Chapter 14: Mechanistic Levels, Reduction, and Emergence[1]

Mark Povich and Carl F. Craver

1. Why Levels?

In *The Sciences of the Artificial*, Herbert Simon (1969) offers the parable of the watchmakers, Tempus and Hora. Tempus builds watches holistically, holding each part in place until it forms a stable whole. Hora builds watches modularly, first assembling stable components and then organizing them into watches. Each is interrupted now and then. When Tempus is interrupted, she loses all her work on the watch and has to start again from scratch. When Hora is interrupted, she loses only her work on a single component. Hora thrives; Tempus struggles. The Moral: Evolved systems are likely to have (nearly) decomposable architectures. Their working parts are likely to be organizations of working parts, which are themselves organizations of parts, and so on. In other words, evolved systems likely exhibit regular mechanistic organization at multiple levels.

One might object to this argument on the ground that watches do not reproduce or that biological evolution does not assemble components sequentially (Bechtel and Stufflebeam 2001). As Simon points out, this would miss his point, which has now been repeated across many domains of science and philosophy: viz., that nearly decomposable (or modular) systems are more stable than holistic systems; they are to be expected when we find improbable stability in the face of random challenges. Contemporary work in network analysis, for example, shows that modular networks (those with high intra-modular connectivity and sparse inter-modular connectivity) are more "robust" than networks that are less modular; e.g., their mean minimum path length remains low in the face of random attacks. This is important, for example, in telecommunication, air traffic control, and neural networks (Albert et al. 2000). Such "modularity" is a widely accepted principle of both development and evolution (see Schlosser and Wagner 2004). Systems with causally independent parts are expected, in part, because nearly decomposable components can be independently assembled, regulated, damaged, and repaired without disturbing the other components. In short, everywhere we look from abstract theoretical considerations to concrete details of biology, we find reasons to expect dynamic but stable systems to be arranged in nearly decomposable hierarchies. Where we find embedded decompositions of working parts within working parts, we find nested mechanistic explanations.

Indeed, the assumption of near decomposability underlies the strategy of reverse engineering, of discovering how something works by learning how its parts interact. Kauffman (1970) calls this practice "articulation of parts explanation;" Haugeland (1998) calls it "explanation by system decomposition;" Cummins (1975) calls it "functional analysis;" Fodor (1966), Craver (2007), and others (Machamer 2004; Glennan 2002; Menzies 2011) call it "mechanistic explanation." Mechanistic explanations explain a phenomenon by situating it in the causal structure of the world. In constitutive mechanistic explanations (as opposed to etiological and contextual mechanistic explanations; see Craver 2001), one looks to a lower level, within the phenomenon, to reveal its internal causal structure. In many such systems, one can explain the behavior of the whole in terms of the organized behaviors of its parts, and one can explain the behaviors of the parts in terms of the organized behaviors of their parts. Levels, on this view, are simply embedded mechanistic explanations.

In this essay, we explicate this mechanistic view of levels (Section 2), contrast it with other senses of "level" (Section 3), and sketch its implications for emergence (Section 4) and reduction (Section 5), for the ontological status of higher-level phenomena (Section 6), and for thinking about the lowest level(s) in such hierarchies (Section 7).

2. Mechanistic Levels

We use the term "mechanism" permissively to describe causal systems in which parts are organized such that they collectively give rise to the behavior or property of the whole in context (cf. the notion of minimal mechanism in Chapter 1). As mentioned above, not all mechanisms are modular at intermediate levels. In regular networks of simple nodes, for example, there is no nearly decomposable structure between the behavior of the mechanism as a whole and the activities of the parts. All mechanisms involve the organization of interacting parts, but not all mechanisms have intermediate levels of organization.

Mechanisms so construed contrast with aggregates (see Section 3 below) in that mechanisms are organized. The parts of mechanisms have spatial (e.g., location, size, shape, and motion), temporal (e.g., order, rate, and duration), and active or otherwise causal (e.g., feedback or other motifs of organization; see Levy and Bechtel 2013) relations to one another such that they work together. This is why mechanistic hierarchies are often called "levels of organization." Aggregates, in contrast, are the lower limit of mechanistic organization: No between-component interactions are relevant to the true aggregate.

