bring its own condition to within tolerance of a prescribed reference point. This is done via *feedback loops*.

In one sort of feedback loop—the *positive feedback* loop—a System responds to changes in a variable (say ρ_j) in the *same* direction as the perceived change: when detection of a drop (for example) in ρ_j occurs, S takes steps to further increase that drop. In contrast, a system set up to respond to changes in any of the variables being tracked (any of σ_i or ρ_j) by making changes that *reverse* a perceived change (by taking steps to move that variable in the direction opposite to the perceived change) is acting in a *negative feedback* loop.

The onset of contractions in childbirth operates on a positive feedback loop: when a contraction occurs, the hormone oxytocin is released into the body, which stimulates further contractions. This results in contractions increasing in amplitude and frequency. Blood clotting is another example. The process is initiated when injured tissue releases signal chemicals that activate platelets in the blood: an activated platelet releases more chemicals to activate more platelets, causing a rapid cascade and the formation of a blood clot. A nursing female's production of milk increases at its young's demand. The "launching" of a physiological function thus involves amplification of an initially minor change (in itself or its environment). But if the activity initiated is not to be ultimately destructive to the System, the process must be halted before its escalation reaches cata-

strophic proportions. How does a System know the right time to arrest a process it has itself launched?

In most cases, once the process' target is reached, a second process comes online to damp further escalations. The damping process requires negative feedback loops: these loops reverse the "launching" mechanism by activating processes that counter changes to a variable: childbirth contractions stop when the young has been expelled; chemicals that break down blood clots are released at a suitable time; lactation stops when the baby no longer nurses. "Control," as Weiner correctly observed, is in the first instance over oneself. Only secondarily or derivatively is it control over the environment.

Physiological control is thus expressed in one part to launch processes to bring about a new target state, and in another part to damp potentially runaway internal processes that threaten a system's integrity. In both cases, robust physiological systems are alert to important "reference points," both internal to themselves and in environmental variables. And they treat different reference points differently to achieve this control. Reference points are fundamental physiological universals.

When control is achievable in the way just described—that is to say, when background equilibrium conditions favor it—the states of S+E can themselves be regarded as a system of interconnected equilibrium conditions that can be reached by fine tuning operating characteristics (for example, the reference points and what they trigger) of the control device S. Cybernetics, then, is the study of the properties of networks of interconnected equilibrium conditions, often in relation to reference points. It encompasses precise treatments of the performance of control apparatus, and features special attention to defective behavior that tends to bring about oscillations (from mild to violent) in control quantities, when control apparatus is poorly designed, mishandled, diseased, infected by a foreign control apparatus, or generally overloaded.

Many biological systems fall squarely and unapologetically within the scope of a Systems analysis. Consider the flocking of birds. Dynamacists have only in recent decades been developing the idea that their remarkable synchrony of motion can be explained by simple local adjustments to motion ("rules of flocking") mediated through sensory modalities such as vision, sound, pressure, or odor detection. Assuming that each organism in a flock can sense local flockmates as well as its en-

vironment, and adjusts its own motion on an ongoing basis, Craig Reynolds devised a computational model of flocking “boids” based upon the following “rules of engagement”:¹⁹

1. Separation: steer to avoid crowding local flockmates in your “near neighborhood.”
2. Alignment: steer towards the average heading of local flockmates.
3. Cohesion: steer to move toward the average position of local flockmates.

In effect, flocking requires uniformly of each flock member only that it reacts to flockmates within a certain small neighborhood around itself, characterized by a *distance* and an *angle* (measured from the organism’s own vector of motion). Flockmates outside this local neighborhood can be ignored. (Similar models have been devised to model the collective foraging behaviors of social insects, for example, ants in an ant colony.)

But details of the large scale features of the synchronized motion matter a great deal—because flocking needs to be very finely tuned if it is to serve the interests of the flock. So how, precisely, does fine tuning of the “rules of engagement” scale to collective properties of the larger motion? For example, how does adjustment of the size of the neighborhood affect sensitivity to environmental features? How does it serve in location of food and avoidance of predators? The answer might surprise: much depends upon features of the flock’s “rules,” and not at all on any feature of any given organism in the flock. The “rules” concern “all-over” characteristics, that might well be shared with flocks of very different species, and in no way depend upon the biology of the species. They are control characteristics, that have to do with how a flock of N boids reduce their degrees of freedom in flight.

