
Mechanisms and Emergence

A Reply to Denis C. Martin

[Carl F. Craver](#)

I respond to some of the major issues in Martin's commentary, particularly (i) his insistence on a robust notion of being "at" a level, and (ii) his desire for mechanistic emergence to explain the genuine ontological novelty of higher level phenomena.

Keywords

Emergence | Levels | Mechanisms

Author

[Carl F. Craver](#)

ccraver@artsci.wustl.edu

Washington University
St. Louis, MO, USA

Commentator

[Denis C. Martin](#)

denis.martin@hu-berlin.de

Humboldt Universität zu Berlin
Berlin, Germany

Editors

[Thomas Metzinger](#)

metzinger@uni-mainz.de

Johannes Gutenberg-Universität
Mainz, Germany

[Jennifer M. Windt](#)

jennifer.windt@monash.edu

Monash University
Melbourne, Australia

1 Introduction

The goal of my target essay is to articulate and recommend a view about what one is and is not committed to when one asserts that the domain of neuroscience or some other special science spans multiple levels of organization. How useful one finds my specific recommendation—that they be understood as levels of mechanisms—will depend on the uses for which one deploys the level metaphor. Martin raises a number of questions about my unpacking of the metaphor that help to clarify what is at stake in this discussion.

2 The levels metaphor

First, a word about the idea that the term level is used as a metaphor. According to the Oxford English dictionary, the term "level" derives from terms describing the idea of being parallel to the ground. From there, the term could easily extend to the idea that a building can have multiple levels, or landings, from the ground up. Something like this spatial arrangement of stacked landings seems to be the basis for many applications of the level metaphor (in other words, landings are the "secondary subject" of the metaphor). Usage of the term "level" increased dramatically in the 20th Century, and surely some of this increase is explained by the

fecundity of the metaphor for ordering items in a set from the highest to the lowest (under some understandings of “high” and “low”). These are the specific topics (the primary subjects) to which the secondary subject (landings) is compared. In asserting that the term “level” has this metaphoric aspect, I do not intend to thereby disqualify the metaphor from also expressing a set of true or false commitments about the structure of the world. Indeed, I think that the mechanistic application of the metaphor more or less accurately describes the structure of at least many systems in the domains of the special sciences.

My point in being a pluralist about the levels metaphor is simply to emphasize that the metaphor of a horizontal landing is used to describe a wide variety of altogether distinct primary subjects. As a result, any understanding of how this level metaphor works, what it is committed to, and whether it works for a particular purpose, must begin with a clear understanding of how the metaphor is unpacked in the given application. (One cannot, to pick a perfectly clear example, equate the great chain of being and levels of mechanisms.) A second point was to emphasize that there are many legitimate applications of the metaphor, and that the utility of the metaphor might vary from one application to the next. I am not arguing for a single, correct way of understanding the level metaphor. However, I am arguing that one prominent use of the level metaphor in sciences such as neuroscience can be usefully unpacked as describing mechanistic levels.

3 Being “at” a level

Martin is particularly concerned with a) the idea that there should be a clear sense of what it is to be “at” a level, and b) the idea there is some significant sense in which being at a level represents a genuine ontological novelty.

The first of these concerns arises from the fact that the placement question cannot be (or has not been) given a satisfying answer within the mechanistic application of the metaphor. The mechanistic account of levels focuses, as Martin correctly notes, on what it means to be

at *different* levels of organization within a mechanism. Concerning what it means to be “at a level,” one can say only that two things are “at” the same level just in case they are not at different levels; if neither thing is a component entity/activity in the behavior of the other, then the two things are not at different mechanistic levels and, in this weak sense, they are at the same level. This provides a sort of answer to the placement question, but not one that will satisfy those, such as Martin, who hanker after some additional factor (such as size or similarity) that unites the items at a level and that explains why all such items are at that level. I take the failure to answer the placement question as one of the key revisionary consequences of thinking clearly about levels of mechanisms. It is a crucial guide to understanding what’s misleading about the monolithic conception of levels.

