

# Mechanistic Hierarchy Realism and Function Perspectivalism

## Abstract

Mechanistic explanation involves the attribution of functions to both mechanisms and their component parts, and function attribution plays a central role in the individuation of mechanisms. Our aim in this paper is to investigate the impact of a perspectival view of function attribution for the broader mechanist project, and specifically for realism about mechanistic hierarchies. We argue that, contrary to the claims of function perspectivalists such as Craver, one cannot endorse both function perspectivalism and mechanistic hierarchy realism: if functions are perspectival, then so are the levels of a mechanistic hierarchy. We illustrate this argument with an example from recent neuroscience, where the mechanism responsible for the phenomenon of ephaptic coupling cross-cuts (in a hierarchical sense) the more familiar mechanism for synaptic firing. Finally, we consider what kind of structure there is left to be realist about for the function perspectivalist.

## 1. Introduction

Mechanistic explanation has emerged as the dominant account of explanation across much of philosophy of science, especially for cognitive neuroscience, biology, and chemistry. This account suggests that to explain a phenomenon is to offer a mechanism that produces it. Crucially, for many mechanists, mechanisms differ from models and other idealized, counterfactual, or instrumental constructs of science, in that mechanisms are actual parts of the world. Models, diagrams, simulations, or other descriptions may represent a mechanism, but the mechanism itself comprises a real structure, independent of our aims and interests, and this real structure plays an explanatory role (either constitutively, for ontic mechanists, or

by reference, for epistemic mechanists – we will return to this point later). Call this position *mechanism realism*.

The exact metaphysical commitments of mechanism realism are nevertheless unclear. The aim of this paper is to clarify these commitments by looking at the tension between two prominent positions within the mechanism literature, realism about the hierarchical arrangement of mechanisms into levels, and perspectivalism about the attribution of functions to mechanisms. We first demonstrate a negative result, that realism about mechanistic hierarchies and perspectivalism about mechanistic functions are inconsistent positions. This is a significant result, as many mechanists endorse at least one of these positions, and at least one prominent mechanist, Carl Craver, explicitly endorses both. We add empirical weight to this *a priori* argument through the example of ephaptic coupling, a functional mechanism in neuroscience that appears to undermine realism about the hierarchical structure of another popular exemplar mechanism, synaptic firing.

We then turn from this negative result to a positive framework for thinking about the exact metaphysical commitments of mechanism realism. A popular view, held by critics and adherents alike, holds that mechanism realism is relatively innocuous, placing minimal and intuitive constraints on the causal and mereological structure of the world. However, once we articulate the precise condition on the interaction between causality and mereology endorsed by mechanistic hierarchy realism, we can see that it constitutes a severe constraint on ontological structure. We argue that there is a large and unexplored space of substantive metaphysical possibilities between a completely deflated view of mechanisms, which places no substantive constraints on ontology, and a full-blown realism about mechanistic hierarchies, which places severe and unrealistic constraints. Our conclusion is disjunctive: one may either (i) maintain realism about the locally hierarchical arrangement of mechanisms, but then one must endorse a more substantive account of mechanistic functions

than perspectivalism can offer, or (ii) maintain function perspectivalism, but then, in order to preserve a non-trivial mechanism realism, one must articulate a substantive constraint on the relation between mereological and causal structure that is weaker than hierarchy realism.

## 2. Function Perspectivalism for Mechanism Realists

Most generic definitions of mechanism realism commit the realist to (i) entities, performing (ii) activities, in a way that is (iii) organized (see e.g. Machamer, Darden, & Craver 2000; Glennan 2002; Bechtel & Abrahamsen 2005). But much turns on how exactly one understands the stipulation that the mechanism be ‘organized’: if organization is meant in a deflated sense of there being some sort of structure, then it seems any arbitrary set of interacting entities would count as a mechanism. Such a deflated view would render mechanisms metaphysically innocuous, but also explanatorily impotent, for attributing a mechanistic structure to a system would amount to no more (or at least not much more) than the attribution of causal structure (see Wilhelm 2019 for a description of a view along these lines). In fact, a central tenet of the mechanist movement is that mechanistic explanation is distinct from, and more metaphysically ambitious than, explanations that merely aggregate causes (see e.g. Craver and Bechtel 2007; Craver 2015; Darden 2002; Raerinne 2011).