We can illustrate this reduction in the context of our present flocking example. Close behavioral coupling among near neighbors in a flock allows a localized change in direction to be amplified and propagated across the flock. This allows each flock member to influence and be influenced by flockmates much farther away than their local neighborhood—it gives each a much larger “effective perceptual range” than their actual sensory range. This scaling is nonlinear. Study of the details of the

scaling relations reveals that it is hard for groups to maintain cohesion if the coupling distance is too short. Longer-range transfer of information is enabled by increasing the coupling distance. Increasing the coupling distance further still creates a cohesive group, but “misinformation” might be propagated (as use of information about motion of distant individuals is in some circumstances less beneficial locally). Damping processes will thus prove useful if long-range transfer of information is absolutely essential.

In addition, coupling can be moderated by context. For example, if individual birds in flight conditioned their reactions upon context (under threat, for example, aligning more strongly with distant flockmates, increasing “system gain”), this could allow for some flexibility, but there is a cost. Heightening sensitivity to weak or ambiguous environmental signals increases susceptibility to “false positives,” and damping response to local fluctuations in less threatening contexts increases “false negatives.” A balance has to be struck.

What has become more and more clear in this discussion is that a flock of hundreds of organisms, operating under a set of “rules of engagement,” is decidedly *not* a system with degrees of freedom on the order of hundreds: it is instead a system with something on the order of a dozen degrees, counting among them rough size, coupling distance, and level of context sensitivity, as well as environmental variables that tend to couple with these features. From a flocking perspective, a flock is an entity with a reduced number of degrees of freedom than there would be without the rules. And some of these degrees lie in the environment itself! These degrees of freedom have displaced many or most of the “micro” variables that reign when the flocks are not flocking, as soon as the flock members begin governing themselves by the rules of engagement.

The flock couples to its environment to achieve flock-specific goals or activities. It is a System with a memory. And so its behavior is modulated by features of the environment, in much the same way that nitrogen modulates plant growth in the example that opened this essay.

One point already mentioned can be usefully reiterated here, by way of repelling potential criticism to the effect that Systems analysis is no more than a certain, possibly distinctive style of causal modeling. It is important to emphasize that Systems analysis leaves out vast amounts of detail, pertaining to quasi-equilibrium quantities, that would be required in any causal treatment of (for one thing amongst many) feedback loops—

such as for instance the size, shape, and other dynamical features of different flocks in that array that constitute the flock. Indeed the equilibrium conditions do not themselves come in for much handling. They are simply taken for granted. And vast amounts of detail are left out too pertaining to the interactions between control quantities (coupling distance, for example) and those that remain largely unchanged (perhaps average distance to nearest neighbors), or simply those that change in ways that do not impact control of the process. This fact testifies to the idea that Systems analysis is supremely uninterested in certain detail—in the *causes* (if you will) that mediate between controlling quantities and controlled. Systems analysis focuses simply on the dynamical aspect of control, not on the (underlying, if it is that) aspect of how control is wrought or maintained.²⁰

6. Complexity

It has been said that we need Systems analysis because the world is complicated. Because there are more particles (and so more quantities) than a mathematical model governed by Newton's laws can handle. Its necessity is a matter of practicality—an “applied” rather than a theoretical matter—not a matter of principle. Systems analysis is thus, according to this idea, a concession to human frailties.²¹ And therefore we stoop to Systems analysis because it's a simple and incontestable fact of life that only superhuman intellect or yet-unattained mathematical facility can handle the computational complexities we face when we attempt to treat in theory the real-life, many-body systems we manipulate on those everyday occasions in which we seek to achieve a measure of control. Complexity is thus an occasion for throwing up one's hands and giving in to a lesser form of analysis.

Our analysis here gives us the resources to repel this argument. For, while the argument insists that complexity is always with us, even when conditions are right (that is, when the special sort of equilibrium conditions obtain) for control, our analysis has insisted that—to the contrary—only a small set of conditions make control possible. So that complexity, if it does anything, can do no other than make control *overwhelmingly unlikely!* For complexity—as such—is precisely what one typically has on hand in all those many circumstances when control is just not in the cards.

Introducing the notion of complexity here thus only draws a veil of mystery about what may not be all that mysterious. It is, ultimately, a red herring.