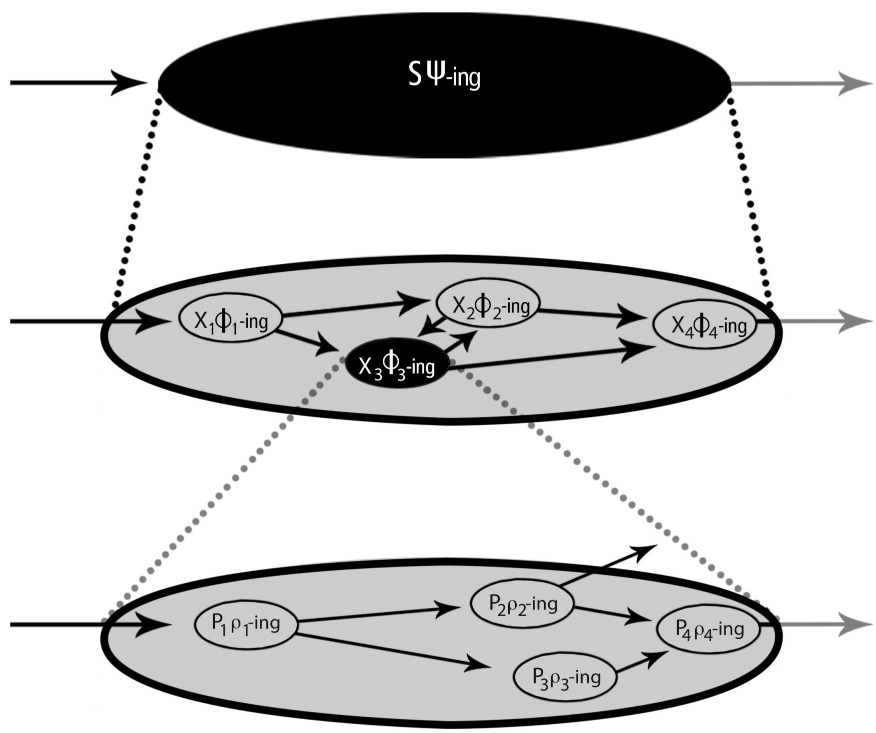

Figure 14.1 depicts three mechanistic levels. At the top is the activity of some mechanism as a whole (S's ψ-ing). S's ψ-ing includes the mechanism as a whole (e.g., the protein synthesis mechanism) doing what it does (e.g., synthesizing proteins). In this diagram, ψ is the topping-off activity of the mechanism; it is the activity to which all of the lower-level components are relevant. Beneath S's ψ-ing are the activities (ϕ_i) and components (X_j) organized such that together they ψ. Beneath that is an iteration of the levels relation, in which one of the X's ϕ-ings is decomposed into the organized ρ-ings

(rho-ings) of Ps. As Simon suggested, this process of decomposition might continue until we run out of either relevant or practically salient decompositions. Individually, the X_j and ϕ_i are organized to compose S's ψ-ing. Collectively, the X_j and ϕ_i exhaustively constitute (or, as is often said, "realize") S's ψ-ing in context.

X's ϕ-ing is at a lower mechanistic level than S's ψ-ing if and only if X and its ϕ-ing are component parts and activities in S's ψ-ing. A component is a part whose activities contribute to the behavior of the mechanism as a whole. The components (X_j) are spatially contained within S (because components are necessarily parts of the things they compose). Furthermore, each component is relevant to S's ψ-ing: X's ϕ-ing should contribute to (or be operative in) S's ψ-ing (see Martin and Deutscher 1966).

There is some debate about how this interlevel relevance relationship should be understood. Craver (2007) defines it in terms of the mutual manipulability of the whole and the part: X and its ϕ-ing are component parts and activities of S's ψ-ing if one can manipulate S's ψ-ing by intervening on X's ϕ-ing, and one can manipulate X's ϕ-ing by manipulating S's ψ-ing. This is a sufficient, but not necessary, condition for componency, as there might be parts of a mechanism that do not change during its operation (consider the walls surrounding pistons in a car engine) or redundant parts that can be independently manipulated with no effect on the behavior of the mechanism as a whole. Couch (2011) argues that componency should be understood in terms of Mackie's notion of an INUS condition. Components are Insufficient and Non-redundant parts of an Unnecessary but Sufficient set of contributors to S's ψ-ing. Still others (such as Harinen 2014 and Romero 2015; see also Craver 2007) understand constitutive relevance as causal betweenness: a component of a mechanism is causally intermediate between the mechanism's start conditions and its termination conditions (see Chapter 9 on parts and boundaries of mechanisms).

Furthermore, Figure 14.1 is ambiguous: Two distinct relations are represented in both the top and bottom (corresponding roughly to what Gillett (2002) calls dimensioned and flat realization). The first is a part-whole relation between S's ψ-ing and each of the X's individual ϕ-ings. We use the term "mechanistic level" exclusively for this relation. The second is a whole-whole relation between two ways of describing the same thing. We can talk about S's ψ-ing or we can talk about the organized collection of ϕ-ing Xs. These are two ways of describing one and the same object. Likewise, at the bottom of the figure, we can talk about a given X's ϕ-ing or we can talk about the organized collection of ρ-ing Ps in virtue of which X ϕs. Again, these are two levels of description that apply to one and the same object. Marr's levels of realization are like this: the computational level, the algorithmic level, and the implementation level are three ways of describing one and the same whole (Marr 1982; see also Chapter 30). We return to this issue in Section 6.

The above explication of mechanistic levels has several conceptual benefits. First, it accurately describes the multilevel explanatory structures one finds in biology and other special sciences, and makes explicit the kinds of evidence required to evaluate such explanations. Second, it satisfies many of our pre-analytic intuitions about the levels of explanation. Things at higher mechanistic levels are typically larger (and necessarily no smaller) than things at lower levels because the latter are parts of the former. Things at lower levels often take less time than things at higher levels because activities at lower-levels compose activities at higher levels. And the idea of mechanistic levels captures the common idea of "levels of organization" because higher-level mechanisms are made up of organized parts and activities. Third, this understanding of levels clarifies why interlevel causation is at least *prima facie* mysterious: causation between mechanistic levels must involve parts interacting with their wholes (see Section 4).