Consequently, mechanism realists do not tend to endorse this deflated view. Rather, they stress that there are no mechanisms *simpliciter*; mechanisms can only be understood as mechanisms *for* the production of some phenomenon (Machamer, Darden, and Craver 2000, 3; Glennan 2002, S344; Bechtel and Abrahamsen 2005, 423; Illari and Williamson 2012, 120). It is only once a target phenomenon has been identified that the investigation of a mechanism can proceed, and correctly characterizing the phenomenon is an important part of mechanistic explanation (Shagrir and Bechtel 2017). This is commonly understood in terms of a mechanism having the *function* of producing the phenomenon of interest (Garson 2013),

and in general, attributing functions to the mechanism as a whole, and/or its parts, seems to be necessary for recovering a non-trivial notion of mechanistic organization.

To see how this works, consider the oft-discussed example of the heart. A mechanistic explanation of the heart decomposes it into component entities and their organized activities. To do so, however, we must start from a (perhaps, tentative) attribution of function. For instance, Galen's theory that the heart served the function of shunting clean blood past waste-carrying blood for expulsion at the lungs drove him to interpret its central wall as a semipermeable filter, i.e. to identify it as a relevant entity organized to engage in a filtering activity. It took William Harvey's insight that the heart could be understood as analogous to a pump to correctly identify the roles played by its component muscles. Different function attributions result in different mechanistic explanations of the same object, and they do so by picking out different features as relevant. Without the guidance of *some* initial function attribution, we would be unable to distinguish those features of the heart that contribute to it as a mechanism within the circulatory system from those that do not (its color, the sound it makes, the heat it radiates, etc.). Typically, function analysis is iterative: a coarse grained attribution of function serves to construct an initial mechanism 'sketch', which in turn motivates a refinement of the phenomenon in an iterative process that increasingly specifies the putative mechanism's parts and activities (Piccinini and Craver 2011; Shagrir and Bechtel 2017).

The notion of function at issue here is *prima facie* normative — if the function of  $X$  is to  $p$ , then  $p$  may not simply be equated with the behavior of  $X$ , as  $X$  may fail to  $p$ , i.e. it may *malfunction*. The possibility of malfunctioning goes hand in hand with the notion of a mechanism *for* some phenomenon — if  $M$  is a mechanism for  $p$ , and  $p$  does not occur, then  $M$  fails to satisfy the relevant success conditions, i.e. it 'malfunctions'. A myocardial infarction affects the ability of the heart to pump effectively, and an adequate mechanistic model must

be able to identify this as a malfunction, on pain of failure to explain healthy heart behavior. The challenge for would-be mechanism realists, then, is to find some way to make this *prima facie* normativity naturalistically acceptable. There are two broad categories of response to this problem: one seeks some sophisticated descriptive feature of the world by which the normative character of functions may be *naturalized*; the second gives up on naturalization and opts for *perspectivalism*, taking the function attributed during mechanistic decomposition as a feature of our explanatory perspective rather than something intrinsic to the mechanism itself.

The most popular naturalization strategies ground functions in the causal history of a mechanism. For instance, etiological accounts (Millikan 1989; Neander 1991) identify function with the role a target object or organ played over evolutionary history in contributing to its ancestors' survival and reproduction. Selectionist accounts (Garson 2017) generalize this strategy, appealing to a wide range of selective processes, including not only evolution, but also learning, synaptic pruning, and technological design, to fix the function of an organism or complex system. These etiological accounts pose a problem for mechanists, as they imply that structurally identical systems may perform different functions if they have different causal histories, which is in tension with the mechanists' commitment to the explanatory primacy of the local organization of entities and activities. A naturalistic account that avoids this worry is Maley and Piccinini's (2017) objective goal account (see also Mossio *et al* 2009, Boorse 1976, and Rosenbleuth *et al* 1943). The basic idea is to identify some states or behaviors of the organism (such as survival or inclusive fitness) that are objective, yet goal-like, and can serve as success conditions for a functional decomposition. Arguably, however, this strategy does little to naturalistically dissolve the seeming mystery of function normativity, but simply kicks the can down the road by *de facto* attributing normativity to objective goals (cf. Dewhurst 2016).