And so, in reply to critics who insist that we need Systems analysis because Systems analysis handles better the complexity of relations amongst highly interactive quantities under control conditions, we ought to reply that we need Systems, not so much *because* the world is so complex (though it is frequently complex), but rather *in spite* of it. Systems does not oversimplify, as those who protest of complexity charge. Indeed, it brings into focus what the complexifiers insist on obscuring: namely, that certain dependence relations, at the heart of certain protocols, are the ground of control. Protocols are (as we noted already) simplifiers and not complexifiers. And so Systems analysis—analysis of higher-order dependence conditions—is analysis of how complexity is reduced. And Systems brings (finally) into clear focus an urgent question not seen before—a question that Systems analysis itself, as such, cannot answer: Why are quasi-equilibrium conditions as ubiquitous as they are? And hence: Why is control as possible as it is? It is a question for which I have seen no answers of any kind suggested, plausible or otherwise. Here, perhaps, is where study of complexity can usefully come in.

7. *Functional Explanation*

Emile Durkheim proposed to explain, holistically, the existence of certain large-scale features of society by way of a proposal to the effect that societies are wholes with needs, purposes, or goals that exert pressure on their parts to behave so as to meet or achieve said needs or goals. (For example, he proposed to explain the existence of criminal activity, as “an act that offends very strong collective sentiments,” by the fact—if it is a fact—that criminal behavior functions to promote social solidarity.) This functional brand of explanation, as now it is called, is regarded as problematic if no feedback mechanism is identified.²² But thinkers on the subject now proclaim that once a feedback loop is specified, this “then turns out to do most or all of the explanatory work in an unpuzzlingly causal way.”²³ For this and other reasons, cause and explanation are now broadly considered synonymous.

I, by contrast, maintain that Systems analysis and causal analysis are *not* synonymous. Indeed, feedback loops are proprietary to Systems analysis, but not to causal analysis as such. Feedback loops of the sort needed to support the persistence of a large-scale feature invoke a notion of equilib-

rium. And so we have to separate conceptually the notion of feedback from the notion of causal explanation: the ideas do not belong to the same family.

But a note of caution must also be sounded: Systems analysis does not guarantee achievement of any goal, however cherished. It *acknowledges* goals, and can describe in abstract terms (that is to say, in terms of the reference points that enter into patterns of stability and quasi-equilibrium) the conditions that make achievement of some of them possible. It is another matter altogether whether the conditions obtain. There may be no explanation of why they obtain when they do so. And this is among the many things one needs to know when one is seeking an explanation of a large-scale phenomenon: one wants to know (for example) which factors in the environment promote, sustain, or thwart criminality (in one instance) or religious fervor (in another). Want of attention to conditions that make achievement of goals possible is precisely what makes functionalism so unpalatable. So a Systems explanation is certainly an improvement upon the functional explanations of old. But of course there is plenty of room for disputing that Systems, all on its own, provides all that is or might be wanted in the way of explanation for large-scale social phenomena. And all this holds true whether we are considering a functional explanation for prevalence of crime, or for a mechanism that maintains bodily temperature within a certain range of tolerance.

Mariam Thalos

University of Utah

NOTES

1. See for example the reaction of P. Roth and T. Ryckman, "Chaos, Clio, and Scientific Illusions of Understanding," *History and Theory*, vol. 34 (1995), pp. 30–44.

2. See books with such titles as *Stairway to the Mind* (New York: Springer-Verlag, 1995), by Scott Alwyn, and *Thinking in Complexity: The Complex Dynamics of Matter, Mind and Mankind* (New York: Springer-Verlag, 1994), by Klaus Mainzer. There are less ambitious treatments, but nonetheless in the same spirit: Mark Bedeau, "Weak Emergence," in *Philosophical Perspectives: Mind, Causation and World*, James Tomberlin, ed., (Blackwell, 1997), pp. 375–99.

3. Introductions to the concepts and language of Systems analysis can be found in Ralph Abraham, *Dynamics: The Geometry of Behavior* (Redwood City, CA: Addison-Wesley, 1992), and H.W. Broer and F. Takens, eds., *Geometry and Analysis in Non-linear Dynamics* (New York: Wiley, 1992).