Finally, this understanding of levels helps to undermine the thought that nature can be usefully divided into monolithic levels of, e.g., atoms, molecules, cells, organs, organisms, and societies (Oppenheim and Putnam 1958). For mechanistic levels, there is no unique answer to the question of when two items are at the same mechanistic level. There is only a necessary condition: X's ϕ-ing and S's ψ-ing are at the same level of mechanisms only if (i) X's ϕ-ing and S's ψ-ing are components in the same mechanism, (ii) X's ϕ-ing is not a component in S's ψ-ing, and (iii) S's ψ-ing is not a component in X's ϕ-ing. Unlike size levels or levels defined in terms of the types of objects found at a given level, mechanistic levels are defined by the componency relationship between things at higher and lower levels. If two things are not related as part to whole, they are not at different levels; if they are in the same mechanism, then they are in this limited sense "at the same level". But one might just as easily say on this basis that sameness of level has no deep conceptual significance for mechanistic levels.

Note, for example, that the levels of nearly decomposable structure in stars do not correspond to the levels in starfish. Each has parts and wholes, but the parts and wholes are of different kinds. It is an empirical question, answered case by case, how many levels there are in a system. It makes no sense to ask whether hippocampi are at a higher or lower mechanistic level than horses. They are not components in the same mechanism; they are not related as parts to wholes. For some, this is tantamount to abandoning the idea of levels (e.g., Eronen 2015). We see our account rather as a distillate of the ordinary scientific concept, an extraction of an explanatorily and metaphysically central idea, leaving behind as residue the problematic commitments inherent in our inchoate, folk talk of "levels."

3. Comparisons and Contrasts

The ordinary concept of "level" is inchoate, in part, because the term is used promiscuously to describe many distinct kinds of relata and relations. Here we distinguish mechanistic levels from aggregates, size levels, causal levels, and Oppenheim and Putnam's levels.

As noted above, mechanisms contrast with aggregates. In aggregates, the property of the whole is literally a sum of the properties of its parts. The concentration of a fluid is an aggregation of particles; allelic frequency is a sum of individual alleles. Aggregate properties change linearly with the addition and removal of parts, they don't change when their parts are rearranged, and they can be taken apart and reassembled without any special difficulty. This is because in true aggregates, spatial, temporal, and causal organization are irrelevant (Wimsatt 1997). Mechanisms, in contrast, are literally more than the sums of their parts: they change non-linearly with the addition and removal of parts, their behavior is disrupted if parts are switched out, and this is because their spatial, temporal, and causal organization make a difference to how the whole behaves.

Though distinct, mechanisms and aggregates are nonetheless species of a genus: each involves a relationship between the properties or activities of wholes and the relevant properties or activities of their parts. Many mechanisms (from steam engines to action potentials) rely on both mechanistic and aggregative relations. As noted above, aggregation is the limit as mechanistic organization goes to zero.

Mechanistic levels also contrast with size levels. The relata in size levels are space-involving entities (like cells), and they are higher- or lower-level than one another because they are bigger or smaller, respectively. They are "at" a level when they have the same (or similar) sizes. Mechanistic levels are also orderable by size, but the size relationship between mechanistic levels follows from the more fundamental, compositional relationship. Like Wimsatt's (1976) classic image of levels as peaks of regularity and predictability, nearly decomposable mechanisms will tend to carve most naturally at interfaces between components, thereby identifying isolable pockets of regularity and predictability. Wimsatt describes but does not explain why these levels of regularity and predictability cluster at different size scales. In mechanistic levels, the clustering is explained by patterns of near decomposability in the mechanism's causal organization.

Sometimes the term "level" is used to describe relations among the stages in a causal pathway, as when one distinguishes "higher-level" and "lower-level" visual areas (e.g., area MT and V1, respectively) or, arguably, when one speaks of "higher level executive functions." These causal relations are clearly distinct from the compositional relation involved in mechanistic levels. MT is not part of V1, but downstream in a causal process from V1; lower-level cognitive functions are not parts of higher-level cognitive functions, but asymmetrically controlled or dominated by higher-level cognitive functions (see Churchland and Sejnowski 1992).

The term "levels" is also associated with Oppenheim and Putnam (1958), who structure their view of the unity of science around a monolithic conception of levels. They carve the world into roughly six ontological strata (societies, organisms, cells, molecules, atoms, and elementary particles). Each stratum corresponds to a distinct field of science, from economics at the top to particle physics at the bottom. Each level has its distinctive theory. The unity of science consists in the explanation of higher-level theories in terms of lower-level theories.

Oppenheim and Putnam could easily have embraced mechanistic levels as the ontic component of their picture; things at higher levels are wholes made up of things at lower levels. As Wimsatt (1986) notes, however, this minor amendment does violence to the simplicity of Oppenheim and Putnam's vision of scientific unity. If levels are compositional, and different mechanisms decompose into different kinds of parts, then we should not expect a uniform decomposition of all things into the same kinds of parts (compare stars to starfish and hippocampi to horses).