The perspectivalist worries that this general line of approach, naturalizing function by reducing it to some supposedly less ontologically suspect part of nature, is doomed to failure. So long as one's account remains teleological, in the sense that it is committed to there being a genuine fact of the matter about what the function of some system is, one has to say that there is some normativity inherent to the mechanism itself. In contrast, the perspectivalist takes function attribution to be a feature of our explanatory practice, playing a heuristic role in identifying and individuating mechanistic structures, but not constituting an objective feature of the mechanism. This is a strategy that has been endorsed by some mechanism realists, and its compatibility with realism will be the focus of our analysis.

Function perspectivalism essentially develops the strategy of Cummins (1975). Cummins argues that the notion of 'function' appropriate for functional analysis is that of disposition, or causal role. On this deflated account, there are no proper or intrinsic functions, strictly speaking; any causal activity of a complex system qualifies as a 'function' in this sense. As Cummins himself realized, for the very reasons articulated above in the heart example, one needs some criteria to distinguish some causal roles from others when constructing a mechanistic explanation — for Cummins, this role is performed by the context of explanation, which typically starts with the goal of explaining a capacity of a larger system, for instance the pump-like entities and activities of the heart may be singled out “against the background of an analysis of the circulatory system's capacity to transport food, oxygen, wastes, and so on, which appeals to the fact that the heart is capable of pumping” (762).

In recent literature, the notion of context has been supplanted by that of perspective, and a recognition that explanation in general, and function attribution in particular, are inherently perspectival. Hardcastle (1999), for instance, presents an early (pre-mechanistic) account of perspectival functions that has been quite influential, and Kästner (2018) has

recently proposed a fully perspectival analysis of mechanistic explanation, which nevertheless claims compatibility with mechanism realism. Glennan (2017) also endorses a perspectival view of mechanistic explanation, but it is less clear whether he is committed to mechanism realism. More generally, many mechanists tend to agree that there is a perspectival component to how the target phenomenon (and thus, mechanistic function) is fixed (see e.g. Darden 2008: 960), but nonetheless want to retain something like a realist attitude towards the mechanism (and mechanistic explanation) itself. For the sake of specificity, we'll focus in the following sections on the detailed and representative account of Carl Craver, who we take to be committed both to function perspectivalism (in his 2013) and to something like realism about the mechanistic hierarchy (see e.g. his 2001, and Craver & Bechtel 2007).

What are the metaphysical implications of function perspectivalism? Kästner claims there are none, that it is “metaphysically neutral” (2018: 76), and thus compatible with full-blown mechanism realism. Craver is explicit that his perspectivalism amounts to anti-realism about proper functions: at the end of the day, there are only causal relationships in the world, and no true functions (or the norms they entail). Nevertheless, he is unclear about the knock-on implications of function perspectivalism for the metaphysics of mechanisms. At times he states that perspectivalism will extend to mechanisms as well (2013: 134), yet elsewhere in the same paper, he writes in a more realist vein, as if function perspectivalism generates mechanism *descriptions*, yet there is still some hierarchical mechanistic organization properly in the world, which these descriptions are perspectives *on* (2013: 153-4). This latter reading is more consistent with Craver's endorsement of the ontic view of mechanistic explanation (Craver 2014; cf. Illari 2013), on which a (mechanistic) explanation is the actual physical system that produces a phenomenon, not just a description or model of that system (the more typical epistemic view of explanation). Whatever his view on the metaphysics of mechanisms

