CHAPTER TEN

The Ontology of Complex Systems

Levels of Organization, Perspectives, and Causal Thickets

Willard van Orman Quine once said that he had a preference for a desert ontology. This was in an earlier day when concerns with logical structure and ontological simplicity reigned supreme. Ontological genocide was practiced upon whole classes of upper-level or "derivative" entities in the name of elegance, and we were secure in the belief that one strayed irremediably into the realm of conceptual confusion and possible error the further one got from ontic fundamentalism. In those days, one paid more attention to generic worries about possible errors (motivated by our common training in philosophical skepticism) than to actual errors derived from distancing oneself too far from the nitty-gritty details of actual theory, actual inferences from actual data, the actual conditions under which we posited and detected entities, calibrated and "burned in" instruments, identified and rejected artifacts, debugged programs and procedures, explained the mechanisms behind regularities, judged correlations to be spurious, and in general, the real complexities and richness of actual scientific practice. The belief that logic and philosophy were prior to any possible science has had a number of distorting effects on philosophy of science. One effect was that for ontology, we seemed never to be able to reject the null hypothesis: "Don't multiply entities beyond necessity."

But Ockham's razor (or was it Ockham's eraser?) has a curiously ambiguous form—an escape clause that can turn it into a safety razor: How do we determine what is necessary? With the right standards, one could remain an Ockhamite while recognizing a world that has the rich multi-layered and interdependent ontology of the tropical rain forest—that is, our world. It is tempting to believe that recognizing such a worldview requires adopting lax or sloppy standards—for it has a lot more in it than Ockhamites traditionally would countenance. Quite to the contrary, I think that the standards for this transformation are not lax, but only different. Indeed, the standards that I urge are closer to our experience and arguably more fundamental than those used during the hegemony of foundationalist methods and values.

In the first section of this chapter, I discuss the criterion for what is real—what I call *robustness*—a criterion that applies most simply and directly, though not exclusively, to objects. In subsequent sections, I use robustness and other information about our world to delineate the major structural features—primarily levels, but with some comments on what I call *perspectives* and *causal thickets*—that dominate our world, our theories, and the language we use to talk about both. These are higher-level ontological features, *Organizational Baupläne,* related to the things that people usually talk about under the topic of ontology (things like objects, properties, events, capacities, and propensities) as paragraphs are to words and phonemes or morphemes. But they are there nonetheless, it is only our concern with the little things, motivated by foundationalist or reductionist concerns, which has deflected our attention from them. This ontology—of levels, perspectives, and causal thickets—is no less required for a full accounting of the phenomena of the physical sciences than it is for biology and the social sciences, but its obdurate necessity has seemed more obvious in these latter cases. This may now be changing. The increased interest in fractal phenomena and chaotic and, more generally, non-linear dynamics emerging from the so-called exact sciences has brought many noisy residua of the ontological scrap heaps of the physical sciences to the center of attention as theoretically revealing data, structures, and objects with new-found status. Most of these things have never before made it into theory—or if so, only into the "theory of observation" under the topic of "error analysis" where they lived in the ubiquitous error term. Messiness—or at least the right kinds of messiness—is now almost a virtue in many of the sciences, as the recent explosion of interest in complexity seems to attest.[1] Levels, perspectives, and causal thickets are major ontological players in these complex areas—domains with significant implications for how to approach many of philosophy's most refractory problems.

Because the aim of this chapter is ultimately taxonomic—to say what there is, or to describe some of the *bigger* things that are—the descriptive sections basically take the form of a list of properties, elaborated either to further explain ideas likely to be unfamiliar, or to explain relations among the properties that help to give the ideas of level and perspective their cohesiveness. Taxonomy may sound boring, but I hope to show you that the description of and relations between a family of newly discovered species can be an exciting task.

I. Robustness and Reality

Before I say what there is in this complex world, I should give my criteria for regarding something as real or trustworthy. Particularly among those of a foundationalist persuasion, it is common to start by providing some criterion, be it indubitability, incorrigibility, or other means of picking out things or assumptions whose veracity is not open to question. One then says that those things are real (true, indubitable, or whatever) if it is either one of these primitive things or if it is derivable from them via a valid series of inferences. Only things admitted in one of these two ways are allowed. I share the foundationalist's concern with securing reliability for our conceptual structures, but I don't think that there *are* any criteria that *both* give indubitability or render error impossible, *and* permit any interesting inferences from that starting point. Thus, I would rather give a criterion that offers *relative* reliability, one that you're better off using than not, indeed better off using it than any other, and that seems to have a number of the right properties to build upon. Rather than opting for a global or metaphysical realism (an aim that bedevils most of the analyses of "scientific realists"), I want criteria for what is real that are decidedly local—which are the kinds of criteria used by working scientists in deciding whether results are real or artifactual, trustworthy or untrustworthy, objective or subjective (in contexts where the latter is legitimately criticized—which is not everywhere). When this criterion is used, eliminative reductionism is seen as generally unsound, and entities at a variety of levels—as well as the levels themselves—can be recognized for the real objects they are, and traditional foundationalism and ontic fundamentalism are in trouble. They will survive, if at all, as a local kind of problem-solving technique of significant but limited usefulness. (But see Chapter 7, on dynamical foundationalism.)

Following Levins (1966), I call this criterion *robustness*. (Chapter 5 analyzes and reviews this concept and methodology; Wimsatt, 1980a, 1980b, has relevant case studies. Campbell's [1966] concept of "triangulation" captures many of the same ideas, and his classic work with Fiske [1959] on the "multi-trait-multi-method matrix" brought this methodology to the social sciences.) *Things are robust if they are accessible (detectable, measurable, derivable, definable, producible, or the like) in a variety of independent ways.* A related but narrower criterion (experimental manipulability via different means) has since been suggested by Hacking (1983), who draws a close link with experiment, and limits his discussions to the realism of entities. But robustness plays a similar role also in the judgment of properties, relations, and even propositions, as well as for the larger structures—levels and perspectives—described below (see also Wimsatt, 1981a, 1974, 1976a). Furthermore, independent means of access are not limited to experimental manipulations but can range all the way from non-intrusive observation or measurement to mathematical or logical derivation, with many stops in between. Experimental manipulation is just a special case. We feel more confident of objects, properties, relationships, and so forth that we can detect, derive, measure, or observe in a variety of *independent* ways because the chance that we could be simultaneously wrong in each of these ways declines with the number of independent checks we have.[2] We can only make the probability of failure decline—though it can get very small, it does not go to zero. This criterion does not give certainty. Nothing does. There are no magic bullets in science—or anywhere else, for that matter. But if that's so, then certainty is not so important as generations of philosophers have supposed.

The independence of these different means of access is crucial. Independence is often not easy to demonstrate, and failures of independence are often well hidden. Cases of pseudo-robustness, while not common, are not truly rare either, and invariably seem to involve unperceived failures of the independence assumption, or—relatedly—not sufficiently broad variation in the means of access.[3] (Wimsatt, 1980b, 1981a, discusses cases of spurious or pseudo-robustness in population biology and psychology, and Culp, 1995, gives a careful and enlightening dissection of degrees of independence and interdependence among experimental techniques in molecular genetics. See contrary arguments by Rasmussen, 1993, and Culp, 1994, about the use of robustness in the analysis of an artifactual "entity," the mesosome, in recent cell biology.) Indeed, if the checks or means of detection are probabilistically independent, the probability that they could *all* be wrong is the product of their individual probabilities of failure, and *this* probability declines very rapidly (i.e., the reliability of correct detection increases rapidly) as the number of means of access increases, *even if the means are individually not very reliable.* This gives us the requisite sense of independence for this criterion—namely, that the *probability of failure* of the different means of access should be independent. Of course, one cannot infer immediately from apparent physical independence of the means of access to their probabilistic independence. That is a further hypothesis that is sometimes false. Probabilistic independence represents a kind of mathematical idealization—a mathematical model of physical processes or, in more complex cases, of a system of interrelated physical, biological, psychological, and social processes.

Although nothing will guarantee freedom from error, robustness has the right kind of properties as a criterion for the real, and has features that naturally generate plausible results. Furthermore, it works reliably as a criterion in the face of real world complexities, where we are judging the operational goodness of the criterion—not its goodness under idealized circumstances. We are judging its performance as well as its competence, as it were. It even has the right metaphysical and epistemological properties. Thus, it is part of our concept of an object that objects have a multiplicity of properties, which generally require different kinds of tests or procedures for their determination or measurement. It follows that *our concept of an object is a concept of something that is knowable robustly.* Indeed, one of the ways in which we detect illusions is that appearances to one sensory modality are not borne out with the appropriate confirmation in the other sensory modalities—confirming, for a visual hallucination or mirage that what we see before us is not an object, not real (Campbell, 1966).

Robustness can wear two faces in a kind of epistemological figure-ground reversal that leads to a kind of almost magical appearance of bringing yourself up by your own bootstraps. Particularly in the early stages of an investigation, we may use *agreement* of different means of detection, measurement, or derivation to posit an object or an objective property or relation that is the common cause of these various manifestations. At a certain stage, we will accept the existence of the entity or property as established—however corrigibly—and begin to use the *differences* observed through the diverse means of access to it as telling us still more about the object. (It is after all *that kind of thing or property that is detectable via these diverse means,* and *shows itself differently through them.*) We will at the same time use these differences to tell about the means of access to the object. (This *one* thing or property appears in these diverse ways *through these different means of access.*)

In this latter stage, we may compare the performance of the different means on a variety of target objects. In so doing, we are both calibrating each means against the others, and learning about their respective limitations.[4] This kind of switching back and forth can lead to considerable successive refinement both in our knowledge of the object(s) in question, and of the characteristics and limitations of the tools we have for accessing them.[5] The fine tuning and power of the refinements are increased if the objects in question turn out to form a class of diverse entities that can all be studied via the same means—as genes did for the Morgan school (Wimsatt, 1992).

Robustness has had a surprising history—it seems to be always there, but seldom noticed. Thus, seventeenth-century philosophers made a distinction between primary qualities (shape, extension, impenetrability, etc.) that they held were really in objects, and secondary qualities (color, taste, sound, etc.) that they held were induced in us by our interactions with the primary qualities of objects. Descartes took the primary qualities of objects as the fundamental properties of matter from which he tried to explain all else through derivation, and it was a general feature of such theories to try to explain secondary qualities in terms of primary qualities. This kind of relationship between primary and derived things became central to and emblematic of deductive and foundational approaches. The ironic fact, not noted at the time, is that the properties that Descartes and others following him chose as primary qualities were all knowable in more than one sensory modality, whereas the secondary qualities were known in just one sensory modality.[6]

Thus, in modern jargon, *the primary qualities are robust and the secondary qualities are not.* The explanatory principle of that period translates as: *Explain that which is not robust in terms of that which is—or, by extension, that which is less robust in terms of that which is more so.*[7] This is still a good principle, and one that is generally followed—it serves equally well in foundationalist and in non-foundationalist camps. It is different from, independent of, and if anything, more basic than anything else in the foundationalist methodology. Ironically then, we see that the paradigm of foundationalist approaches is simultaneously a paradigm use of robustness as a criterion for the real, and that the best applications of the deductivist paradigm occur when the foundational assumptions, objects, or properties are robust.

This indicates a coincident starting point for deductivist and robustness paradigms. There are other ways—elaborated in Chapter 5—in which they diverge. Thus, on the deductivist paradigm, the length of derivations doesn't matter (as long as they are finite), and additional derivations of the same conclusion through different means are redundant and unnecessary. But if overall reliability is the primary concern, and one has at each stage a small but finite chance of misapplying valid inference rules, then the length of serial deductive arguments does matter. Furthermore, in a world where failure is possible, multiple derivations of a result by different paths are no longer otiose as a way of checking or providing further support. One can stray still further from foundationalist values: *with parallel independent means of support available and the net reliability of the conclusion as the only concern, there is no longer any reason to limit inferences to truth-preserving ones, and the use of good inductive, abductive, or more generally, heuristic principles may have a place in the construction of exemplary arguments—in philosophy as well as elsewhere.*[8] Indeed, robustness as a criterion of superiority among arguments can and should cast a very broad and long epistemological shadow, once we get away from the unrealistic assumptions about human reasoning that have anchored 350 years of foundationalist thought.