Oppenheim and Putnam's vision of scientific unity is descriptively inadequate, as they would likely acknowledge. Both sciences and theories run rough-shod over levels. Models and theories often span levels of organization, linking phenomena studied by different fields of science (Darden and Maull 1977; Craver 2005; 2007). Scientific fields are also increasingly transdisciplinary (e.g. neuroscience, ecology). Things that are the same size can usefully be investigated by altogether distinct fields of science: Cytologists, anatomists, and electrophysiologists all study cells with different tools. The relation between fields of science, levels of theory, and ontological levels is thus many to many to many.

The notion of mechanistic levels provides a compelling alternative to the Oppenheim and Putnam model. Mechanisms span multiple levels of organization. Scientists approach these structures and processes with diverse tools and from diverse theoretical vantage points. The unity that results from this collaboration is more like a mosaic than a layer cake: it is achieved when different scientists with different instruments and modeling tools, use their diverse expertise to understand the same mechanism (Craver 2003; 2007).

4. Levels, Emergence, and Interlevel Causation

Do things at higher levels "emerge" from things at lower levels? And can things at different levels causally interact? If we think of levels as mechanistic levels, these questions have clear content, and it is clear what is at stake in answering them.

Mechanisms are not aggregates. A property or activity at a higher level of mechanistic organization (e.g., S's ψ-ing) is literally more than the sum of the properties of its parts. So, if emergence is defined as the failure of aggregativity (call this *organizational emergence*; Wimsatt (1997)), then things at higher mechanistic levels emerge organizationally from things at lower levels. No ontological extravagance is required: two stacked toothpicks have the organizationally emergent capacity to catapult raisins. Neither toothpick can catapult raisins alone: so the whole has capacities the parts alone do not possess.

Often the term "emergence" is used in an *epistemic* sense to refer to the inability to predict the properties or behaviors of wholes from properties and behaviors of the parts. Epistemic emergence can result from our ignorance, such as failing to recognize a relevant variable, or from failing to know how different variables interact in complex systems. It might also result from limitations in human cognitive abilities or current-generation representational tools (Boogerd et al. 2005; Richardson and Stephan 2007). The practical necessity of studying mechanisms by decomposing them into component parts raises the epistemic challenge of conceptually putting the parts back together so they work (Bechtel 2013). Epistemic emergence results from the limits of our knowledge or of our representational capacities, not from discontinuity in the world's causal structure (see Chapter 28).

Ontic emergence is suspect or promising (depending on one's perspective) precisely because it involves such discontinuity: there are higher-level properties and capacities that have no sufficient (ontic) explanation in terms of the parts, activities, and organizational features of the system in the relevant conditions. Some say life, consciousness, or intentionality are emergent in this sense. Mechanistic levels, in contrast, are levels of mechanistic dependence: they are defined in terms of componency (Craver 2014). If that ontic relationship is severed, then the sense in which emergent properties are at a “higher level” must be different than the sense in mechanistic levels. Indeed, it is unclear why properties that ontically emerge should be thought of as higher-level at all (rather than, e.g., effects of a cause). Ontic emergence, whatever its virtues, is mysterious precisely because it is distinct from, and so gains no plausibility from verbal association with, organizational or epistemic emergence. Organizational and epistemic emergence are unmysterious both in scientific common sense and common sense proper (Van Gulick 1993; Kim 1998). Appeal to ontic emergence, on the other hand, arouses suspicion because it is committed to the existence of phenomena that have no sufficient (ontic) mechanistic explanation.

Can things at different mechanistic levels properly be said to causally interact with one another? Many common assumptions about causation appear to block this thought. For Hume and Lewis, the relata in a causal relationship must be distinct, “and distinct not only in the sense of nonidentity but also in the sense of nonoverlap and nonimplication. It won't do to say that my speaking this sentence causes my speaking this sentence; or that my speaking the whole of it causes my speaking the first half of it; or that my speaking causes my speaking it loudly, or vice versa.” (Lewis 2000, p. 78). Wholes and parts overlap; a token S's ψ-ing as a whole includes every part of S's ψ-ing, and that includes all the X_j, ϕ_i, and organizing relations. On Salmon's process view of causation, the relata in causal relations must intersect in space-time, exchange conserved quantities, and maintain those changes after the intersection ends. But parts and wholes are always everywhere together; the whole has no additional conserved quantity to pass to its parts. The relationship between LTP and the opening of NMDA receptors during LTP induction is directly analogous to the relationship between speaking the whole of a sentence and speaking its first half. The induction of LTP is partly constituted by the opening of the NMDA receptor. The would-be cause in this top-down causal claim already contains the would-be effect; talk of causation across levels of mechanism thus appears inappropriate (Craver and Bechtel 2007).