One apparent problem for this consequence (the absence of a satisfying answer to the placement question) was first raised by Lindley Darden (personal communication): X’s -ing might be a component in (and so at a lower level than) S’s ψ -ing, *and* yet both X’s -ing and S’s ψ -ing might be at the same level as (i.e., not at a different level than) some altogether distinct P’s β -ing. There is a failure of transitivity. This, it seems to me, is simply a consequence of the idea that the application of the level metaphor to levels of mechanisms breaks down when one is not talking about relations between parts and wholes. This is unproblematic; it is simply an alternative way of expressing the idea that being “at the same level” is of no additional metaphysical significance within the mechanistic application of the metaphor than simply not being at different levels. This does not prevent one, of course, from using some other ordering criterion for this expressive purpose; one might lump things together on the basis of their sizes or perhaps the instruments used to detect them. But one should be aware that at this point one has left the mechanistic application of the metaphor.

Another consequence of the mechanistic view is that one might equally correctly carve the boundaries of mechanisms in any number of

ways (see e.g., Craver 2004, 2009, 2012). If so, there will not be a uniquely correct answer to the question of how many levels a given mechanism has. Our decisions to privilege some grains in the decomposition of a mechanism as an appropriate place to locate a level depends on our techniques, our theoretical background assumptions, our characterization of the phenomenon, our representational tools, and perhaps certain features of human psychology. As Simon argued, these systems are only *nearly* decomposable; how a system is decomposed depends, for example, on the relative strength of intra- versus extra-system causal interactions that one takes to be appropriate for carving the system at its near-joints. This is another reason I am hesitant to pin too much on the idea of being “at” a level.

4 Ontological novelty and emergence

I am not hostile to the second idea (b) that levels of mechanisms exhibit a kind of ontological novelty, depending on how this novelty is understood and how one thinks that it is achieved. I suggest in the target article that a mechanism as a whole can do things that its parts (taken individually) cannot do. Lawnmowers mow grass; spark plugs do not. But I also claim that a mechanism as a whole cannot do things that its organized and interacting parts cannot do. This is because the behavior of a mechanism as a whole just is (or at least is ontologically intimate with) the organized interaction of its component parts. I don't know how to make the notion of “ontological intimacy” precise here (though the behavior of a mechanism in context will surely supervene on the organized collection of interacting components in that context). In my hands, “ontological intimacy,” is meant to denote an exhaustive ontological grounding of the behavior of a mechanism as a whole in the organized interactions of its components in a given causal context, however that is properly to be unpacked. Everything has an ontic explanation in terms of the organized activities of the mechanism's parts.

The term emergence strikes me as suspicious, and I hesitate to use it even in the con-

text of levels of mechanisms, precisely because it suggests a severing of this ontological intimacy, a slide from the banal fact that mechanisms behave as they do because they are organized arrangements of interacting parts to the ontologically esoteric thought that something comes into being with causal powers ontologically inexplicable in terms of the organized interactions of the parts. Such claims about ontological novelty are almost always accompanied by claims that emergent things have a “downward” causal influence. I have labored to drive a wedge between levels of mechanisms and such ideas.

This connects with Martin's description of Thompson's view. According to Martin, Thompson requires that the parts in a system must be coupled and dynamically interacting to produce emergent properties and, further, that the properties of the whole should act “downward” on the dynamics of the components (2007). Take these in turn.

The idea of levels of mechanisms, by itself, requires only that the parts be organized and interact with one another; it does not require that the components interact non-linearly. Mechanisms with components that interact non-linearly are a subset of mechanisms more generally. It should not make any difference with respect to the novelty of higher-level properties that the parts interact non-linearly. Such dynamical interactions might make the behavior of the whole harder to predict on the basis of our understanding of the parts; but this is an epistemic, not an ontological, observation irrelevant to the ontological question to which I am responding. To be ontologically novel is not merely to be surprising.

Perhaps systems with dynamically interacting components are harder to idealize into separable interacting components; again, this would appear to be an epistemic issue. Thompson, at least in Martin's description, is drawing a more fine-grained distinction than I draw. This is perhaps a terminological matter. The pressing question between us is whether non-linearity (or perhaps some other form of complexity) makes any ontological difference to the kind of thing that “emerges” from the “organized interactions” of the parts. I'm inclined

to say no, but a full consideration of the issue would require a more detailed treatment than I can give it here.