I intend to apply these methodological lessons right here. Throughout this chapter, I not only use the concept of robustness as a tool in the analysis, but I also employ it in the structure of the argument by using multiple concepts and arguments that individually have a heuristic character—having less than deductive analytical force. There are lots of characterizations that represent strong tendency statements, which can be cashed out in terms of statistical rather than universal claims. This is data that can't by the nature of the objects be formulated or used in arguments that require necessary and sufficient conditions. Attempts to tighten them up would only render formulations that are too narrow in scope or fail to capture most of the interesting phenomena. It is suggestive of the situation for which "fuzzy set theory" was invented, though the present character of that theory makes no allowance for the systematic character of biases and exceptions (Wimsatt, 1985, 1992). This is a common pattern for entities, regularities, mechanisms, and explanations involving complex systems, yet we shouldn't refuse to discuss them for that reason. They are too important for their reality to be denied, or rendered suspect by false simplifications or idealizing assumptions. We should value *for that reason* an analysis that recognizes the centrality they have in everyday life.

In a way, then, this analysis has something in common with folk psychology and some of the basic assumptions of ordinary language philosophy—like them it takes for granted that the world we see, live in, respond to, and act upon is too important, too central to our way of being, to be dismissed. But this much is not just anti-scientific sloppiness (ordinary language philosophers went much further). For all of the ontological radicalism of quantum mechanics, Niels Bohr felt the need to postulate his "correspondence principle"—that an adequacy condition for quantum theory was that it had to produce (in the right limits) the macroscopic phenomena we observe everyday. The approach advocated here proceeds more like Bohr (in spirit, if not in content), and less like ordinary language philosophy in trying to suggest the outlines of a more realistic scientifically motivated epistemology and metaphysics for approaching these problems. But before attending to the ordinary phenomenology of this new taxonomy, a bit of abstraction is necessary to see where we are going in this new philosophical landscape.

Ontologically, one could take the primary working matter of the world to be causal relationships, which are connected to one another in a variety of ways—and together make up patterns of causal networks. (I won't address problems with causality in this chapter. Those who favor "Humean skepticism" will also find lots else to object to here, and can stop reading now unless they want to see how far you can get without it!) These networks should be viewed as a sort of bulk causal matter—an undifferentiated tissue of causal structures—in effect the biochemical pathways of the world, whose topology, under some global constraints, yields interesting forms. Under some conditions, these networks are organized into larger patterns that comprise *levels of organization,* and under somewhat different conditions they yield the kinds of systematic slices across which I have called *perspectives.* Under some conditions, they are so richly connected that neither perspectives nor levels seem to capture their organization, and for this condition, I have coined the term *causal thickets.* Much of psychology and the social sciences, for all the appearances of local order and local approximations to levels and perspectives, when looked at more globally and once the various idealizations of our theories are recognized, seem to be in this third state, or in a hybrid mixture that contains elements of all three. These three kinds of structures are rich in methodological and philosophical consequences for understanding the strengths and limitations of different approaches to studying problems and phenomena in systems characterized by one of them. We now turn to the first of these *Organizational Baupläne*—levels of organization.

II. Levels of Organization

The analysis presented here elaborates on parts of two earlier papers on reductionism and levels of organization (Wimsatt, 1976a, and Chapter 9). There has been a fair amount of work on levels since, in which they are taken to mean an astounding variety of things. Much of it, though relevant to the analysis of some complex systems, leads in the wrong direction for present purposes. Thus, I agree with McClamrock's argument (1991) that Marr's (1982) three levels (algorithmic, computational, and hardware) are better viewed as levels of analysis or of abstraction, or as kinds of functional perspectives on a system, than as compositional levels of organization. This conflation is apparently a common kind of mistake among philosophers of psychology.

More generally, people sometimes talk as if the material, psychological, and sociocultural realms constitute monadic levels (e.g., as in Popper's first, second, and third worlds). These rough distinctions are of major importance because they delimit regions where different major concepts, theories, methodologies, and explanatory strategies dominate, but they are larger heterogeneous aggregates spanning multiple levels and including also other less well-ordered structures rather than single individual levels of organization. Thus, by *any* criteria, there are obviously multiple compositional levels of organization within the material realm: elementary particle, atom, molecule, macro-molecule, and so forth, or, within the biological realm, as units of selection, for example, selfish genes (transposons), some kinds of supergenes (chromosome inversions), selfish gametes (the t-allele case in mice), selfish cells (cancer), selfish organisms, and selfish groups—all of which would fit into the material realm, traditionally conceived.[9] Similarly, most current cognitive theories recognize multiple levels of a compositional character[10] within the mental realm: structural representations of belief or planning, linguistic structure, or hierarchical representations of features in a classification system. Atomic families, small groups, mobs, speakers of a local dialect, social classes, sectors of the economy, and citizens of a nation-state are all obviously social, or sometimes sociocultural units at diverse levels of organization—whose interactions follow diverse dynamics.

By level of organization, I mean here compositional levels—hierarchical divisions of stuff (paradigmatically but not necessarily material stuff) organized by part-whole relations, in which wholes at one level function as parts at the next (and at all higher) levels. Though composition relations are transitive (so one could collapse the highest level systems to the smallest parts), levels are usually decomposed only one level at a time, and only as needed.[11] (Thus, neurons are presumably composed of parts like membranes, dendrites, and synapses, which are in turn made of molecules, which are in turn made of atoms, and so forth down to quarks, though to the connectionist modeler, neurons are adaptive modules with properties like incoming and outgoing connections and thresholds, and which might as well be indivisible atoms for all of the use that is made of their still lower level properties.) Most of what I say below relates to material compositional hierarchies and levels, because I utilize constraints characteristic of the physical world—which also includes the physics of biological, psychological, and social objects.

Nonetheless, this is not a reductionist analysis in the sense in which a philosopher might use that term. (I would urge, however that it *is* reductionist, or at least broadly mechanistic as those terms would be understood by most scientists. See Wimsatt, 1976b, 1979.) Nor should it be taken as implying, either in evolutionary history, or in current state-of-the-art genetic engineering, that usually or always, the preferred, most effective, or (stepping back to punt) even a *practically possible* way of making a given upper-level object is by assembling a bunch of lower-level parts. This over-extension of what I have called (1976a) the "engineering paradigm" is one of the things that have given reductionism and materialism bad names. (I remind the reader that the paradigms of genetically engineered molecules are not examples of *ab initio* constructions, but rather examples of the conversion of naturally occurring organic factories to the production of other products.) There is some assembly to be sure, but it is assembly of the jigs on the production line, and sometimes rearrangement and redirection of the line—not construction of the factory. To believe otherwise is to mistake arguments *in principle* for arguments *in practice*. (For the limitations and interpretation of such in principle claims, see Chapter 11.) Ultimately, we sometimes just have to stop promising and deliver the goods.

One of the reasons that it is important to look at material compositional levels more closely is that a number of properties of higher-level systems, which are treated as if they were emergent in some non-reductionist sense, follow directly from rather general properties of purely material compositional levels.[12] Thus, there is nothing intrinsically mentalistic (or social or cultural) about multiple-realizability, or the dynamical autonomy of upper level phenomena, or the anomalousness of higher-level regularities relative to the lower-level ones. Though each of these traits has been taken by some philosophers to be characteristic of the mental, they are actually characteristic of any move from a lower compositional level to a higher one. That goes for the theory of chemical bonding relative to fundamental quantum-mechanical theories of the atom no less than for "the" relation between the neurophysiological (*which* neurophysiological level?) and "the" cognitive (*which* cognitive level?). *These traits are features that always accompany the emergence of a new stable level of organization.*

As a kind of reductionist, I want to get as much as I can about higher levels from the properties of lower ones. As a kind of holist, it is tempting to try to do the reverse. For evolving systems, it is not controversial to argue that the arrangement of lower-level parts (and consequently the appearance of certain higher-level phenomena) is a product of higher-level selection forces (Campbell, 1974b). And you can do both at the same time (and we do) as long as you don't commit yourself to saying that the system you study is to be exhaustively characterized by one approach or the other, but regard them as complementary. So it is possible to be a reductionist and a holist too—but not *any* kind of reductionist, or holist. Unlike an eliminative reductionist, I think that we add knowledge of both the upper level and the lower level by constructing a reduction. We add to the richness of reality by recognizing these linkages—not subtract from it. Eliminativists generally worry too much about the possibility of error at the upper level, and not enough about how stable and resilient—how robust—most upper-level phenomena are, a fact that can make the upper-level details more revealing under some conditions than the lower-level ones.

The notion of a compositional level of organization is presupposed but left unanalyzed by virtually all extant analyses of inter-level reduction and emergence. A pioneering and important attempt to deal with levels of organization (and even more with the naturally resulting concepts of hierarchy) is Herbert Simon's (1962) classic "The Architecture of Complexity," which contains both useful conceptual distinctions and arguments of absolutely central importance. The views expressed here show Simon's influence strongly, but go further in other directions. I urge a view that Simon would share: *that levels of organization are a deep, non-arbitrary, and extremely important feature of the ontological architecture of our natural world, and almost certainly of any world that could produce, and be inhabited or understood by, intelligent beings.* (This gives levels an almost Kantian flavor.) *Levels and other modes of organization cannot be taken for granted, but demand char-*

acterization and analysis. If I am right (Wimsatt, 1976a), compositional levels of organization are the simplest general and large-scale structures for the organization of matter. *They are constituted by families of entities usually of comparable size and dynamical properties, which characteristically interact primarily with one another, and which, taken together, give an apparent rough closure over a range of phenomena and regularities.* (For anyone who still believes in "necessary and sufficient conditions" style analyses, I note at least five qualifiers in this sentence—all apparently necessary—that would be difficult at best to deal with, and the referents of these qualifiers are also often disturbingly general, and correspondingly unclear. Note also, that I said that levels are "constituted by," not "defined in terms of." Definitional language is notoriously unhelpful in contexts like these. Broad-stroke characterizations, focused with qualifications and illuminated with examples, are more useful.)

Levels are in many ways the ontological analogues of conceptual schemes—though without the difficulties said (e.g., by Davidson, 1973) to attend the supposition that there is more than one of them. We live in or at one, and most of our important everyday interactions are with other entities at our level of organization—i.e., with people, tables, chairs, cars, dishes, or computers. We don't normally interact with a person's cells, or with a computer's memory chips. Persons and computers are designed to be opaque with respect to the operation of their lower-level hardware—we don't usually "see" such hardware details unless they cause a macroscopically observable malfunction, or unless we take the deliberate and special additional steps to allow us to observe things at different levels. Most of the explanations of the behavior of an entity, and most of the means for manipulating, causing, or modulating its behavior, will be found and most naturally expressed in terms of entities, properties, activities, and regularities at the same level. Our level is our common world of folk psychology, or more broadly, of the objects that populate Sellars' "manifest image" or its scientific same-level descendants.[13]

A number of other levels are also accessible to us—in part because their effects occasionally leak up or down to our level (through those few interactions that fail to be characteristically level-bound),[14] and in part because we have actively searched for and exploited these few direct connections with other levels to enrich and expand our awareness of and control over these other domains of phenomena within and around us.[15] [Author's note, 2003: In doing so, we are "extending our senses," a particularly apt description since our senses at the one end—

and developmental adaptations of cognition and physiology at the other—already are designed to stretch the range of size and time scales over which we can perceive changes in and act upon nature. See the discussion of environmental grain in Chapter 12 for further explication of these ideas.] Because any complex material objects can be described at a number of different levels of organization, identity, composition, or instantiation, relations must hold between descriptions of the same object at different levels. These provide additional important means of accessing the different levels and calibrating relations between them, and the inspiration for explanatory reductionist mechanistic theories of the behavior of the systems in question.[16]

At lower levels of organization (those of the atom and molecule) we tend to have well-defined types of definitely specified composition and, at least in principle, an exhaustively specifiable range of possible states. At higher levels of organization (from our anthropocentric perspective, but definitely middle-range on a cosmological scale) levels become less well-defined in terms of size scale and other properties (see the top row of Figure 10.1). Higher-level types of entities may no longer have crisp compositional formulae,[17] but cover a range and, in some cases, composition may no longer be a primary individuating characteristic.[18] They must do so for two connected reasons: (1) the disparately composed entities at a given level may nonetheless show multiple similarities in their behavior under similar conditions—all to be covered by multiple regularities (thus engendering at least rough multiple-realizability as the rule rather than the exception), and (2) these similar entities found at higher levels, despite their similarities, become occasions for an increasing number of exceptions to whatever regularities we can construct (see Wimsatt, 1972) because of the increased richness of ways entities have of interacting with one another (due in part to the increasing number of degrees of freedom and of emergent properties).