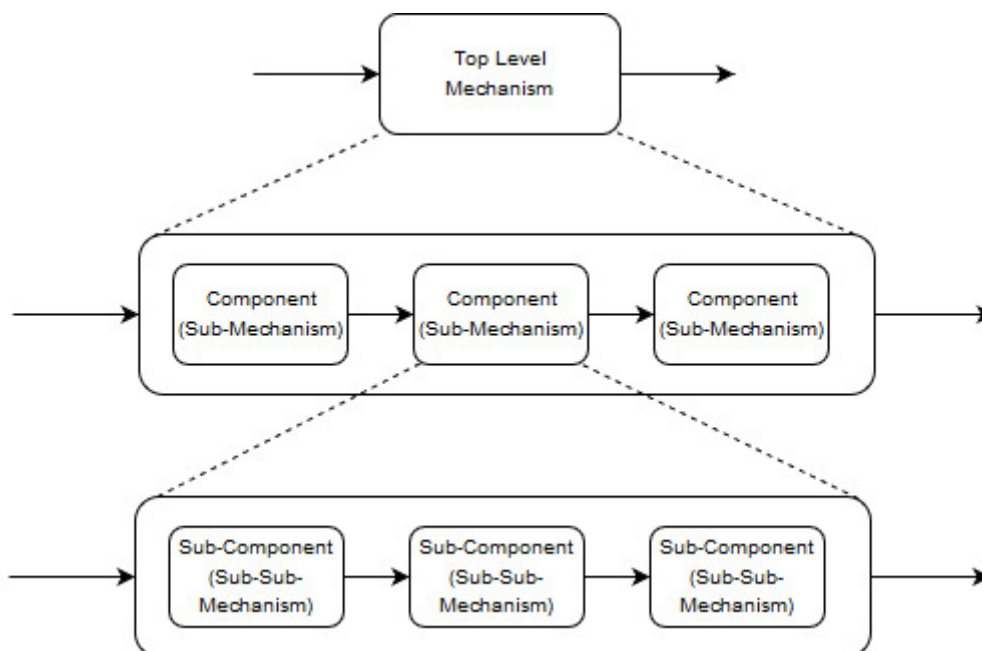
*per se*, Craver explicitly argues that the *hierarchical organization* of a mechanism into levels is “defensible as a metaphysical picture of how phenomena studied in the special sciences are constituted” (2015: 1), and it is realism of this kind that we will argue is incompatible with function perspectivalism.

Before moving on, we should first say a little more about the vexed distinction between epistemic and ontic explanation (see e.g. Illari 2013 for an overview). The tension we present below, between what we call ‘mechanism hierarchy realism’ and function perspectivalism, is most obviously applicable to ontic mechanists, who are committed to the actual mechanism itself being *constitutive* of a mechanistic explanation. Perspectivalism about the hierarchical structure of a mechanism would therefore constitute a more general explanatory perspectivalism. Insofar as our exemplar target, Carl Craver, is also an ontic mechanist (see e.g. his 2014), this might be sufficient for our argument to have bite. However, we also think that a version of the argument should concern epistemic mechanists, who say that it is a *description* of a mechanism that is explanatory, rather than the actual mechanism itself. They should be concerned because, if the mechanism hierarchy is perspectival, then the description of that hierarchy will also depend on our explanatory perspective, again resulting in an explanatory perspectivalism. In both cases maintaining a realist stance will involve either rejecting function perspectivalism, or accepting something less than full-blown realism about the mechanism hierarchy (an option we will explore further in section 6). We turn now to an elucidation of what we are calling ‘mechanism hierarchy realism’, bracketing for the time being the question of whether mechanistic explanations are epistemic or ontic.