4. Hubert Dreyfus has recently revived an argument he originated with Stuart Dreyfus against a computational model of human expertise, drawing on Walter Freeman's dynamical-

systems treatment of the psychological processes of perception, as well as the phenomenological work of Merleau-Ponty ("Intelligence without Representation—Merleau-Ponty's Critique of Mental Representations: The relevance of Phenomenology to Scientific Explanation," *Phenomenology and the Cognitive Sciences*, vol. 1 (2002), pp. 367–83; cf. *Mind Over Machine* (New York: Free Press, 1986) and Walter Freeman, "The Physiology of Perception," *Scientific American*, vol. 264 (1991), pp. 78–85). The new argument allies itself, at least for purposes of handling human expertise, with Tim van Gelder's model of cognition as a dynamical 3-way interaction amongst brain, body, and world. Whereas for Dreyfus the guiding examples have to do with expert performance such as driving, playing tennis, or chess, or even adjusting a vantage point on a piece of artwork, the guiding example in van Gelder's thinking is the so-called *Watt governor*, a device that automatically adjusts, through feedback mechanisms, the throttle valve on a steam boiler so as to maintain uniform speed of the flywheel despite changes in steam pressure or workload. (T. van Gelder, "What Might Cognition Be, If Not Computation?" *Journal of Philosophy*, vol. XCII (1995), pp. 345–81.) In seeking to replace the reigning paradigm of the computer with that of the Watt governor, van Gelder is tapping into a movement at the forefront of which figure certain cognitive scientists (primarily neuroscientists, roboticists, and ecological psychologists) dispute whether there is anything performed by human brains that deserves to be called representational. Sometimes invoking Piagetian themes of development, these scientists maintain that the best means to understand and explain the operations of embodied minds are *not* the folk-psychological instruments of mental concepts, mental contents, and mental representations but instead dynamical ideas, expressible in the language of differential equations to which physics also appeals, that interrelate changes in time amongst control variables in brain (as they call them), body, and world. See especially, E. Thelen and L. Smith, *A Dynamic Systems Approach to the Development of Cognition and Action* (Cambridge: MIT Press, 1994), but also J.A.S. Kelso, *Dynamic Patterns* (Cambridge: MIT Press, 1995); F. Keijzer, "Doing without Representations which Specify What To Do," *Philosophical Psychology*, vol. 9 (1998), pp. 269–302; H. Maturana and F. Varela, *Autopoiesis and Cognition* (Dordrecht: Reidel, 1980). Thelen and Smith, for example, write: "Explanations in terms of structure in the head—beliefs, rules, concepts and schemata—are not acceptable. . . . Our theory has new concepts at the center—nonlinearity, reentrance, coupling heterochronicity, attractors, momentum, state spaces, . . ." (p. 339). Kelso writes: "The thesis here is that the human brain is *fundamentally* a pattern-forming, self-organizing system governed by non-linear dynamical laws. Rather than compute, our brain dwells (at least for short times) in metastable states" (p. 26). Dynamical-systems accounts of the mind are at this point in time broadly behaviorist accounts. (Indeed one gets the distinct impression that at least some of the apostles of dynamicalism, notably Thelen and Smith, are still very much in the grip of the romantic notion of science as observation and deduction from the phenomena.) They may or may not acknowledge "internal structure" to the cognitive agent, but their overriding research goal is interrelation of the outputs of cognition with the stimulus conditions that elicit them.

5. *Cybernetics* (Cambridge: MIT Press, 1948), especially chapters I and III.

6. See J. Pearl, *Probabilistic Reasoning in Intelligent Systems* (San Mateo: Morgan Kaufman, 1988); C. Glymour, P. Spirtes, and R. Scheines, *Causation, Prediction and Search* (Springer-Verlag, 1992).

7. Hitchcock, Christopher, "Probabilistic Causation," *The Stanford Encyclopedia of Philosophy* (Fall 2007 Edition), Edward N. Zalta, ed., <http://plato.stanford.edu/archives/fall2007/entries/causation-probabilistic/>.

8. Much of what Nancy Cartwright says in *Nature's Capacities and their Measurement* (Oxford, 1989) can be read this way. See also her "Where is the Theory in our 'Theories' of Causality?" *Journal of Philosophy*, vol. CIII, no. 2 (2006), pp. 55–66; and "From Metaphysics to Method: Comments on Manipulability and the Causal Markov Condition," *British Journal for the Philosophy of Science*, vol. 57 (2006), pp. 197–218.