As the richness of causal connections within and between levels increases, levels of organization shade successively into two other qualitatively different kinds of ontological structures that I have called, respectively, "perspectives" (Wimsatt, 1974) and "causal thickets" (Wimsatt, 1976a). Objects whose mode of organization is characterized by the three distinct types of structures (levels of organization, perspective, and causal thickets) have interestingly different consequences for the methodology of sciences that study them. Below I describe some properties of levels of organization, and then say rather less about perspectives and causal thickets. These remarks are intended less as an analysis (in terms of necessary and sufficient conditions) than as a characterization of some of their most important properties (many of which are discussed further in Wimsatt, 1976a). The complex interplay of these various criteria and forces that mould levels of organization is one of the main things that give the complex sciences their richness and texture.

Levels of organization have a variety of properties that make them rich in ontological and epistemological consequences. Taken individually, these properties seem to be almost accidentally associated—important but merely empirical or contingent properties. Looked at more closely, their *merely empirical* status is probably more a product of the fact that they haven't yet been taken seriously by any of the dominant philosophical views. In fact, these properties of levels are closely connected in ways that make the features of levels and their analysis not just a contingent empirical matter. (For further discussion of some topics not found below—including the role of first- and third-person perspectives in an account of levels of organization and further remarks on the degree to which levels of organization are inevitable features of nature and of our conceptual scheme—see Wimsatt, 1976a.) In the following section I discuss these contingent properties, tying them together with a network of further empirical and conceptual facts as I go.

1. Compositional Levels of Organization: The Role of Size

a. *Successive levels of organization represent a compositional hierarchy.* If one entity is a part of another it is characteristically at a lower level of organization than the other, though in some cases and for some purposes, parts of roughly commensurate sizes as the whole system are treated as being at its level. Entities at the same level of organization are usually of roughly the same size, though there tends to be greater size variance (even proportionally) at higher levels of organization, largely due to the increasing number of degrees of freedom and ways of interacting characteristic of larger systems. With the "engineering paradigm" (Wimsatt, 1976a)—that we normally assemble complex systems out of simpler parts, a process that can be iterated—entities at successively higher levels of organization tend to show roughly geometric increases in size (see also Simon, 1962).

b. *Size and surface/volume ratio, which is a function of size, are major factors in determining which physical forces are most central to the explanation of behavior* (see Haldane, 1927), *so the size of characteristic objects at a level is not an accidental feature of this analysis.* Changing size is a necessary consequence of compositional hierarchies (given the old saw about how two [simple] objects can't occupy the same place at the same time), but changing size is also central to how different level entities get their different properties. The size-scaling factor between adjacent levels is not arbitrary—if so it would have a simpler solution. To see this, let's suppose it were arbitrary. Why not arbitrarily pick, for example, a binary aggregation scale in which every time two similar (same-level) objects are aggregated, it involves going up a level of organization? This would surely be both possible and preferable if levels were determined by convention, or by a search for the most algorithmically economical generating relations.

Nor is it entirely without a physical basis. Binary aggregation seems natural for the architecture of computer memory, and binary doubling is naturally inherent in cell replication. In fact, starting with the same elementary particles, this scheme would produce an organizational hierarchy of all nature as regular as a giant fractal lattice. (This would be both simpler and far more elegant than what actually happens.) But, pursuing the cell-division example for a minute, this line does not produce natural vertebrae in the search for nature's joints for more than the first few cell-divisions past the zygote. Then differentiation begins, and other properties become more important, such as which cells are inside and which are on the outside of the developing cell-mass. Cell divisions in different lineages lose their synchrony fairly quickly in most metazoans. Some cell-types die and are continuously replaced by others of the same type, while others go on dividing with no significant mortality in their lineages. Consequently, organisms with a large number of cells show no tendency greater than random to have their cell-numbers be at or close to integral powers of two, and the relevant functional units don't show bottom up binary regularities either. The basic problem with binary aggregation is that this aggregation mode does not track the regularities found in nature—the entities thus produced would seldom be those with any broad natural significance.

This idea of binary aggregation was introduced as an aggregative mode which—despite occasional significant pairing—is so obviously not an architectural principle for the natural world to demonstrate that the problem has a natural rather than a conventional or purely formal solution. (One might ask social constructivists why this is so!) Although size scale is an important causal determinant of levels of organization, it is not the only one. The relevant (and highly variable) geometric scaling factor between successive levels is itself a complex function of the interplay of different physical forces on relatively stable structures at the different levels, and the kind of system in question.

c. Size is a relevant, and in many cases a good criterion because a number of causal interactions characteristically become significant or insignificant together for things in a certain size range. *Size is thus a robust indicator for many other kinds of causal interactions.*[19] *This should be one of the reasons why physics has so many straightforward and simple applications to aspects of our macroscopic world.* Dust particles and bacteria are not *prima facie* good choices to be functional analogues for anything, but their common size and mass range nonetheless create strong similarities across whole arrays of their behavior. They both make excellent Brownian motion particles—and indeed the discoverer of Brownian motion made the plausible assumption that all such particles were alive. (After all, how else could small entities move around apparently actively in an obviously inert fluid!?) Size has further consequences for the design of means of locomotion in bacteria that have to deal with the fact that at their size scale, it is not a trivial matter to move in ways that are not both reversible and reversed—and thus for their movements to actually take them anywhere (Purcell, 1977)!

d. *Size is not a sufficient indicator of level*—consider bacterium-sized black holes. These definitely would not exhibit Brownian motion, at least not for conditions found in our part of the universe because they would be incomparably too massive. This is not (just) a philosopher's silly hypothetical example, though it may have been a physicist's game. An extended series of letters in the journal *Nature* in 1974–1975 discussed the existence and properties of black holes in the size range of 10^{-2} to 10^{-4} mm. in diameter. Cosmological debates had suggested that the creation of such microscopic black holes in the early history of the universe was a possibility. The discussion in *Nature* considered whether one of them could have caused the gigantic explosion over Tunguska in Siberia in 1908 (the standard candidate is a meteor some 40–50 meters in diameter). Debate ceased when it was pointed out that on the black hole hypothesis there should have been a comparable exit hole and explosion in the Baltic Sea shortly thereafter. Such a black hole (1) would not show Brownian motion, or behave in any other way like a Brownian motion particle; and (2) things around it would respond to it in a bulk, aggregate, or an "average" way—for example, the rate at which it will accumulate mass and emit radiation is a function of the net disposition of mass around it, not of the detailed organization of that mass or how it is grouped into particles or chunks. (It is so much more massive than they that its trajectory and relative rate of mass accretion—over short periods of time—is also virtually independent of them and their velocities, but only depends on where its trajectory passes relative to them. However, the objects close to the black hole are dominated in their behavior by its presence—they behave to it as an individual: individual details of its motion, size, and so on do matter for them.)

e. The example of the dust particle–size black holes suggests a natural criterion in addition to composition for ordering entities by level of organization—probably a sufficient criterion, but alas not a necessary one:[20] *Of two entities, if one relates to the other's properties as part of an average, but the second relates to the first as an individual, then the first is at a higher level of organization than the second.* This is of somewhat broader applicability in characterizing levels than compositional relations because it enables one to order entities that are not above and below one another in the same compositional hierarchy. It indicates a kind of individuation asymmetry relating to scale that is generally true of things found at different levels in compositional hierarchies, but is not limited (as the part-whole relation is) to things in the same hierarchy. In addition, it seems plausible to say that two things that relate to one another as individuals are at the same level, and two things that relate to each other as parts of averages are both embedded in larger systems, but may vary relative to each other with respect to level.[21]

2. Levels and the Simplicity of Stratification: A Layered Tropical Ontology and the Consequent Development of Language Strata

f. *Levels of organization can be thought of as local maxima of regularity and predictability in the phase space of alternative modes of organization of matter.* This is the closest I will come to a definition, because this characterization has rich connections with a number of other important properties of levels. The levels must be viewed as occupying a remapped space of reduced dimensionality relative to this enormous phase space of all physically possible states of matter, since in the levels-oriented ontology, there are strong interactions and similarities among quite diverse kinds of systems.[22] Because they are compositionally very diverse, these systems will tend to be far apart in the embedding phase space, but because they are similar in terms of the variables appropriate to the levels description, they must be close together in the reduced-dimensional projection of that space in terms appropriate to that level.[23] Almost all entities are at levels. *Since most direct interactions of something at a level of organization are with other things at that same level of organization, regularities of behavior of that entity will be most economically expressed in terms of variables and properties appropriate to that level.*[24]

In talking about these as local maxima, I mean to imply that entities with modestly larger or smaller values of key properties (think of size) would show messier regularities than and key into fewer regular relationships with the other entities and each other than is true for the entities we have. The larger number of regularities or stable patterns involving the larger number of relatively stable entities, both concentrated at or near levels of organization, makes the characterization of levels as local maxima of regularity and predictability correct. This is analogous to a kind of "fitness maximization" claim for ontology, springing from a deep embeddedness of our world in a spectrum of different equilibrating and selection processes acting on different size and time scales (see also Dennett, 1995, for convergent "deep" claims about an evolutionary ontology and dynamics).

g. The fact that most direct interactions of something at a level of organization will be with other things at that level means that detectors of entities at a level will be or will have parts that are at the same level as the target entity, and that will interact with it via properties characteristic of that level. This has several direct implications:

1. The theory of instruments for *us* to detect properties or entities at level x will involve causal interactions, mechanisms, objects, properties, generalizations, and regularities of level x.
2. If we are at a different level, this theory of instruments will also involve causal interactions, mechanisms, objects, properties, generalizations, and regularities at *our* level, since we need to be able to detect and record their output. For these reasons, and for others, eliminative reduction is often not possible, necessary, or desirable—our very instruments anchor us at our level, as well as at the level we are observing. Such instruments are *inter-level transducers*.
3. The entities of a level will be multiply anchored through causal interactions to other entities at that same level, and will therefore show substantial robustness at that level.
4. Many of the properties attributed to entities at a given level (or sometimes attributed to the instrument used to detect them) will in fact be disguised relational properties—properties of the interaction between target entity and instrument. (This, or something like it, should be the correct move for the classical secondary qualities, but it also occurs for many other theoretical properties—perhaps most notoriously fitness, which is a relational property of phenotype and environment, but is misleadingly attributed without qualifications to organisms, traits, and genes.)

5. Many of the apparent ontological paradoxes characteristic of different level accounts of a system—paradoxes which may appear to require the elimination of upper-level properties and entities to a zealous reductionist—arise from forgetting this relational character. In Eddington's "two tables" paradox, there is nothing contradictory in saying that this table is both continuous, colored, and solid (when using my fingers and eyes as probes) and at the same time mostly colorless empty space (when using a beam of electrons as a probe.

h. Theories come in levels (to analogize an observation of John Dillinger) because that's where the entities are. Simpler theories can be built with those entities (and their major interactions) than with slightly larger or smaller or otherwise different ones. On this account of the theorist as bank robber (or forager, or economist), theories of entities at levels provide the biggest bang for a buck. These entities will be theoretically fruitful because of their many causal interactions, and the appropriate choice of entities at levels will more often produce naturally segmented systems that are nearly decomposable—which "cut Nature at its joints" (Wimsatt 1976a). Thus language (in which concrete nouns—entity words—are learned first) and theories constructed using and refining this language are in this way *responses to* rather than *determiners of* the structure of the world.[25] A causal asymmetry is asserted here that runs counter to most recent linguistic or social-relativist views of the world. During the heyday of linguistic philosophy one might almost have had the impression that nature came in levels because language came in strata—a kind of theory dependence or conceptual scheme dependence of our ontology.[26] For most of the natural world, this has it exactly backwards: language is a tool for dealing with problems in the environment (including the human environment, and including the environment of different levels of organization accessed by our ever-further-reaching and multi-faceted instrumentation). For the most part, language has the macroscopic structure that it does because of the structure of the environment, and only relatively rarely is it the other way around. If most of the robust entities are at levels (as they are),[27] then the levels will themselves be robust—they will be relatively stable and multiply detectable. Theories are tools for representing, explaining, and dealing effectively with Nature. If they deal whenever possible with objects and properties that are at levels, they will be simpler, and will deal with things that are stabler, and (for that reason), also more common and persistent.