### 3. Mechanism Hierarchy Realism

The standard definition of mechanisms as sets of entities engaging in organized activities does not explicitly take a stand on how different mechanisms may be related to each other, but in practice, all mechanists have emphasized the arrangement of mechanisms into hierarchies, such that the components of any mechanism may themselves be treated as mechanisms at a finer grain, and decomposed into lower level entities and activities. Likewise, any mechanism may itself be treated as a unitary component, or an entity in a higher-level, more coarse-grained mechanism. The picture is one of arbitrarily many nested levels of entities, each engaging in characteristic activities at some fixed level of the hierarchy (Craver 2001; 2015; Glennan 2005; see fig. 1). Call realism about the hierarchical arrangement of mechanisms into levels *mechanism hierarchy realism*, or MHR. MHR is the view that, once we have specified which mechanism we are interested in, there is some fact of the matter about the hierarchical structure of that mechanism, i.e. how it ought to be decomposed and which components and subcomponents will appear at which level of the hierarchy.



**Figure 1.** A schematic version of the hierarchical structure of mechanisms. Black arrows indicate causal interactions, whereas dotted lines indicate the composition relation between levels. As typically understood, there can be no causal interaction *between* levels, and a level is defined in terms of the causal interactions between its parts (components).

It is important to emphasize the MHR is a popular position, endorsed by almost all mechanism realists. None of the paradigmatic examples of mechanistic explanations discussed in the literature remain contained on a single level. Rather, the construction of an effective mechanistic explanation is taken to involve both ‘looking down’ into lower levels of a mechanistic hierarchy and ‘looking up’ into its higher levels (Bechtel 2009). In general, analyses of mechanistic function cascade through these levels, as the attribution of a functional role to an entity at one level of the hierarchy will determine which entities and activities are identified as components of that entity’s mechanism at the next level down. This point was already made by Cummins: identifying the heart as an entity performing the function of pumping within the circulatory system is precisely what allows one to pick out some subcomponents and activities within the heart as relevant in a functional analysis.

This seeming complementarity between function attribution and hierarchical (de)composition will play a key role in our argument below. First, however, we must flesh out the notion of MHR with two clarifications.

It is critical to recognize that MHR does not imply *universal* hierarchy realism (UHR). UHR holds that the entirety of nature is neatly organized into a hierarchy of levels of scale and interaction, a view most famously endorsed by Putnam & Oppenheim (1958). Although once widely held, this view has been largely discredited by detailed scientific examples, especially from the biological sciences (cf. Potochnik 2017). All sophisticated forms of MHR clearly emphasize that mechanisms are a purely local affair, and thus their hierarchical structure is also purely local. In fact, many mechanists are strongly opposed to making any claims at all about global hierarchies, or about how the levels of one mechanistic hierarchy might relate to those of another (spatially distinct) mechanism (see e.g. Povich & Craver 2017, and see Eronen 2015 for some critical discussion). Consequently, MHR is not committed to any claims about how the levels of mechanisms in one spatiotemporal region, or performing one type of function, may or may not ‘match up’ with the levels of another, spatiotemporally distant mechanism, performing a radically different function. Our discussion throughout concerns only MHR in this localized sense.

Second, MHR is a thesis about *both* mereology and causation. At first blush, it may seem that the talk of entities decomposing into mechanisms at lower levels implies only a mereological hierarchy. However, the ‘activities’ within each level of a mechanism should be understood as highlighting the relevant aspects of its internal causal structure (i.e., for producing the target phenomenon), and because mechanists typically take causation to be strictly intra-level (see below), this structure is also constrained by MHR. This means that a mechanistic analysis of causal structure is also going to be an analysis of hierarchical structure, and vice versa (we make this claim more precise and explore its metaphysical

implications in Section 6). The strict commitment to accepting only intra-level causes is confirmed in the complete rejection of inter-level causation by many mechanists (Craver and Bechtel 2007, Romero 2015, Baumgartner and Casini 2017, Kaiser and Krickel 2017). Most notably, Craver and Bechtel (2007) argue that any putative cross-level causation is always reducible “without remainder” to a combination of composition relations across levels and intra-level causation. For instance, the saturation of photopigment in cone cells does not *cause* the eye to transduce light into neural signals; rather, photopigment is a constitutive, lower level component of the mechanism of the eye, which as a whole, at a higher level, participates in an intra-level causal interaction between light and electrical activity in the brain (Craver & Bechtel 2007, 555). We call this the Craver-Bechtel (C-B) principle. The C-B principle provides MHR with a constructive strategy for subsuming putative counterexamples into a single mechanistic hierarchy, by explaining away apparent inter-level causes as mere compositional relations. However, when coupled with function perspectivalism, it threatens MHR from another direction, as mechanistic levels defined in terms of causal relations may also come to be perspectival.