What about the bottom-up case? Here we must guard against an ambiguity. Sometimes mechanists describe the "phenomenon" as an activity or process that starts with the mechanism's setup conditions and ends with its termination conditions (Machamer et al. 2000). For example, Long-Term Potentiation (LTP) is described as a process beginning with rapid and repeated stimulation of the presynaptic neuron and ending with enhanced synaptic transmission. Other times, mechanists describe the phenomenon as the product of that process (as one of its termination conditions). For example, the mechanism of LTP produces a potentiated synapse. This second phrasing leads us to seek the antecedent causes: the tetanus and the subsequent changes in the NMDA receptor. But if we think about the phenomenon as a process, an input-output relation starting with the tetanus and ending with a potentiated synapse, then we should not say the NMDA receptor is a cause of that, if we wish to speak accurately about the ontic structure of the situation at hand. The NMDA receptor is a part of that causal process; it causes neither itself nor its antecedents in the mechanism.

5. Reduction

Like "level," "reduction" is used many ways. Sometimes, it expresses a thesis about *explanation*, about how one theory or law is explained by other theories or laws. Other times, it is a thesis about *scientific integration* and unity, about the relationships among the diverse branches of science. And still other times, it is a *metaphysical thesis* about the structure of the world (Section 6). The idea of mechanism offers considerable insight into explanation and scientific integration. Although the mechanistic view of levels alone does not resolve the deepest ontological puzzles, it nonetheless offers a descriptively adequate image of the kind of ontological structure by which the manifest phenomena of our world might be (and arguably are) related to its most fundamental parts and activities.

According to the covering law (CL) model, reductive explanations involve deriving higher-level theories from lower-level theories with the aid of bridge laws connecting their distinct theoretical vocabularies (Nagel 1961; Schaffner 1993). This model of explanatory reduction serves as a valuable ideal in the search of mechanisms. A description of the behavior of the mechanism should be shown to *follow from* a description of the parts and their interaction, as one understands them. For example, Hodgkin and Huxley showed that one could recover the shape of the action potential from measured values of ionic concentrations and experimentally determined ion-gating functions coupled with laws governing electrical circuitry. Yet this epistemic test of the sufficiency of an explanation should not be confused with the explanation itself.

Not all predictively adequate models are explanatory. We can, as Hempel (1965) argued, derive Boyle's law from the conjunction of Boyle's law and Kepler's laws, but not in a way that explains it. We can predict the behavior of a part from the description of the behavior of the whole and the behaviors of a few other parts, despite the fact that many expect such micro-reductive explanations to work only from the bottom up. We can predict the behavior of a system by tracking the right correlations in the system irrespective of their causal relations to one another (as in diagnosis and prognosis) and would not thereby explain the system's behavior. Finally, the classical model of reduction is not equipped with an account of explanatory relevance, a way to determine which predictively relevant features are explanatorily relevant and which are not. If one thinks of reduction not as primarily an epistemic, deductive relationship between descriptions of the behavior of a mechanism and of the organized activities of its component parts but as a matter of learning and revealing the causal structure of a mechanism, these problems for the classic theory of micro-reductive explanation are directly addressed (Craver 2007).

Given the multilevel structures found in mechanistic explanations, it's clear that scientific integration will not conform to the layer cake pattern Oppenheim and Putnam describe across science as a whole. Interlevel linkages will be of more local significance, relativized to particular explanandum phenomena. Different fields of science collaborate with one another both to bridge across mechanistic levels and to bring their unique perspectives to bear on one and the same thing. A kind of mosaic unity results from having the same mechanism as the target of one's scientific investigation; integration occurs when the findings from different fields and perspectives mutually constrain our understanding of how the mechanism works.

This mechanistic perspective provides an attractive successor to empiricist views of explanation as a model of interfield integration. First, it is based on a more accurate and informative view of levels. Second, it provides significantly more insight into the diverse evidential constraints by which interlevel bridges are evaluated and into the forces driving co-evolution across different levels (Craver and Darden 2013). Constraints on the parts, their causal interactions, and their spatial, temporal, and hierarchical organization all help to flesh out an interfield theory. Finally, many mechanists argue that classical reduction models are myopically focused on downward-looking explanations at the expense of a fuller account of the details by which interlevel and interfield theories more generally are constructed and evaluated. This is especially true when one must look up to higher levels to see how a part contributes to a mechanism of which it is a component and in which two fields combine findings "at" rather than across mechanistic levels.

6. The Ontological Status of Higher-Level Phenomena

We turn now to the metaphysical questions central to many philosophical discussions of reduction: what is the relationship between higher-level phenomena and things at lower levels? We disambiguated two questions implicit in Figure 14.1. The first is: Is the behavior of the mechanism as a whole greater than the sum of the behaviors of its parts? We defend an affirmative answer above. The second question is: Is the behavior of the mechanism as a whole greater than the organized interaction of its parts in context (commonly called the *realizer* of the mechanism's behavior)?

Simple examples suggest that the relationship between the phenomenon and its realizer is more intimate than correlational, causal, and nomologial relationships are. In a standard mousetrap, if the trigger is depressed, the catch slides, and the impact bar completes its arc (collectively, M), then the trap necessarily fires (F). We are precluded from imagining that all the stages of M have occurred but that the trap has not fired precisely because there is nothing more to be done, no activity or property to add, to transform M in context, into an instance of F. We are similarly precluded from imagining in a population of mousetraps composed of identical parts organized and interacting identically with one another in the same conditions, that only 30% of these mechanisms should fire and 70% should not. Correlational, causal, and nomological relationships are not necessarily like this. The intimacy between a phenomenon and its realizing mechanism thus appears to be stronger, more metaphysically necessary, than a simple matter of regularity, causation, or law. The behavior of the whole contains the behaviors of the parts, and the behaviors of the parts collectively and exhaustively constitute the behavior of the whole. Dualism, parallelism, and emergentism do not share this commitment and allow that there are properties and activities of wholes that are not exhaustively constituted by lower-level mechanisms.