3. The Coevolution of Levels and Their Entities

i. Richard Levins (1968) argues that *organisms evolve in such a way as to minimize the uncertainty in their environments.* This is an important truth—but only half of the story: organisms will try (1) to be as *unpredictable* as possible to their predators, while (2) trying to render the behavior of resources they need, including prey, as *predictable* as possible! This selection for unpredictability (together with selection to respond adaptively to energetically tiny informational cues in the environment) introduces a level of predictive complexity in aspects of the detailed behavior of biological systems that seems to have no parallel in the inorganic world.[28] These kinds of interactions should lead naturally to positive feedbacks, non-linear dynamics, and chaotic behavior. *This interdigitating web of designed predictabilities and unpredictabilities, together with the consequent selection for heightened sensory acuities, probably serve more than anything else to make the regularities of the biological natural order so conditional, so context-sensitive, and so complex.* It leads to the exploitation of sources of information, good predictors of fitness-relevant parameters, wherever they can be found—including at other levels of organization. *Thus organisms, just like human scientists, sometimes have reasons for developing interactions that are not level-bound, and these opportunistic inter-level connections make higher-level phenomena less well-defined with respect to level, and levels themselves more diffuse.* The fact that these trans-level interactions for such things as functional organization (Wimsatt, 2002) can themselves sometimes be described in a systematic way that is not level-bound is ultimately what makes what I describe as *perspectives* below so important for the analysis of biological systems.

j. More generally, considering Levins' original insight, *as stable foci of regularity and predictability, levels should act as attractors for other systems changing under selection pressures.* These evolving systems will do so by plugging into regularities in as many levels as are accessible to them—in effect by matching levels, where possible, with their environments.[29] When they do so, then their own regularities of behavior become part of the context to which other organisms adapt. This insight is a major feature in most or all concepts of the ecological niche (see Schoener, 1989, for a review), and is further generalizable.

k. *Levels themselves evolve over time,* with higher levels becoming occupied and lower levels becoming more densely occupied, while the biological objects comprising them and their interactions change on still faster dynamics. The temporal course of levels thus mimics the ecological phenomena of succession, and the stratified and rich ontology of the tropical rainforest rather than that of a Quinean desert. This is a perspective seemingly more appropriate to modern cosmology (which is a story of the successive occupation of higher and higher of the lower "physical" levels up through the atomic and molecular scale—and paradoxically, the differentiation of lower and lower of the higher physical levels on the astronomical scale) than it is to modern ontology, but it is also profoundly evolutionary. The level of organization is more like an ecosystem than a species—it evolves as a product of the evolutionary trajectories of the entities that compose it, and provides selection forces that guide their evolution (by affecting what is stable). *From the evolutionary perspective, levels define niches for their composing entities, but these are coevolving niches that are products of the entities that make up the levels.* (Compare the "constructional" view of the relationship between organism and environment of Levins and Lewontin, 1985; the concept(s) of the ecological niche by Schoener, 1989; and, for an important and instructive extension of the concepts of niche and species to the evolution of theories and research traditions, see Allchin, 1991.)

Note—as philosopher Chuck Dyke has urged upon me—that this last observation places an important constraint on the ways in which levels or their entities can be regarded as compositionally defined. In Section II I noted that while levels were compositional, this should not lead one to the mistaken view that the best way to make a higher-level entity (according to the engineering paradigm) was to assemble it out of lower-level parts. On the view advocated here, within the organic and social realms (I won't speak for large "merely physical" aggregates), levels are for many purposes co-evolved, generated, or developed, rather than aggregated. It is still true that in a relevant sense, any higher-level entity will be composed (without remainder—I still believe in the conservation of mass) of its lower-level parts, but it will be a (mechanically explicable) non-random generated complex of those or other lower-level parts, which may have required a diversity of "chaperones" (as molecular biologists call other molecules designed to facilitate a given reaction) and other same and higher-level co-generating complexes for its construction or development. But if this is true for many of the entities at a level, and if the entities at a level act as coevolutionary forces on one another, it is also true for the level itself, and the description of the level as a compositional entity will—to that extent—be misleading.

4. Levels, Robustness, and Explanation

l. There is a general *level-centered orientation of explanations* that can be explained in terms of the greater stability and robustness of entities at levels of organization, and probably more globally, in terms of the consequent robustness of levels themselves. This is a general and important meta-principle for the organization of explanations that is usually taken for granted and seldom commented on. It facilitates explanatory clarity, but occasionally misfires (see the discussion of perceptual focus in the last two sections of Wimsatt, 1980b, where I discuss the biasing effect of the tendency to refer group phenomena down to the individual level of description in the units of selection controversy). The robustness of levels tends to make them stable reference points that are relatively invariant across different perspectives and therefore natural points at which to anchor explanations of other things. *Explanations of the behavior of between-level entities tend to be referred upwards or downwards in level, or both—rather than being pursued in terms of other between-level things. Even the fine tuning of the exact "altitude" of the between level entity—its size and thus the distance it is above the lower and the distance it is below the upper levels—is motivated by concerns originating at one or the other of the levels. The robustness of levels makes the level-relativity of explanations a special case of the phenomenon referred to in the preceding section—the explanation of that which is not robust in terms of that which is robust.* I will consider the case of Brownian motion as a between-level phenomenon, which, by its very nature requires very special relations to the level below and the level above. (For a more technical exposition of some of the details, see Jeans, 1940.)

A good Brownian motion particle must be small enough that sampling error effects in molecular collisions produce temporally local imbalances in change of momentum between colliding molecules and the particle—giving net random fluctuations in the motion of the particle. In effect, it is enough larger than the colliding molecules that it jiggles relatively slowly (the law of large numbers works pretty well), but not so much larger that it works perfectly (that the jiggles are too small to detect). In a gas, the colliding molecules are moving at a mean speed equal to the speed of sound (of the order of 1100 ft./sec. in air at room temperature at sea level—so-called standard temperature and pressure). The Brownian motion particle must be enough larger than the gas molecules that individual collisions do not move it too fast or far before the next collision (or actually, the next significant failure in local averaging of collisions), so that we can continue to track it visually. Increased size of a particle (relative to its molecular drivers) acts in four ways to facilitate tracking: (1) it slows down motion in response to a collision with a particle of a given momentum; (2) the larger cross-section gives more collisions per unit time, giving temporal averaging in a shorter distance and decreases the expected absolute path length (or time) until the next perceived change in direction; (3) the increased size also decreases its *relative* path length (the ratio of path length to diameter), increasing the perceived *relative* stability of its position and motion—an important variable in our perceptual ability to track it; and (4) the Brownian motion particle also has to be large enough to reflect light in the visible spectrum, or else we couldn't see it (but if the particle is too large, it will not move enough for us to be able to detect the motion).

Individual "jaggies" in the Brownian motion particle's trajectory do not generally correspond to individual molecular collisions, but rather to local imbalances in collisions that force a distinguishable change in its velocity in times short enough to be perceived as instantaneous. Our visual system reifies paths between these super-threshold changes as straight-line trajectories, with piecewise constant velocities, but the value of that threshold is a complex function of illumination level, our static and dynamic angular resolving power, flicker-fusion frequency, and the wavelength of the reflected light—not to mention the magnification and optics of any instrumentation we use to watch it. (It is this fact that is responsible for the frequent claim that Brownian motion is a fractal phenomenon: changes in the magnification of the scene, or of the motion sensitivity characteristics of the detector will change the length scale over which velocity changes are detected.) If there are entities causing the changes in direction that *we* notice, *as we reify these changes,* they are *clusters* of collisions, rather than individual collisions, and the character and size of the clusters that we will reify as a group is a function of our perceptual parameters. (Other organisms would see it differently—possibly resolving a fractal pattern on a different scale determined by the relevant parameters of *their* visual systems.)

The colliding molecules are below the Brownian motion particle in level, and we are above it, but *there are no levels in between for the Brownian motion particles to occupy.* If anything is at its level, it is these clusters of molecules, whose grouped collisions cause noticeable changes in velocity or direction of the particle. We do not recognize these clusters as entities for at least two reasons: (1) the perceiverdependent and thus subjective time and size scale fractal characteristics of the Brownian motion—changes in which would change the temporal boundaries of the relevant clusters, and (2) the lack of unity of the cause of these motions—because the clusters are mere temporary assemblages that have no stability—they don't make "good" objects.[30] Explanations are, as here, referred downwards and upwards in level.

Another revealing indicator that Brownian motion particles are between levels is that they are given no intrinsic characterizations—as is indicated by the fact that things as diverse as dust motes and bacteria can all be Brownian motion particles. *Between-level entities tend to be defined functionally rather than in terms of their intrinsic properties—it is almost as if they have no intrinsic properties to use in such a definition.*[31] If so, this suggests the paradoxical conclusion that we may recognize the intrinsic properties of things, at least in part, due to characteristic interactions they have with other same-level things, since only levels have the intensity of different kinds of interactions among entities to fix unique sets of intrinsic properties as being causally relevant.[32] Multiple realizability in between-level contexts washes out the causal salience of most specific intrinsic properties.

m. It is also true that in our world, the dominant methodology is reductionist—we tend to explain features of the behavior of an entity in terms of its internal features, rather than how it relates to its environment. This implies a kind of explanatory priority, that *things not explicable at a given level are to be referred to the next lowest level, rather than to the next highest level.* This is a contingent, but very deep feature of our methodological world—sufficiently so that we tend to be suspicious when we are called on to explain phenomena by going up a level (as with functional explanations), or even by staying at the same level (as with phenomenological causal theories). These suspicions are frequently unjustified, and there are situations where explanations in terms of other same-level or higher-level entities are exactly what is required. Different aspects of the reasons for and character of this bias are discussed at length in Wimsatt, 1976a, part III; Chapter 11 in this volume; and Wimsatt, 1980b (the section on reductionist problem-solving heuristics and their biases), and I will not discuss them further here.

5. Time Scales, Multiple Realizability, Stability, and Dynamical Autonomy

n. As noted by Simon (1962), *processes at higher levels* (with a few important exceptions) *tend to take place at slower rates than processes at* *lower levels as measured by their "relaxation times"—the time it takes a reaction to go a certain fraction (usually one-half) of the distance to equilibrium.*[33] This phenomenon would certainly follow from the fact that it takes longer for causal effects to propagate larger distances. The coupling of size and time scale might look suspiciously like an application of relativity theory to physical processes, but it is not that simple. Most causal effects propagate at speeds that are a negligible fraction of the speed of light—governed by different processes that have more to do ultimately with quantum mechanics than relativity (the rate of propagation of disturbances of various energies in various solid, liquid, and gaseous media). Even if these processes are rooted in quantum mechanics, they would be so via pathways that—at least in the organic realm—are sometimes torturously indirect. (Consider the rate of propagation of membrane depolarization pulses in nerve fibers, and locomotion speed in all types of animals—both of which *increase* for larger structures, but in ways that lead to decreases in the frequency of repetitive actions for larger animals. Thus, an elephant runs much faster than a mouse, while its legs move at a much lower frequency. Bearing this in mind, I was astounded to discover that my expensive SLR camera did not have a lens speed fast enough to stop an ant in motion!) The net effect is to make one chary of any simplistic explanation for this probably very heterodox phenomenon.

o. *The multiple-realizability of higher-level properties or types is a general fact of nature, and applies to any descriptions of entities at two different levels of organization.* (it is thus entertaining to see philosophers of psychology act as if this characteristic is a special property of the mental realm). Multiple-realizability is entailed jointly by (1) the astronomically larger number of possible distinguishable micro-states than possible distinguishable macro-states—a ratio which (assuming that micro- and macro-variables have equal numbers of allowable states) grows roughly as an exponential function of the ratio of sizes of characteristic entities at the two levels, and (2) the numerical identity of the upper-level system thus described with the lower-level system thus described. Given that relatively many states at the micro-level must (because of the numerical identity) map into relatively few at the macro-level, the multiple-realizability of the few by the many follows (Wimsatt, 1981a).

p. More importantly, the *dynamical autonomy* of upper-level causal variables and causal relations—their apparent independence of exactly what happens at the micro-level—is entailed by this multiple-realizability and two further facts: (3) the relative stability of macro-level features (which persist for a characteristically longer time than micro-level features as a joint result of longer relaxation times and multiple realizability—items *n* and *o* above) in the face of (4) a constant flux of micro-level changes on a smaller size and shorter time-scale. (These items can be collapsed into a single assumption by taking the relative character of the stability claim seriously.) *The stability of macro-states in these conditions further entails that the vast majority of neighboring (dynamically accessible) micro-states map into the same or (more rarely) into neighboring macro-states.* To suppose otherwise would require at least a tremendously convoluted and radically improbable mapping from micro-states to macro-states—if it were even consistently possible. It is dynamical autonomy, more than anything else, which makes room for higher-level causal phenomena and theories, and the causal effectiveness of macro-level manipulations.

q. *Dynamical autonomy in turn entails that most (and in simple multi-level systems, an astronomical majority of) micro-level changes don't make a causal difference at the macro-level,* and that, except for cases of causal divergence (such as are found widely in chaotic dynamical systems, but are still presumably relatively rare since they would be selected against in most circumstances), most macroscopically causally efficacious factors will correspond to major global and often structural differences at the micro-level. The possibility of micro-level chaos shows that most macro-systems that show stability (or the respects in which they show stability) are tuned in such a way that the micro-level changes do not cause deviation amplifying (and therefore unpredictable) changes at the macro-level in those respects. In many simpler systems (for example, the mappings between micro-states and macro-states for a gas under conditions in which it does not show turbulence) we get this easily, but it applies to more complex systems as well if the systems are to show distinguishable macroscopic order.