#### **4. Levels Perspectivalism**

At first pass, function perspectivalism appears to be a new take on a common position in the epistemology of mechanisms: we cannot access mechanisms directly because they are too complex, but we may build models of them, employing simplifications and idealizations, and then infer features of the larger mechanism in the world from the (e.g. predictive) success of these models (Glennan 2005; Darden 2007; Bechtel 2016). Yet function perspectivalism is distinctive in explicitly rejecting the reality of a feature of the model that plays a constitutive role in its construction. Functions are not introduced late in the modeling process as simplifications, idealizations, or useful fictions, they are assumed at the start and guide

decisions about its ontology and scope. It is fair to ask, then, if perspectivalism about functions implies perspectivalism about other features of the model, and if it does, whether it follows that we should be anti-realists about these features as well. Here we establish that perspectivalism about functions implies perspectivalism about the hierarchical organization of mechanisms. Consequently, we cannot infer from the success of a mechanistic model that the hierarchical structure it posits is a feature of its target in the world, thus threatening MHR.

The basic argument is as follows. Function attribution is required to specify intra-level causal structure. Once intra-level causal relations are fixed, inter-level composition relations are constrained to induce a hierarchy of levels. However, since function attribution is perspectival, a different perspective may fix different causal relations as intra-level, which in turn induce their own hierarchy of composition relations. Since hierarchical structure is derived from intra-level causal structure (and not vice versa), and since intra-level causal structure is fixed by perspectival function attribution, then hierarchical structure is perspectival as well. Furthermore, nothing about the procedure for constructing mechanistic models ensures that their respective, perspectival levels will ‘match up’ — in principle, hierarchies may cross-cut each other. Consequently, the perspectivalist should be anti-realist about levels (just as she is anti-realist about functions).

*Function attribution specifies intra-level causation.* Craver identifies three kinds of function attribution in mechanistic model building: etiological, input-output (I-O), and contextual. It is the second of these that is critical for determining intra-level causation, namely “function [as] a mapping from inputs to outputs in conformity with a rule” (2013: 149). To attribute an I-O function to a putative mechanism component is to identify some causal influences (‘inputs’) and some characteristic activities (‘outputs’) as constitutive of that component’s role in producing the explanandum. This I-O attribution thereby assigns these putative causal relations a special status in the context of model construction, because

the C-B principle defines mechanistic hierarchies in terms of intra-level causation. Since the entities that input to a component, the activities that component outputs, and the entities impacted by those activities all interact causally, they are thereby constrained to fall on the same level of the mechanism, as ontological (and compositional) equals.

*Intra-level causal relations induce mechanistic hierarchy.* Once a single level of a mechanism is fixed, the complementary processes of ‘looking up’ and ‘looking down’ elaborate it into a hierarchy, in a manner heavily constrained by the causal structure functionally attributed to the reference level. Looking down, each component of the reference level that performs an I-O function may itself be broken down into subcomponents performing subtasks of that function. Crucially, the spatial and functional scope of these subcomponents within the mechanism is constrained by the initial analysis:

Isolated descriptions of an X’s  $\phi$ -ing specify the activity for which a lower-level mechanism will be sought and so fix the active, spatial, and temporal boundaries of that mechanism. (Craver 2001: 65; cf. 2013: 151)

Likewise, looking up, one may treat a pattern of causally interacting entities at one level as constituting a single entity (component) at a higher level. This higher-level entity itself performs an I-O function, analyzable in terms of intra-level causal relations at that higher level (Craver 2013: 152). So, composition determines a hierarchy:

levels of mechanisms are defined fundamentally by the relations question: by the componency relationship between things at higher and lower levels. (Craver 2015: 19)