This metaphysical commitment functions as a constitutive ideal in the search for mechanistic explanations (Haugeland 1998). If one knows all of the relevant entities, activities, and organizational features, and knows all the relevant features of the mechanism's context of operation, and can in principle put it all together, then one must know how the mechanism will behave. This was Hempel's important insight about the epistemology of explanation. It is an epistemic warning sign if features of the mechanism's behavior cannot be accounted for in terms of our understanding of its parts, activities, organization and context. Mechanists thus operate with a background assumption that the phenomenon is exhaustively explained (in an ontic sense) by the organized activities of parts in context.

One is tempted to say that higher-level phenomena (in a given context) *are identical to* the organized activities of the mechanism's parts (in that context). By this one might assert a type identity, according to which ψ-ing in general is identical to the organized ϕ-ing of X's, or a token identity, according to which each instance of ψ-ing is identical to some organized ϕ-ing of X's. Classical reductionists favor type identity, though this thesis is largely out of fashion (though see Polger 2004). As Putnam (1967) and Fodor (1974) argued, such type-identity statements are false when the phenomenon is or can be multiply realized in distinct mechanisms. There are many ways to trap a mouse and many ways to produce an action potential. If so, the same type of phenomenon might be instantiated on different occasions in different mechanisms; the realized type is not the same as the realizer type.

Bechtel and Mundale (1999) counter that this argument rests on an illusory matching of grains between the characterization of the phenomenon and the characterization of the mechanism. If one characterizes the phenomenon abstractly, then many mechanisms can or could give rise to it. If one characterizes the phenomenon in fine detail (including for example, all its excitatory, inhibitory and modulatory conditions), then the class of actual and potential realizers might shrink to one. If the phenomenon and the mechanism are characterized at comparable degrees of abstraction, then the thesis of type-identity begins to look more promising.

But perhaps a deeper problem for type-identity is lurking in the background and not so easily addressed: the question of whether there is, in general, an "appropriate" grain of analysis that adequately captures the phenomenon and the mechanism at once. Even some token things are multiply realizable: One and the same rook in a chess game might be instantiated at different times by different figurines (imagine replacing a lost plastic castle midgame with a Fig Newton; Haugeland 1998). One and the same kidney is made up of different parts over a person's lifetime. To answer whether ψ and the organized collection of ϕ-ers are the same thing, we need to be able to answer the question "The same what?" with a non-dummy sortal (Marcus 2006). The conditions under which something is the same ψ-er might be different from the conditions under which it is the same organized collection of ϕ-ers, in which case there are no individuation conditions that cover the thing both as a ψ-er and as an organized collection of ϕ-ers, and there is no sense to the question of whether they are identical. Items at different mechanistic levels are no-doubt intimately related, but our ordinary notions of token and type identity appear poorly equipped to describe that relation. For these reasons, we prefer to speak of a particular organized ϕ-ing of X's as exhaustively constituting a particular ψ-ing at or over a time.

Instead of asking about identity, one might ask rather about whether phenomena have causal powers in a given context that the organized collection of interacting parts do not have. As noted above, it is relatively uncontroversial that wholes can do things their parts cannot do (see Kim 1998, 85). Do realized phenomena have causal powers their realizing mechanisms do not have in the same contexts? If one thinks of the phenomenon as identical to the behaving mechanism, then the answer is no: the two have all and only the same powers. Non-reductive physicalists have traditionally attempted to argue that, though higher-level phenomena are all ontically explainable in terms of more fundamental mechanisms, there remains a sense in which the higher-level phenomena have causal powers of their own.

Again, issues of multiple realization arise. Suppose a neuron generates an action potential, causing neurotransmitters to be released from the cell. Whenever it does, we suppose there is a determinate set of entities (ions, channels, etc.) organized and interacting in determinate ways that fully constitutes it at that time. Now suppose we ask: which feature of the world is relevant to the release of neurotransmitters: the action potential or the determinate organization and interaction of ions and channels? One standard test of causal relevance is to ask what would happen if the cause had been different. Experimental investigation will reveal that the action potential makes a substantial difference to the probability of neurotransmitter release by raising membrane voltage. The precise arrangement of ions and ion channels makes no difference except insofar as it affects membrane voltage. Many differences among the realizers make no difference to the effect in question. If causal relevance is a matter of making a difference (Woodward 2003), the fine details about the realizing mechanism are not causally relevant. Each action potential is exhaustively constituted by a mechanism, but the causally relevant difference is the occurrence of the action potential, not the precise mechanism that produced it (see Craver 2007, Chapter 6).