An example may help, and we have a particularly important one at hand, for *the genetic system is a paradigmatic example of a striking kind of paradox frequently found in evolving systems. It is systematically tuned (as a matter of design) so that small differences can have effects on a variety of size scales including the very large, in which context dependence of effects is a common phenomenon, but where it is crucial that most differences do not have significant effects most of the time.* (I suspect that most people used to inter-level relations of the sort characteristic of classical statistical mechanics, where "law of large number" averaging is a reasonable mode of moving from one level to the next, will find the complex interplay of sensitivities and regularizing equilibrations of the relations between genotype and phenotype to be quite remarkable.)

Consider the following: We are given the genetic variability at many loci characteristic of virtually all species of organisms, and the scrambling effects of genetic recombination, so that each offspring is essentially without precedent in the specification of its genotype. Offspring of the same parents (save for identical twins) should characteristically differ at thousands of loci. Furthermore, we know that small genetic changes can and often do have large effects, and that interaction between genes in producing their effects is the rule rather than the exception.

Given these facts, if we didn't know any better, it would be plausible to expect almost no correlation in phenotypic properties between different members of a species (within the range of properties defined by that species), and between parents and their offspring. Yet offspring commonly inherit their parents' traits, as well as their fitnesses—not perfectly, but much better than random. The stability of the phenotype at many levels is essential for the heritability of fitness required for the evolutionary process to work. Not only must elephants breed elephants, humans humans, and *Drosophila Drosophila,* but the variability and systematic and independent inheritance of individual survival-relevant characters from parents to offspring within each species must be preserved—not glued together with a thicket of epistatic and linkage interactions—if temporally and spatially local adaptation to changing environments is going to be possible. We are constantly told by geneticists of cases where a single base change in a gene or a single amino acid change in a protein has enormous consequences for adaptation and function at a variety of higher levels of organization. But this has to be the exception rather than the rule for evolution as we know it to be possible. (Sickle-cell anemia remains the classic case here, and there still aren't many cases known as yet, though these should increase with our knowledge of developmental genetics.) Nonetheless, the plain fact remains that most genetic changes that happen under biologically normal conditions have no readily discernible effects (see Lewontin, 1978, on "quasi-independence," and Wimsatt, 1981b, for further discussion). Wagner (2005) provides a superb review and analysis of this "designed neutrality" at multiple levels of organization.

Therefore, most small micro-state changes do not make a difference at the macro level—even in systems that are characteristically sensitive *to small changes.* The converse does not follow: as pointed out above, closely related or identical macro-states may be realized by widely disparate kinds of micro-states, as illustrated by the Brownian motion of dust motes and bacteria!

r. *For instantiations of stable macro-level properties, in a sense there is no micro-level explanation for why they have happened, since changes in these properties, even if characterized at the micro-level, are macroscopic in scope.*[34] In giving extensive micro-level detail in an explanation, there is an implication that the detail matters—that the event or phenomenon in question would not have happened but for the cited details, that if just one detail were different, the outcome would have been significantly different. But if a process shows multiple realizability and dynamical autonomy this is just what is denied for the relation of most microscopic events to their macroscopic descriptions. There is, however, a crucial related question—namely, *why are these macroscopic states, properties, and relations stable?* This question will require an answer that is at least partially anchored in lower-level mechanisms—though not in large numbers of context-sensitive micro-level details. (If selection processes are involved in the explanation, it may also require reference to events at higher levels as well.)

s. The operation of *evolutionary and differential selection processes should tend to expand the scope of dynamical autonomy—increasing the range of multiple realizability—still further in cases where a macro-level property contributes positively to fitness.* Mutations will accumulate, which make its realization more likely and easier (this is a kind of generalized "Baldwin effect" response to selection). (Now, a decade after this was written, Wagner [2005] has provided robust empirical and theoretical support for this conclusion.)

Dynamical autonomy begins with the stability of properties of physical systems, but as the systems get larger and more complex, and their behavior more potentially variable, selection can breed stability of these usually more complex and contextual properties. Even in cases where the environment is unstable, making different properties desirable for fitness in different environmental contexts, evolution should select for context-sensitivity and conditional developmental programs—which tend to make the right things in the right contexts—all thereby increasing the heritability or stability of fitness across different environments (Wimsatt, 1986a). The only fly in this ointment is the increasing capabilities of the predators, parasites, and competitors of each species—referred to in item i above, and enshrined in Leigh van Valen's

(1973) Red Queen hypothesis—that even though each species is evolving, because of the co-evolution of others, you have to run as fast as you can just to stay in the same place. This should simply serve to generate increasing complexity and context-sensitivity of at least some organic interactions, and ultimately lead to the breakdown—through interpenetration and demodularization—of well-defined levels, and the emergence of other modes of organization in the ontology of complex systems.

One might think that one could go up indefinitely, successively aggregating and composing larger and larger systems into entities that occupy still higher levels of organization, but—whether as empirical fact, robust statistical regularity, or nomic necessity—other things emerge as salient cuts on natural processes and systems as these systems become more complex. My best guess is to think that the systems for which these other relevant modes of organization emerge are all products of biological or cultural evolution, since these are processes that tend to produce complex, contextually conditional, systematic, and characteristically adaptive behavior (see item i above), which has to simultaneously meet a variety of constraints at different levels of organization. But in lieu of more robust arguments for this conclusion, we must beware of overgeneralizing from the cases that our theories (and our interests) have given us the greatest reasons to consider. In the next section I try to characterize the conditions leading to the breakdown of well-defined levels and the emergence of perspectives. For an important review and analysis of levels and explanation in neuroscience that leads in complementary directions, see Craver (2007).

6. *From Levels to Perspectives: The Breakdown of Levels*

As long as there are well-defined levels of organization, there are relatively unambiguous inclusion or compositional relations relating all of the things described at different levels of organization. In that case, inter-level identificatory hypotheses are an important tool of explanatory progress in localizing and elaborating lower-level mechanisms that explain upper-level phenomena (Wimsatt, 2006a). There are relatively unproblematic assignments of all entities and properties with respect to level, and often systematic theories of phenomena at the respective levels. At this stage, theories are either directed to phenomena at specific levels or (for inter-level theories) acting to tie levels together by elaborating inter-level mechanisms or connections (see Maull, 1977; Wimsatt, 1976b; Darden and Maull, 1977). But, conversely, when neat compositional relations break down, levels become less useful as ways of characterizing the organization of systems—or at least less useful if they are asked to handle the task alone. At this point, other ontological structures enter, either as additional tools or as replacements. These are what I have called *perspectives*—intriguingly quasi-subjective (or at least observer, technique, or technology-relative) cuts on the phenomena characteristic of a system, which needn't be bound to given levels. Since the discussions of perspectives in Wimsatt (1974, included here as Chapter 9), and of the relation of levels, perspective, and causal thickets in Wimsatt (1976a), an even broader diversity of different perspective-like things have appeared in the literature of the last 30 years, and have been invoked to solve a similarly broad range of problems. This characterization of perspectives is tentative, incomplete, and still unsettled even on such major questions as to whether they are a unitary kind of thing. Nonetheless, there is a class of such things that do have a lot in common. Below I provide a tentative list of properties of these strange objects, and a set of examples suggesting some of their differences as well as their similarities. Further refinements will have to await another occasion.

The transitions suggested here—from levels to perspectives to causal thickets—characterize systems in terms of increasing complexity and context-dependence, and lower modularity and degree of regularity. This is an ordering in terms of kinds of complexity. It is not a natural evolutionary trajectory for systems, or any other kind of natural dynamical transition. Although, if I am right, systems later in this sequence first appear after systems earlier in the sequence (as a result of the continuing action of biological and sociocultural evolutionary and developmental processes), there are specifiable circumstances in which selection processes favor simplicity, modularity, near-decomposability, increased regularities of behavior, and well-defined compositional relations. Thus, with few exceptions, the order given here should be regarded as taxonomic, rather than temporal. Given the taxonomy, we may later wish to argue about temporal trends.

t. As higher levels get more complex (they have more degrees of freedom), they get more diffuse, and they overlap more in size scale and other related properties with neighboring levels, and engender perspectives and thickets. With more molar properties at the higher levels, and each one a potential pathway for causal interaction with entities that have or respond to that property, there are more ways to "plug into" a level. With more degrees of freedom, higher-level objects get potentially richer in their budget of properties, more multi-dimensional. At their best, they should thus be capable of higher degrees of robustness than lower-level entities. (There should be more ways of interacting with a spouse than with a quark!) There should *also* be more ways of being *not* very robust, of being only marginally connected to the causal processes of a level, and also more ways in which objects could interact simultaneously with or *bridge* two or more neighboring levels. These last two kinds of cases would increase the diffuseness of the levels associated with the entities. Thus, as levels get higher and more complex (up, roughly through the level of the ecological community or ecosystem, and perhaps on up to the biosphere), we should expect them to get more diffuse, for levels to overlap more, and for it to get more difficult to localize an entity or phenomenon by level unambiguously and for all contexts (see Figure 10.1).[35]

u. As objects find new ways to bridge levels, fluctuations at the lower level, which without the bridge average out at the upper level, are now transmitted directly (as we can observe Brownian motion with the aid of a microscope, and through that, the effects of micro-level events), generating the possibility of macro-level amplification of these micro-level events, creating a kind of sensitive dependence on initial conditions that will tend to increase the number of circumstances under which macro-level regularities will break down. Thus, we should expect that the maximum degree of regularity of upper-level phenomena for complex organized systems would be less than that for simpler systems composed of more homogeneous parts. This is the complement to the "diffusion" of levels: as they come to span a broader range of sizes, the maximum predictability decreases, almost as if the area under the level waveform for each level is a constant[36] (see level c in Figure 10.1). One must remember that small differences—fluctuations or signals—make a difference when they are detected by a system designed to respond to them, and for which the pattern is significant. The human eye can detect a single photon—a micro-level event to be sure, but not yet a pattern. The number of photons, if appropriately distributed in space and time, necessary to convey information is larger than this (probably of the order of 10), but still astoundingly small.[37] Detectable information can lead to macroscopically major (and, with modern technology, even further divergent) behavior.

v. At the same time we should also expect (ultimately, for reasons of increased dimensionality) to find more frequent, obvious, and severe context-dependence of the behavior of our entities at higher levels of organization. This would most often be expressed via systematic and not so systematic exceptions to simple generalizations involving these entities. This is one of the reasons why it is better to think of regulari-