9. This question is directly related to the question of whether a process is *centralized* or *distributed*, and this is a broadly logical or computational feature of the process—it pertains to a design attribute, rather than to features special to physical implementation. A centralized process is a hierarchical (or top-down) process, with a single "processor" residing at the top-most level of control, whereas a local or distributed process is one in which control is distributed horizontally at the top-most level amongst a plurality of processors equally ranked from the point of view of control. We might say of distributed processing that the initiative is "grass-roots." And obviously control can, at least logically speaking, be distributed at intermediate levels as well: a process can be centralized at the top-most level but distributed at lower levels, and vice-versa. D.A. Norman and T. Shallice, "Attention to Action: Willed and Automatic Control of Behavior," in *Consciousness and Self-Regulation: Advances in Research and Theory*, Richard Davidson, Gary Schwartz, and David Shapiro, eds. (New York: Plenum Press, 1986), pp. 376–90, offers an intersection model of top-down and bottom-up processing without a ranking central processor. Valerie Gray Hardcastle, "A Critique of Information Processing Theories of Consciousness," *Minds and Machines*, vol. 5 (1995), pp. 89–107, makes a beginning at a philosophical discussion of these topics. A fuller account of the difference between control structures and causal structures is discussed in Thalos, M. "Sources of Behavior: Towards a Control Account of Agency," *Distributed Cognition and the Will*, Don Ross and David Spurrett, eds. (Harvard: MIT Press, 2007), pp. 123–67.

10. It originated in the 1950s in the work of Eliot Stellar, in the form of dual-center theory, now generally thought overly simple.

11. The metaphysics of systems theory is currently in its infancy. Thalos (2007), *op. cit.*, makes some progress in articulating the metaphysical differences between systems theory and causal analysis.

12. "The Organization of Complex Systems," in H. Pattee, ed., *Hierarchy Theory: The Challenge of Complex Systems* (New York: Braziller, 1973), pp. 1–27.

13. W. Wimsatt, "Forms of Aggregativity," in Donagan, Perovich and Wedin, eds., *Human Nature and Natural Knowledge* (Dordrecht: Reidel, 1976), pp. 259–91; and more recently in "Aggregate, Composed and Evolved Systems," *Biology and Philosophy*, vol. 21 (2006), pp. 667–702, and *Re-engineering Philosophy for Limited Beings* (Harvard, 2007), makes reference to many of these conditions as well in describing the difference between aggregates (on the one hand) and composed or evolved Systems (on the other), and W. Bechtel and Robert C. Richardson, "Emergent Phenomena and Complex Systems," in Anskar Beckerman, Hans Flohr and Jaegwon Kim, eds., *Emergence or Reduction* (New York: de Gruyter, 1992), pp. 257–88, illustrate the ways in which scientific methodologies attain what they refer to as emergent phenomena utilizing these criteria. Wimsatt also warns that our appreciation of the value of frequency information (with a bearing on causation, of course) is hampered by failures to appreciate the ways that the objects in the regime we are studying are composed or organized.

14. The Lego model is a helpful visualization of very simple coalition structures. But we must not be misled into thinking that the spatial properties of Legos must carry over into the world of agents.

15. The outline draws on Thalos, "Degrees of Freedom: An Essay on Competitions between Micro and Macro in Mechanics," *Philosophy and Phenomenological Research*, vol. 59 (1999), pp. 1–39, and Thalos, "Nonreductive Physics," *Synthese*, vol. 149 (2006), pp. 133–78.

16. The conceptual details are spelled out in Thalos: "In Favor of Being Only Humean," *Philosophical Studies*, vol. 93 (1999), pp. 265–98.

17. As argued in Thalos (1999) and (2006), op. cit.

18. C. Lanczos, *The Variational Principles of Mechanics* (Toronto: University of Toronto Press, 1949) puts it this way: "The analytical approach to the problem of motion is quite different [from the Newtonian approach]. The particle is no longer an isolated unit but part of a 'system.' A 'mechanical system' signifies an assembly of particles which interact with each other. The single particle has no significance; it is the system as a whole which counts."

19. "Flocks, Herds, and Schools: A Distributed Behavioral Model," *Computer Graphics* 21(4), 1987, 25–34.

20. A third point can be added, though it might not carry authority with all readers: there are no grounds in a Systems analysis for prohibiting faster-than-light signaling, which some have compellingly argued is foundational to the notion of causation. And once again this is testimony to the difference of topic.

21. Something of this complaint enters into Paul Teller's analysis of emergence in "A Contemporary Look at Emergence," in Anskar Beckerman, Hans Flohr and Jaegwon Kim, eds., *Emergence or Reduction* (New York: de Gruyter, 1992), pp. 139–53.

22. This style of analysis has been popular on and off in the analysis of meaning. See, most recently, Ruth Millikan, *White Queen Psychology and Other Essays for Alice*, (Cambridge: MIT, 1993).

23. Martin Hollis, *The Philosophy of Social Science* (Cambridge University Press, 1994), p. 113.