Thompson also requires that the higher-level phenomenon should have “determinative” influence on things at lower-levels. “Determinative” surely cannot mean the existence of a universally quantified material conditional with the behavior of the mechanism as a whole as the antecedent and the organized interactions among the components as the consequent. Multiple realization precludes such an analysis. But “determinative” also cannot be understood in the causally productive sense in which something (such as charge or momentum) is passed from whole to part. And for reasons expressed in the target article, there are many widely shared assumptions about the independence of causes and effects that stand in the way of understanding such top-down relations in causal terms. This is why I suggest (as in Craver & Bechtel 2007) that the language of downward causation is misleading in the context of token mechanisms. We recommend that claims about downward causation are really shorthand ways of expressing an often complex web of constitutive (about relations between things at different levels of mechanisms and levels of realization) and causal claims (about things that are not at different levels from one another) about the system. However this is to be cashed out, the term “determinative” seems misleading.

5 Levels and techniques

Finally, I’ll offer a brief remark concerning Martin’s positive suggestion for reifying levels. Martin’s central idea—that we often pick out things as being on a level because they are detectable with the same or a similar kind of apparatus or with the same or a similar set of procedures—seems to me correct and important. Among the many possible ways of decomposing a spatio-temporal whole into component parts, some of these correspond to items that for one reason or another are readily detectable as such by one or more experimental apparatus. This way of putting things highlights how pragmatic factors, such as available apparatus, determine what we

will take to be an appropriate “landing” in a hierarchy of levels. And it raises useful questions about, for example, when two different apparatus, or two different tasks, or two different procedures in fact target the same items, or items at the same level. On a fine-grained characterization of our experimental instruments, no two experiments are the same. On a coarse grain, even superficially quite distinct experiments can be targeting the same phenomenon (consider, e.g., implicit bias or spatial memory). As noted above, it seems to me that many other epistemic, theoretical, and psychological factors enter into these decisions as well.

This good point, however, is obscured by an ontology in which properties apparently come into existence during acts of detection. Martin’s view is that objects have disposition profiles, and these profiles are turned into properties when they are measured. It is unclear to me, however, why one would not want to say that properties are there to be detected all along, or perhaps that properties just are the dispositional profiles of things. Martin doesn’t do much to motivate this experimental idealism about properties, but the thought seems hard to motivate, at least for macroscopic phenomena. It is an apparent consequence of Martin’s account that levels don’t exist until they are detected. And it is an apparent consequence of his view that new levels come into existence when we develop new instruments to detect them. And, to reiterate the thought in the last paragraph, we will have to wrestle with the question of when two techniques detect the same thing or different things. Martin’s positive proposal makes this question much more pressing, given that, for Martin’s view, the structure of the world—specifically, the distribution of properties in space and time—apparently hangs in the balance. But if we abandon the idea that the placement question must have a uniquely correct answer, these questions are less pressing for thinking about ontology; nonetheless, questions about how we coordinate different experimental tasks, protocols, and procedures are at the heart of the epistemological challenge faced by any experimentalist (see Sullivan 2009, 2010).

References

- Craver, C. F. (2004). Dissociable realization and kind splitting. *Philosophy of Science*, 71 (5), 960-971.
- (2009). Mechanisms and natural kinds. *Philosophical Psychology*, 22 (5), 575-594. [10.1080/09515080903238930](https://doi.org/10.1080/09515080903238930)
- (2012). Functions and mechanisms: A perspectivalist account. In P. Huneman (Ed.) *Functions* (pp. 133-158). Berlin, GER: Springer.
- Craver, C. F. & Bechtel, W. M. (2007). Top-down causation without top-down causes. *Biology and Philosophy*, 22 (4), 547-563. [10.1007/s10539-006-9028-8](https://doi.org/10.1007/s10539-006-9028-8)
- Sullivan, J. (2009). The multiplicity of experimental protocols: A challenge to reductionist and non-reductionist models of the unity of neuroscience. *Synthese*, 167 (3), 511-539. [10.1007/s11229-008-9389-4](https://doi.org/10.1007/s11229-008-9389-4)
- (2010). Reconsidering spatial memory and the Morris water maze. *Synthese*, 177 (2), 261-283. [10.1007/s11229-010-9849-5](https://doi.org/10.1007/s11229-010-9849-5)
- Thompson, E. (2007). *Mind and life: Biology, phenomenology, and the sciences of the mind*. Cambridge, MA: Belknap Press of Harvard University Press.