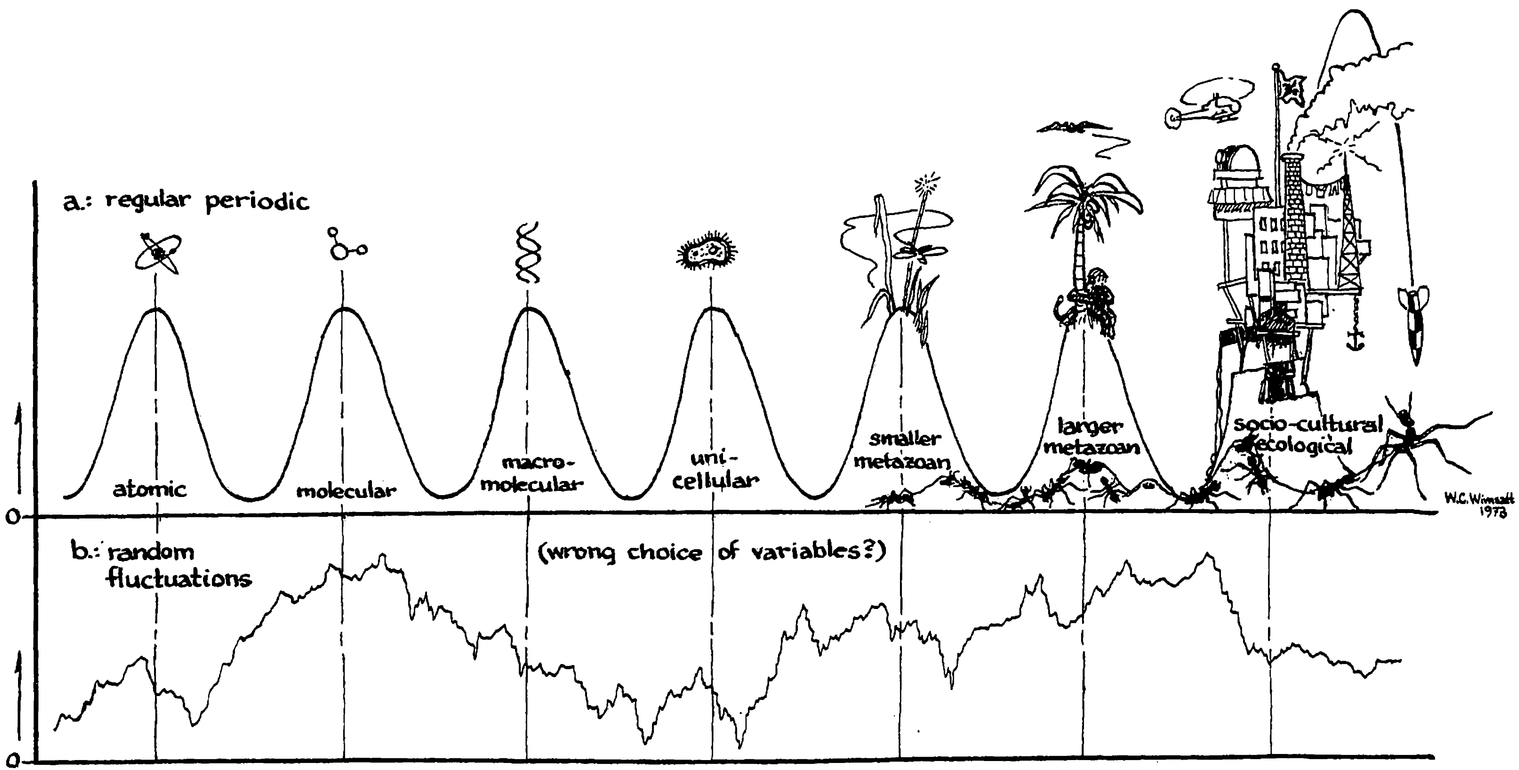
a.: regular periodic
atomic
molecular
macro-molecular
uni-cellular
smaller metazoan
larger metazoan
socio-cultural ecological
b.: random fluctuations
(wrong choice of variables?)
0
0
W.C. Wimsatt 1973

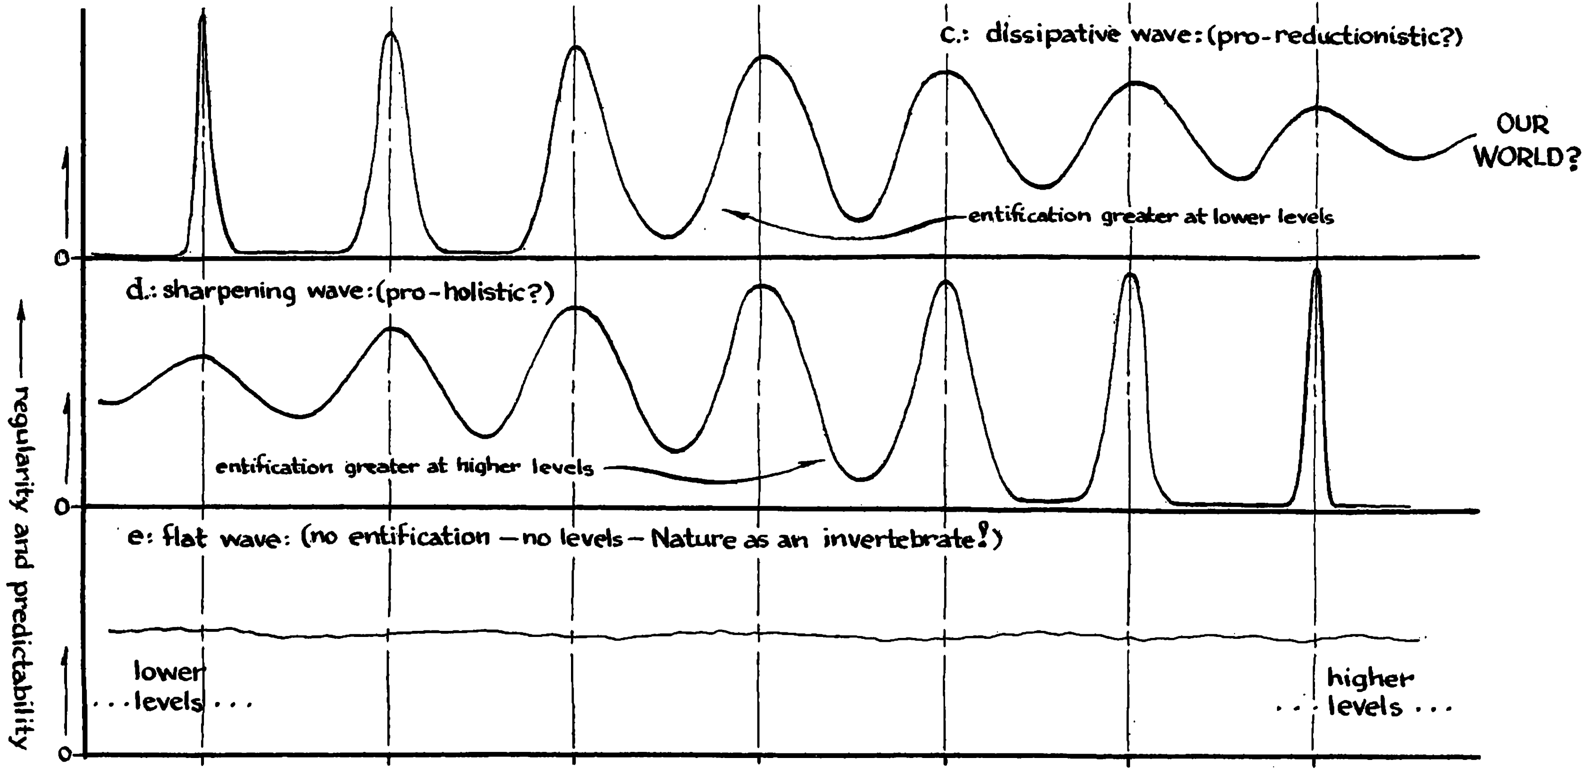
c.: dissipative wave: (pro-reductionistic?)
OUR WORLD?
entification greater at lower levels
d.: sharpening wave: (pro-holistic?)
entification greater at higher levels
e: flat wave: (no entification — no levels — Nature as an invertebrate!)
regularity and predictability
0
0
0
lower
... levels ...
higher
... levels ...

ties in complex systems in terms of mechanisms rather than laws. The latter, but not the former, suggests a search for exceptionless generalities and explanatory completeness, whereas the former fit naturally into a scheme that is satisfied by providing a characteristic *ceteris paribus* qualified articulation of causal factors (see Chapter 11).

w. Finally, as it becomes more common for entities to interact directly with other entities only through a subset of the properties that are causally relevant at that level, with different entities responding to different subsets, the notion of a *niche* (derived originally from ecology, cf. Schoener, 1989, and also as applied to research programs and theories by Allchin, 1991) becomes more relevant to the characterization of their behavior. This notion of a niche makes it clear and naturally explicable how different systems could act upon and react to the "same" environment in fundamentally different ways (Wimsatt, 1976a). The fact that the niche must be characterized relative to the organism—it is mutually defined by the organism and its "objective" environment (the

Figure 10.1 Waveform representation of compositional levels of organization as they might occur in different conceivable worlds—not all of which are physically *possible* worlds. In each row, the vertical axis is the degree of regularity and predictability—or, in more modern terms, the degree of pattern—for objects of different sizes. Size is represented logarithmically along the x-axis, so that regular periodic maxima would represent patterns found at geometrically increasing size scales. (Such scales would be expected if objects at each level were aggregates of roughly commensurate numbers of objects from the level immediately below.) It is argued in the text that the diagrammatic top row (a) and the second row below it (c) are the best representations of levels of organization in our world—(a) for its periodic character spilling over in an unruly fashion increasingly at higher levels, suggesting (c) for the greater diffuseness of the higher levels of organization (in the middle range of size scales that we occupy). The levels diagrammed here are really only the middle ones. Presumably, quantum mechanics renders the very small again diffuse, and astronomical scales again produce well-defined objects interacting in a relatively limited number of well-defined ways. I believe that waveforms (d) and (e) are *not* found in our world. As discussed in Wimsatt (1976a), a waveform like (d) would favor holistic over reductionist methodologies, and non-periodic forms like (b) or (e)—where there are no levels of organization—are ruled out by Simon's arguments concerning the role of evolution via stable subassemblies. Given the obvious existence of levels of organization over the range sampled and the random excursions in the sampled variable, the reasonable assumption for (b) is an incorrect choice of variables, or perhaps diagnosis of a causal thicket.

configuration of physical and biotic factors affecting its evolution)—introduces a feature of subjectivity that I explore further in the next section.

III. Perspectives: A Preliminary Characterization

What I call *perspectives* is a diverse range of things that nonetheless appear to have at least some of the properties of being "from a point of view" or to have a subjective or quasi-subjective character. In spite of that, perspectives differ substantially in terms of their other properties, and in terms of their relative objectivity. Their subjective character is because of the properties that they do share, which are discussed below. (The parenthetical remarks in these paragraphs are usually further elaborations of how this is so.)

1. *Perspectives involve a set of variables that are used to characterize systems or to partition objects into parts, which together give a systematic account of a domain of phenomena, and are peculiarly salient to an observer or class of observers because of the characteristic ways in which those observers interact causally with the system or systems in question.* (So far, this does not distinguish a perspective either from a methodological approach, or from the ecological niche of a species—two things that both have a kind of observer-relativity, and also have the curious objective-subjective duality I think characterizes a broad range of perspectives.)

2. *The set of variables in question is recognized not to give a complete description of all aspects of the systems that they are used to investigate.* Thus there is an explicit denial of a closure clause. (If this captures an important aspect of subjectivity, which I think it does, it is the recognition that it makes no sense to speak of something as subjective [or as objective] without the other category—which at this stage [from the subjective side] involves at least the recognition that there is something outside of the boundary of the subjective.)

3. In spite of this, there may be a restricted closure of the following sort: *there is a reasonably well-defined class of problems that can be solved without bringing in information from outside the perspective.* These are treated as paradigmatic problems for that perspective. These may also be problems that cannot (or cannot plausibly) be solved in any other way. So there are paradigmatic anatomical, physiological, and genetic problems, though (cf. (2) above), no one believes that these approaches individually exhaust what may be said about the organism. (This suggests a kind of unity and systematic problem-solving utility to

the subjective. There are things one can accomplish wholly within the subjective perspective, and things that only can be plausibly solved from within the subjective—or a particular subjective—perspective.) In effect, this says that perspectives partition problem-space in a nearly decomposable fashion.

4. Indeed, it is commonly taken for granted that *multiple perspectives can be applied to different aspects of the behavior of a system.* (Without this, there is not yet a recognition of the objective—a recognition of the robustness of the system accessed by the different perspectives.) I refrain at this stage from saying that the objective requires the existence of other subjectivities (thus perhaps characterizable as the intersubjectively accessible, or interpersonally robust), or merely the applicability of other perspectives (which could still be true in a Robinson Crusoe universe with plenty of external robust objects but no other persons).

I won't say any more here about the personal, interpersonal, and material realms, but to note that robustness, levels, and the idea of a perspective, together with an account of what it is to have a shared perspective are useful tools in characterizing our objective, mental, and social worlds. (In the last section of Wimsatt [1976a] I note and exploit parallels between the kinds of access we have to things at our own level and the less direct access we have to things at other levels, and the dichotomy between first-person and third-person perspectives.) But that is for another time and place. I want to consider particularly the kinds of complexities that make levels break down. The next two properties of perspectives were described more fully in Chapter 9 and in Wimsatt (1976a).

5. *Simple systems as well as complex ones can be described from a variety of perspectives, but will differ in the degree to which they have problems that are trans-perspectival—which require the use of information from more than one perspective for their solution.* Simpler problems are bounded and solvable with the resources of a single perspective. Simpler systems have more of their problems (or more of their problems for the purposes at hand) bounded within individual perspectives. Note that since problems usually arise out of purposes, a system can be simple for some purposes, and complex for others.