Just how this mechanistic perspective best fits with metaphysical views about, e.g., identity and causal powers, remains to be decided. A determinate view about multiple realization and its implications requires metaphysical commitments that go beyond the skeletal framework of levels presented here. Views on multiple realization differ depending, *inter alia*, on how one thinks about properties (Heil 1999), the realization relation (Shapiro 2004), and theories (Klein 2013). Similarly, the prospects for token physicalism depend, *inter alia*, on one's account of events (Kim 2012) and identity (Marcus 2006). One goal of future work is to integrate the mechanistic perspective with a metaphysics that produces a coherent and plausible view of the structure of the mechanistic world.

7. The Lowest Level

Questions also arise concerning the bottom of a hierarchy of mechanisms. Two options seem to be available: either 1) the hierarchy has a bottom or 2) it does not, i.e., it continues ad infinitum. Though both options are metaphysically live, some facts about quantum mechanics seem to favor the first.

It has been argued, for example, that the non-separability of ontological units in the theories of quantum mechanics entails that there are no distinct objects to be identified as components beneath a certain level of organization. Kuhlmann and Glennan (2014) argue that for any local, classical phenomenon, there is a lowest classical level at which mechanistic explanation bottoms out. Precisely where it bottoms out varies across phenomena because decoherence – the physical process by which the quantum realm becomes classical – is itself local. Below this level, where properties like spatial location are indeterminate, prospects for mechanistic explanation are murky. Mechanistic explanation may still be possible if quantum phenomena can be explained by causal organization among non-local "objects" (for example, see the discussion of laser light [Kuhlmann and Glennan 354]), but in quantum holistic cases where the properties of subsystems don't appear to determine the properties of the whole system (for example, EPR experiments), the constitutive ideal may be unreachable.

8. Conclusion

Our goal has been to sketch the mechanistic approach to levels, to contrast it with near neighbors, and to explore some of its metaphysical implications. This perspective provides some important content missing from standard metaphysical discussions of levels. Most notably, it offers a clear sense of what it means for things to be at different levels. It also allows us to distinguish mechanistic levels from realization relations. And finally, it provides a clearer vision of the structure of multilevel explanations, the evidence by which such explanations are evaluated, and the form of scientific unity that results from building them. The mechanistic approach does not, by itself, solve all the metaphysical questions one might have about physicalism. Yet it does provide a skeletal framework for thinking about how the macroscopic phenomena of our world are or might be related to its most fundamental parts and activities. Simon's parable and its many contemporary analogues give us reason to think that this skeletal picture captures a central and important pattern in the causal structure of our world.

1. Thanks to Mike Dacey, James Gulledge, Ben Henke, Eric Hochstein, Emily Prychitko, and Felipe Romero for invaluable comments and to Pamela Speh for design of Figure 14.1.

References

Albert, R., H. Jeong, and A.-L. Barabási. (2000) "Error and Attack Tolerance of Complex Networks" *Nature*, 406, 378–382.

Bechtel, W. (2013) "Addressing the Vitalist's Challenge to Mechanistic Science: Dynamic Mechanistic Explanation," in S. Normandin and C. T. Wolfe (eds), *Vitalism and the Scientific Image in Post-Enlightenment Life Science 1800-2010*, Dordrecht, NL: Springer.

Bechtel, W. and J. Mundale. (1999) "Multiple Realizability Revisited: Linking Cognitive and Neural States," *Philosophy of Science*, 66, 175–207.

Bechtel, W. and R.S. Stufflebeam. (2001) "Epistemic Issues in Procuring Evidence About the Brain: The Importance of Research Instruments and Techniques," in W. Bechtel, P. Mandik, J. Mundale and R.S. Stufflebeam (eds), *Philosophy and the Neurosciences: A Reader*, Blackwell: Malden, MA, 55–81.

Boogerd, F.C., F.J. Bruggeman, R.C. Richardson, A. Stephan and H.V. Westerhoff. (2005) "Emergence and Its Place in Nature: A Case Study of Biochemical Networks," *Synthese*, 145, 131–164.

Churchland, P.S. and T. Sejnowski. (1992) *The Computational Brain*, MIT Press: Cambridge, MA.

Couch, M. (2011) "Mechanisms and Constitutive Relevance," *Synthese*, 183, 375–388.

Craver, C.F. (2001) "Role Functions, Mechanisms, and Hierarchy," *Philosophy of Science*, 68, 53–74.

Craver, C.F. (2007) *Explaining the Brain*, Oxford University Press: Oxford.

Craver, C.F. (2014) "The Ontic Account of Scientific Explanation," in M.I. Kaiser, O.R. Scholz, D. Plenge, and A. Hüttemann (eds), *Explanation in the Special Sciences: The Case of Biology and History*, Springer: Netherlands, 27–52.

Craver, C. F. and W. Bechtel. (2007) "Top-down Causation without Top-down Causes," *Biology and Philosophy*, 22, 547–563.

Craver, C.F. and L. Darden (2013) *In Search of Mechanisms: Discoveries Across the Life Sciences*, University of Chicago Press: Chicago.

Cummins, R. (1975) "Functional analysis," *Journal of Philosophy,* 72, 741–765.

Darden, L. and N. Maull. (1977) "Interfield Theories," *Philosophy of Science*, 44, 43–64.