6. The complexity of trans-perspectival problems also varies from simpler to more complex with whether they decompose systems in ways that (a) are spatially coincident (in which case the different perspectives must also be either at the same level, or span the same range

of levels); (b) are hierarchically rationalizable relative to one another, so that the parts of one perspective are all whole systems in another (in which case the perspectives are related to one another as different level descriptions of the same system); or (c) overlap in arbitrary ways. The last case produces an enormous increase in complexity, but is common in the biological, psychological, and social worlds. (This is called *descriptive complexity* in Chapter 9, and the preceding kind of complexity is called *interactional complexity.*)

7. *Note that levels come out as a kind of special case of perspectives on this analysis—a class of perspectives that map compositionally to one another so that their entities are related without cross-cutting overlaps in a hierarchical manner.* It is tempting to say that we need to require also that the entities/parts at levels are especially robust, though that may come out for free given that hierarchical (and modular) compositionality will tend to require or entail substantial robustness of the systems and parts at all levels. Note that thus far, I have introduced nothing that a hard core materialist could not accept. (Indeed, I believe that all that I have introduced so far a hard-core materialist *must* accept.) Given this, hierarchical compositionality suggests a number of further interesting (but at this time still speculative) connections: (a) The "nearly sealed" aspect of living at a level of organization (the fact that level-leakage is relatively rare), and the comparatively torturous and indirect paths to systematic access to another level can at least help to explain qualitatively the first-person/third-person dichotomy between subjective and objective modes of access indicated in Wimsatt (1976a); and through that (b) it may suggest naturally how subjectivities can be seen to be anchored in a natural world. (c) Also, if "level leakage" is just a variety of "perspectival leakage," it suggests that and how modest amounts of comparability or leakage between subjectivities may be essential both to the recognition of other subjectivities and the reality anchoring of our own (necessitated by the private language argument). It also predicts (d) that, how, and why the breakdown of levels with increasing complexity can come to create problems for the localization and bounding of subjectivities as well as for the bounding of well-defined perspectives. This latter problem I take to be connected to new wave contextual embedded and distributed theories of consciousness.

I now wish to consider perspectives that are not levels. They may fail to be levels either because they are too small, they are located mostly at levels, or because they aren't of sufficiently broad span to count as levels. Or they may fail to be levels because, in a way, they are too big—they cross-cut levels: they are transverse sections that do not include more than a small fraction of the phenomena at any given level, but span phenomena at more than one (usually at several) levels. It is these two basic kinds of entities that allow us to go beyond levels to importantly different kinds of entities.

8. *The smaller kind of perspectives are those things that look most subjective, since they are most explicitly keyed to the point of view of a particular kind of organism or observer.* When objectively characterized without regard to other than physical or biological properties, I call these *niches,* because I think that the ecological niche of a biological species is the prime exemplar of this[38] (on niches, see Schoener, 1989). When characterized explicitly cognitively and subjectively, with respect to the cognitive and sensory capacities for and from the point of view of an animal, I call this the *subjective niche,* or *Ümwelt,* to use von Üexkull's (1934) term (von Üexkull and Nagel, 1974, are the best, and remarkably close, exemplars of this position). This notion of perspective naturally suggests further subdivisions that are psychological or cultural rather than biological in character, and how to make these further subdivisions (and how many) is an important question, although I do not address it further. Is there a paranoid schizophrenic's perspective? One or many? Is there a female perspective (is it cultural or biological)? Is there a feminist one? An upstate New Yorker's or a Manhattanite's perspective? An only child's (first child's, second child's, etc.) perspective? Does each new interest group or reference group individuate a perspective or a component of a perspective? Does every person? Does every life stage? How has my perspective changed since I was an assistant professor? Got married? Became a father? Learned how to program in *Pascal?* Even we must follow Quine's desert intuitions for ontology here in recognizing that there may be too many potential perspectives standing in the doorway! So how should we decide? This domain is the topography of many of the most important battles in the social sciences.

9. What I called perspectives in my 1974 paper is not usefully captured by any of the notions of perspective discussed so far. It is the *larger* variety of perspectives promised above. It is a more robust ontological category than they are, since it is not essentially defined by the relationship of a single kind of entity with its environment. Perspectives in the 1974 paper (a) spanned more than one level, and thus could not

be ordered as higher and lower or more primary and secondary than one another; (b) gave criteria for decomposing systems into parts using the properties and tools appropriate to that perspective; (c) were manifestly incomplete descriptions of their objects; (d) were such that different perspectives (for complex systems) could cut up systems in quite different ways that were not easily comparable to one another; (e) had a class of problems that they could solve in isolation; and (f) (for complex systems) had other problems that could not be solved without bringing in the resources of another perspective or perspectives.

Anatomy, physiology, and genetics are different perspectives on an organism in this sense. Perspectives may sometimes correspond loosely to disciplines, but need not. They may be either larger or smaller. Thus, the adaptationist perspective, in which the parts of an organism are all analyzed in terms of their evolutionary function—those aspects of behavior responsible for their selection, elaboration, and maintenance—which is larger than a discipline, unless disciplinary lines are drawn extremely broadly to include it; the discipline of evolutionary biology, for example. Fate maps also seem plausible as perspectives in this sense, in which the cells of a developing embryo (or layers, or regions—so this is not confined to a level) are marked to indicate what they will become are a specialized representational tool within classical developmental biology, and thus much smaller than a discipline. There are specialized tools for revealing these (such as radioisotope labeling, which can give an iconic representation of the fate of a cell, layer, or region through development).

If I were to rename this kind of perspective now (as I probably should), I would call them *sections*—short for cross sections (or perhaps sometimes transverse sections in messier cases!)—views chosen by architects, engineers, and anatomists to give particularly revealing aspects of their complex structures; *views that can cross-cut one another in various ways, and at various angles; views that are individually recognized as incomplete; views that may be specialized for or better for representing or for solving different problems; and views that (like perspectives) contain information not only individually, but also in how they articulate.*

10. Important ontological features of perspectives are captured in Figure 10.2,[39] which indicates that *perspectives cannot be ordered compositionally relative to one another—you cannot say that the objects or parts of one perspective are "really" composed of the objects or parts of another—or if you could do so, that a corresponding claim could be*

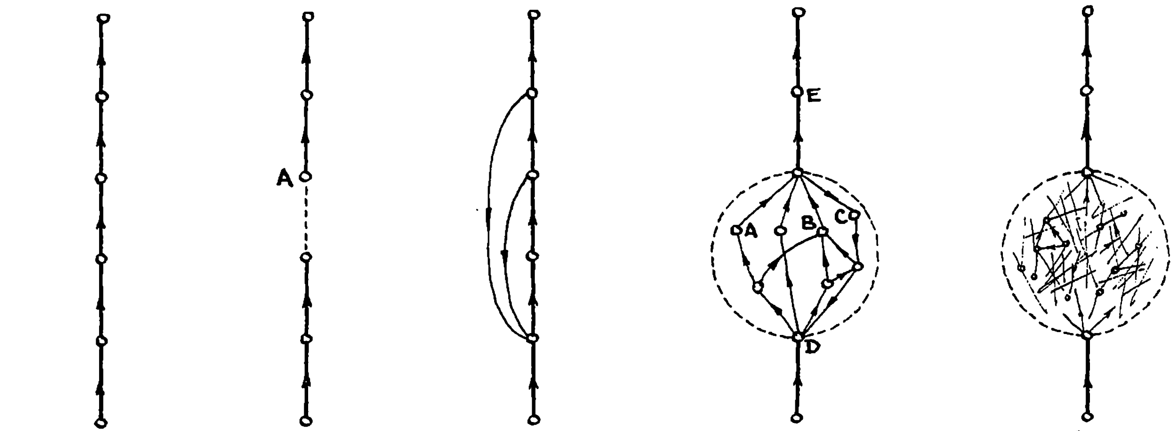

(a) simple reduction

(b) emergent at A.

(c) selection and explanatory "feedbacks"

(d) descriptive and interactional complexity

(e) defineable perspectives break down.

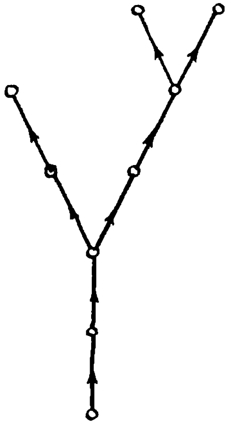

f) divergent branching yielding diverse products.

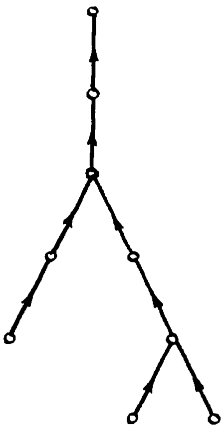

g) convergent branching yielding heterogeneous products.

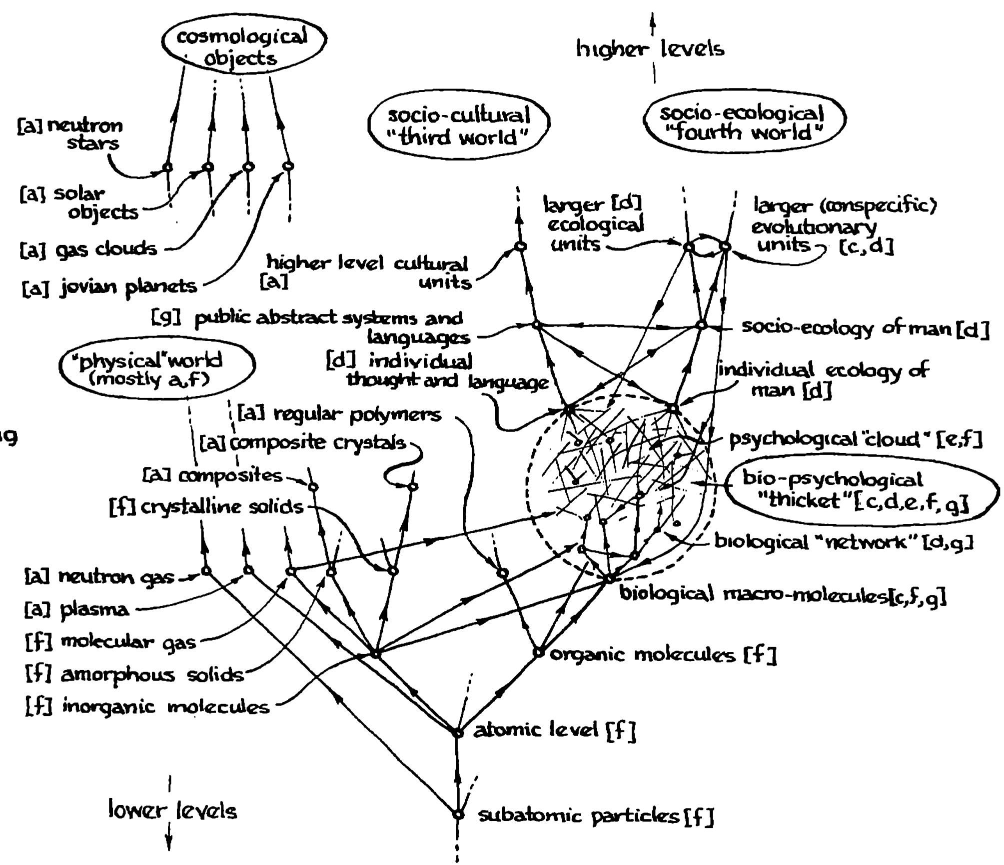

(h) A reductionistic (?) illustrative phylogenetic ontology of our world as we see it.

made in the other direction with equal justice. (Are anatomical features composed of physiological processes or conversely? The question doesn't make sense, but information from each perspective is relevant to the solution of at least some problems in the other.) But compositional talk is not forbidden within perspectives, even putting levels

Figure 10.2. Complex orderings of levels and perspectives. This figure depicts modes of composition of aggregate and complex systems, ordered in terms of the direction of explanatory relations. (Nodes are levels or perspectives; arrows give direction of explanation.) In simple systems this follows compositional relations, with behavior of the wholes explained in terms of the properties and relations of the parts. Thus, the simple reduction picture of the "unity of science" movement is given by (a), in which each level explains the one above it—a picture which, as Roger Sperry complained, "seeks to explain eventually everything in terms of essentially nothing" (quoted in Wimsatt, 1976a). The classical picture of emergence (as a failure of reduction) introduces a gap, as in (b). (This account is rejected in Chapter 12.) Explanatory feedbacks from higher to lower levels are introduced by selection processes (Campbell, 1974b), diagrammed in (c). Complex organization of the phenotype (as a product of selection processes) builds on explanatory feedbacks from higher to lower levels, creating further ordering problems with the emergence of perspectives (d), and increased interactional complexity producing cross-perspectival problems, and ultimately breakdowns of and ambiguities in the boundaries between perspectives, resulting (e) in "causal thickets" (Wimsatt, 1976a).