Eronen, M.I. (2015) "Levels of Organization: A Deflationary Account," *Biology and Philosophy*, 30, 39–58.

Fodor, J. (1974) "Special Sciences: Or the Disunity of Science as a Working Hypothesis," *Synthese*, 28, 97–115.

Gillett, C. (2002) "The Dimensions of Realization: A Critique of the Standard View," *Analysis*, 62, 316–323.

Glennan, S. (2002) "Rethinking Mechanistic Explanation," *Proceedings of the Philosophy of Science Association*, 3, S342–S353.

Harinen, T. (2014) "Mutual Manipulability and Causal Inbetweenness," *Synthese*, 1–20.

Haugeland, J. (1998) *Having Thought: Essays in the Metaphysics of Mind*, Harvard University Press: Cambridge, MA.

Heil, J. (1999) "Multiple Realizability," *American Philosophical Quarterly,* 36, 189–208.

Hempel, C. (1965) *Aspects of Scientific Explanation*, Free Press: New York, NY.

Kauffman, S.A. (1970) "Articulation of Parts Explanation in Biology and the Rational Search for Them," *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association,* 1970, 257–272.

Kim, J. (1998) *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*, MIT Press: Cambridge, MA.

Kim, J. (2012) "The Very Idea of Token Physicalism," in C. Hill and S. Gozzano (eds), *New Perspectives on Type Identity: The Mental and the Physical*, Cambridge University Press: Cambridge, 167–185.

Klein, C. (2013) "Multiple Realizability and the Semantic View of Theories," *Philosophical Studies*, 163, 683–695.

Kuhlmann, M. and S. Glennan. (2014) "On the Relation Between Quantum Mechanical and Neo-Mechanistic Ontologies and Explanatory Strategies," *European Journal for Philosophy of Science*, 4, 337–359.

Levy, A. and W. Bechtel. (2013) "Abstraction and the Organization of Mechanisms," *Philosophy of Science*, 80, 241–261.

Lewis, D. (2004) "Causation as Influence," reprinted in J. Collins, N. Hall and L.A. Paul (eds), *Causation and Counterfactuals*, MIT Press: Bradford, 75–106.

Machamer, P. (2004) "Activities and Causation: The Metaphysics and Epistemology of Mechanisms," *International Studies in the Philosophy of Science*, 18, 27–39.

Machamer, P., L. Darden and C.F. Craver. (2000) "Thinking about Mechanisms," *Philosophy of Science*, 67, 1–25.

Marcus, E. (2006) "Events, Sortals, and the Mind-Body Problem," *Synthese*, 150, 99–129.

Marr, D. (1982) *Vision*, Freeman Press: San Francisco.

Martin, C.B. and M. Deutscher. (1966) "Remembering," *Philosophical Review*, 75, 161–96.

Nagel, E. (1961) *The Structure of Science: Problems in the Logic of Scientific Explanation*, Harcourt, Brace and World, Inc.: New York.

Oppenheim, P. and H. Putnam. (1958) “Unity of Science as a Working Hypothesis,” in H. Feigl, M. Scriven, and G. Maxwell (eds), *Concepts, Theories, and the Mind-Body Problem, Minnesota Studies in the Philosophy of Science II*, University of Minnesota Press: Minneapolis, 3–36.

Polger, T. (2004) *Natural Minds*, MIT Press: Cambridge, MA.

Putnam, H. (1967) “Psychological Predicates,” in W.H. Capitan and D.D. Merrill (eds), *Art, Mind, and Religion*, University of Pittsburgh Press: Pittsburgh, 37–48.

Richardson, R.C. and A. Stephan. (2007) “Emergence,” *Biological Theory*, 2, 91–96.

Romero, F. (2015) “Why There isn't Inter-level Causation in Mechanisms,” *Synthese*, 192, 3731–3755.

Salmon, W. (1984) *Scientific Explanation and the Causal Structure of the World*, Princeton University Press: Princeton.

Schaffner, K. (1993), *Discovery and Explanation in Biology and Medicine*, University of Chicago Press: Chicago.

Schlosser, G. and G. Wagner (eds), (2004) *Modularity in Development and Evolution*, University of Chicago Press: Chicago.

Shapiro, L. (2004) *The Mind Incarnate*, Oxford University Press: Oxford.

Simon, H. (1969) *The Sciences of the Artificial*, MIT Press: Cambridge, MA.

Van Gulick, R. (1993) “Who's in Charge Here? And Who's Doing All the Work?” in J. Heil and A.R. Mele (eds), *Mental Causation*, Oxford University Press: Oxford, 233–56.

Wimsatt, W.C. (1976) “Reductionism, Levels of Organization, and the Mind-Body Problem,” in G. Globus, G. Maxwell, and I. Savodnik (eds), *Consciousness and the Brain: A Scientific and Philosophical Inquiry*, Plenum Press: New York, 202–267.

Wimsatt, W.C. (1997) “Aggregativity: Reductive Heuristics for Finding Emergence,” *Philosophy of Science*, 64, 372–384.

Woodward, J. (2003) *Making Things Happen*, Oxford University Press: Oxford.