Section (h) of the figure is a compound diagram illustrating the composition modes of various kinds of physical, biological, psychological, and social systems (letters in square brackets refer to the local character of the network around that node). It is illustrative: I will not argue for its detailed architecture, and it may be wrong in representing complex physical systems in the biological, psychological, social, and cultural realms. I see no obvious errors, however. The biological organism (a developed language using socialized human) has perspectival structure (actually at its lower levels of biological organization, merging continuously with causal thicket structure as we get into the internalized psychological and social realm). Two ontological lineages emerge from this: those of cultural objects (abstract objects, presumably also viewable as abstract relational properties of objects in the second lineage), and socio-ecological objects (kinds of complex material systems having the whole range of social, ecological, biological, cultural, and psychological properties). I believe that the connectivity patterns relating these various realms inside and outside the individual are much more complex than represented here. Thus, social institutions obviously are complex hybrids of objects at a variety of levels from both of these lineages. There are causal thickets above and outside the individual interacting rather directly in various ways with causal thickets inside the individual, and an embodied socialized theory of consciousness is required.

aside as a special case. For a perspective, you may (and usually will, if you are a materialist!) be able to find lower-level objects (indeed, a greatest lower bound, or *GLB*, of largest common parts) such that all the entities in the perspectives are composed of them—de facto atoms, as it were. There may also similarly be higher level objects (correspondingly, a lowest upper bound, or *LUB*, of smallest common systems), such that the objects in the perspective are all parts of those objects (e.g., organisms), but for the regions in between the GLBs and LUBs, there is at most local orderability of compositionally ordered parts *within* each of the perspectives. If they exist, GLBs and LUBs of a set of perspectives are rich in implications. The GLB and LUB decompositions of embedded and embedding systems will both be robust, because they will be level-descriptions and will be orderable relative to one another. Given their unambiguous robustness and status as entities at levels, there will be a tendency to regard them as more important or as ontologically more central than descriptions derived from the perspectives in between. Reductionists will tend to favor the GLB descriptions on down through lower levels of organization, and functionalists or holists will tend to favor the LUB descriptions, and possibly on up through higher levels. If we accept the objectivity of the GLB and LUB descriptions, this will *tend*[40] to fix all the perspectives between them within the objective realm (or, more generally, to give them any ontological properties common to the two bounding levels). If so, these properties will be aggregative rather than emergent properties for that class and within that range of descriptions of systems (see Wimsatt, 1986a, and Chapter 12 in this volume).

The most interesting thing about perspectives follows from this ontological feature:

11. If compositional ordering relations break down, as they may between descriptions of the same object in different perspectives, then above a GLB and below an LUB traditional formulations of materialism are inadequate for ontological reasons *because you can't say what is composed of what,* although your complex system contains nothing immaterial. If this is right, then in that interesting size range in between atoms and organisms (or perhaps in many regions in between atoms and societies) you will often find a situation for parts or properties where neither type-identities nor token-identities appear to be of much use (Wimsatt, 1976a). Token identities aren't of much use anyways, beyond expressing advocacy of a token materialism. Nancy Cartwright said in a recent lecture[41] that token identities are too weak—they ignore the systematic regularities that are there, even in messy

cases. The problem (as she also noted) for type identities (and also for laws as they are normally conceived of by philosophers) is that the systematic regularities aren't exceptionless either. And you can't make them exceptionless without introducing so many qualifications as to make them essentially useless (I urged similar views in 1976—see Chapter 11). But we can't even get to this juncture if we can't specify composition relations, and in this interregnum of multiple partial incomplete perspectives, we can't. This might seem to be the death knell for any possible reductionisms—as it clearly is for any formalist or deductive accounts of reduction. It is also clearly at least highly problematic for any identity-based accounts like that urged in Wimsatt 1976a and 1976b (see also related discussions of what happens when localizations break down in Wimsatt, 1974, and in greater depth—using connectionist models as an example—in Bechtel and Richardson, 1993) But note also that this breakdown occurs without really doing anything to compromise the spirit of materialism because we can understand *in materialistic terms* why compositional relations are problematic, and a variety of general structural and methodological features about the situation, and can do so without admitting that there are any phenomena (or regularities) that we cannot explain. This is a remarkable situation, but one that characterizes, for at least some problems and properties, all naturally evolved systems.

12. But even with the varieties of incomparability suggested above, *organisms can share dimensions of niches in that some causal factors can be causally important to all, or to an important subset of them. This makes these dimensions or causal factors particularly important in explaining their behavior, and also particularly real, objective, and robust.* One way for perspectives to emerge (in the sense of sections, above) would be around causal clusters of variables that are robust niche dimensions—as sets of descriptive variables whose analysis generates adequate solutions to classes of correlative problems. The primary qualities would be good examples here, and statics (for physical structures) and anatomy (for biological ones) stand as good correlative perspectival theoretical structures. Within biology, similar or shared niche dimensions may be important causes of convergent evolution.

Finally, insofar as a theory deals with only a subset of the causally relevant properties of an object, it has a perspectival character, but if the properties it deals with are sufficiently robust and fruitful, it may be easy to forget this fact. It is worth considering (on another occasion) whether and when theories (in general or in particular; folk psychology or Kuhnian paradigms) should be viewed as being or as providing a perspective.

13. How do we judge whether perspectives are real? I think that there are two ways. First, when there is agreement across perspectives in identifying or saying things about objects they access in common, this judgment not only recognizes the robustness of the object, but—indirectly—confirms the means of access. Second, we can treat the perspective as object, rather than as means of access to other objects. But then the same criteria of robustness should apply—the extent to which the perspective is multiply detectable, in this case by being articulatable with other perspectives, affects the degree to which it is real to us. We can do this in various ways, with different ways appropriate to the kind of perspective it is. I will mention only one here, because it is already commonly recognized in methodological discussions in the social sciences. It is this activity we are practicing when we practice *Verstehen* (seeing the act in the way the agent did, and judging it to be rational or otherwise explicable from their perspective) to understand action. The target here is not action, or its justification, but the explanation of the action. And we can provide an explanation by putting ourselves in the other agent's shoes, and see that the action *is* rationalizable from their perspective. (Of course, it doesn't justify the action, the perspective could be that of a heinous fellow.) If we understand the action, in the sense of explaining it, then by taking on the perspective, and successfully practicing *Verstehen,* we have not only explained the action, but also confirmed the existence of that perspective, and its salience to the action.

IV. Causal Thickets

I noted above that each perspective will tend also to contribute to the solution of some problems that it cannot solve by itself—and that for more complex systems, this would tend to happen more frequently. With increases in the complexity of objects, and in their number and variety of degrees of freedom, they can interact with one another in more varied and complex ways, and more problems involving their behavior require the use of two or more perspectives for their solution. Sometimes, when there is a range of problems that can characteristically be solved using two or three particular perspectives or disciplines together, a new subdiscipline gets formed (e.g., psycholinguistics, or even developmental psycholinguistics). Sometimes problems are fought over by practitioners from two or more different perspectives. And sometimes problems appear to be big enough, or generally enough stated (e.g., the mind-body problem), that they seem to be intrinsically multi-perspectival. Since a perspective maintains its identity in part by having problems that its corresponding discipline can characteristically solve by itself, the characteristic identification of important problems with certain perspectives and the identity of perspectives tend to break down simultaneously. When the relative frequency of such problems gets too high (either as a function of the way the world is or as a function of the inefficiency of our conceptualizations in organizing our problem-structures), the boundaries of perspectives begin to break down and it becomes more difficult to decide which perspective (or perspectives) a problem belongs to. (Correspondingly, as the preceding parenthetical remark might suggest, it becomes harder to tell when we are talking about our world and when we are reflecting only or primarily on our own conceptualizations. Thus the "perspective," and many of the claims of the new deconstructionists and sociological relativists are in a way predictable and explainable in this situation—which I remind you, is still characterizable within a broadly materialistic perspective!)

This breakdown of boundaries induces competition among the different methodologies associated with the different perspectives, and so we should expect that methodological disagreements would proliferate, along with disputes about how to fragment systems into parts and how to best define key terms. As the boundaries break down this far, not only is it true that others' perspectives intrude on the one you wish to argue for, but also that your perspective can seem to reach legitimately to the horizon. Paradoxically, as the perspectives weaken in their own domain, they don't retreat, like good scientific theories, but their generality appears to increase without bound. (Deconstructionism is not the only banner to have claimed the whole field—witness methodological individualism under the banner of rational decision theory [fighting mostly prisoners' dilemmas], or the self-reinforcing behaviorisms of a generation ago.) At that point, philosophers may rush in where scientists fear to tread—or perhaps have done so and stubbed their toes! Here, if anywhere, philosophers may be useful if they know the lay of the land.

Perspectives have now degenerated into a *causal thicket.* This term is intended to indicate a situation of disorder and boundary ambiguities. Perspectives may still seem to have an organizing power (just as viewing a thicket or shrub from different sides will reveal a shape to its bushy confusion), but there will be too many boundary disputes. Claims may be made that phenomena are at a given level, or are to be viewed from a given perspective, and any level of analysis or perspective that has successful associated theories will attempt to claim disputed territory. But that is just the point—there will be a lot of disputed territory—and the disputes will often turn on how the system is to be cut up for analysis—or even (to those of a holistic persuasion) whether it can be cut up for analysis at all. (Some connectionists seem to expect that local analysis will fail for all interesting mental properties, which will therefore be holistically distributed, while others are busy denying that we will have to recognize any mental properties because they don't find them at any locations!) Most complex biological problems involve levels, perspectives, or a combination of both—except in neurophysiology and some areas of developmental biology. The neurophysiological, psychological, and social realms are mostly thickets, which are only occasionally well-ordered enough for local problems to be treated as perspectival or level-relative problems. All of this enormously complicates talk of reduction, because with such multiply connected entities, and the failure of the ability to say what is composed of what, it may now be almost impossible to determine what is being reduced, what is doing the reducing, and what even is the proper scope of the system under analysis and the problem we are being asked to solve.

The proliferation of disputes of this form involves an unusually large proportion of conceptual issues, methodological arguments, and boundary disputes. This phenomenon is predictable simply from looking at the form of complexity such systems take, and the form disputes should take when boundaries break down. Some of these disputes are likely to indicate sources of genuine disagreement, but this can't be determined when so many things are up for grabs. Moreover, the natural tendencies of most theorists toward expansionist territorial claims, and of all of us to understand the merits of our own positions better than those of our opponents, makes frequent disagreements seem inevitable where there are boundary ambiguities. Localization of problems with the existing conceptual structures, and of disputes to the right trouble spots will have to await the development of conceptual structures, methodologies, and new explanations of mechanisms in terms of them. If this explanation for their occurrence is correct or nearly so, an unusually large fraction of the disputes should be resolvable as people from the different groups learn and work out how to talk with one another, if (and it is a big sociological *if*) they maintain a commitment to try to understand one another rather than bloating their reputations by taking cheap shots at the opposition. This is perhaps the deepest pragmatic commitment of science—that it is in one's interest to come to understand differences, and then to resolve them. This yields an ultimately realist picture only because the world has an indefinitely large number of constraints for acceptable theories, if you know where to look. But you'd better get an overall sense of the geography before you decide on your colonizing strategy. This has a lesson as well, of which eliminativists should beware: you don't make friends with the natives (folk) by denying their legitimacy (psychology), and you can't tell what's in the territory without a native guide. You can play imperialist without heeding these warnings, but it usually requires more resources, costs a lot more, and takes a lot longer. And you may end up having to grant them autonomy anyway!

So far, we seem to have defined causal thickets as a kind of wastebasket category. They needn't be. On a priori grounds, considering the possible connectivities of causal networks, shouldn't causal thickets be the norm, and relatively insulated levels or perspectives the rare cases? Wouldn't causal thickets be, as it were, the high entropy or generic states of the causal structure of the universe—sort of an ontological primal slime? This is to exchange assumptions of simplicity and order in the universe for assumptions of randomness in causal connection—a kind of structural disorder. An absurd view, one might say, but not a priori absurd. To be sure, we wouldn't exist, and couldn't survive in such a universe, but considering it provides a useful kind of change in perspective. One of the remarkable things about our universe is the degree of order we find in it. To be sure, it is not an exceptionless static order—crystalline without flaw. There are regularities at all levels, and mechanisms tying them together, and perspectives that give cross-sectional cuts on the phenomena for a range of problems. And then there are some things that are just too multiply-connected to fit exhaustively into any of these ontological categories. And we can say something about the conditions in which we expect each of these to arise, and their methodological consequences. This looks a lot more complex than the old story, but it provides tools and ways of thinking and talking that seem a lot closer to the truth. And, as I've been trying to tell you, that's the way the world is.