Nevertheless, since compositional relations at lower and higher levels are constrained by the entities at some initial level, and since these initial ontological boundaries are determined by a perspectival functional analysis, this hierarchical structure is ultimately constrained by that perspective as well. The upshot of this line of reasoning is that function perspectivalism implies level perspectivalism, because different functional perspectives will privilege different causal relations, and thus (following the C-B principle) change the relative levels of the hierarchy at which components are placed.

It's worth reflecting briefly on why functional analysis is prior to mereological analysis: couldn't one first parse the putative mechanism into relevant parts, and only after attribute functions to those parts? If that were the case, then compositional structure would be independent of functional analysis, and so nothing about the perspectivalism or not of hierarchical structure would follow from function perspectivalism. There is a strong intuition that beginning from mereological analysis should be possible. Consider a paradigmatic mechanism, e.g. a grandfather clock: aren't the gears, chains, weights, springs, face, and cabinet all easily identifiable as hierarchically arranged components, even if we have no idea whatsoever how a clock works, or what it is for? But this intuition is misguided, for it ignores the substantive role played by judgments of similarity and relevance in any such decomposition. Why is the suggested breakdown different from one that decomposes the clock into left and right halves, or metal components and wood components, or small parts and large parts? In fact there are too many 'natural' ways to decompose the clock into a hierarchy of parts, each assuming a different answer to the question of what features of its physical composition are relevant for determining part boundaries and composition relations — mereology is wholly unconstrained without some principle of relevance by which to distinguish those properties that contribute to part individuation from those that do not. Functional analysis provides this principle of relevance, at least so far as mechanistic

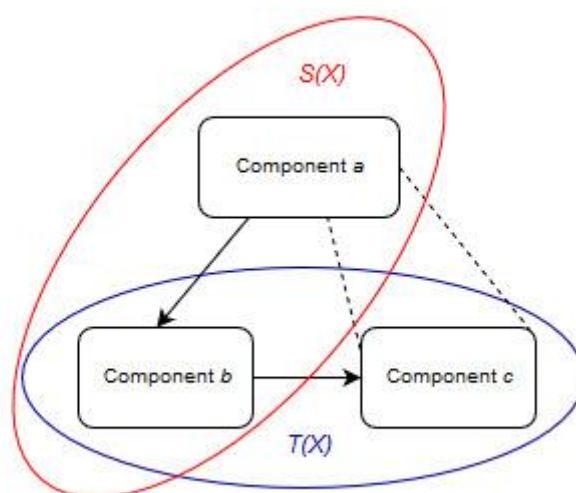
explanation is concerned. Attributing the function of doorstop, or *objet d'art*, or timekeeping device to the clock will distinguish different properties as relevant, thereby initiating different mereological decompositions.

*Hierarchy perspectivalism implies levels anti-realism.* So, functional analysis is prior to, and substantively constrains, hierarchical structure; consequently, function perspectivalism implies hierarchy (or levels) perspectivalism. Nevertheless, it does not necessarily follow that we should be anti-realists about hierarchical structure. For instance, we might have reason to think that the many perspectival hierarchies we derive from different function attributions to the same system are all consistent, i.e. that they may be embedded or otherwise combined to form a single, more elaborate hierarchy, one that is ultimately non-perspectival. We take this to be the position articulated by Craver when he asserts that “mechanisms are susceptible to multiple nested decompositions” into levels (2015: 19), i.e. there may be more than one way of specifying the mechanistic hierarchy, but each will be nested within one another. This seems to rule out the possibility that levels may *cross-cut* each other, i.e. that entities on two distinct levels in one mechanism may occur on the same level of a different mechanism, since cross-cutting hierarchies are not embeddable, or ‘nested’.

The possibility that mechanisms for spatially or functionally distant systems may cross-cut is already acknowledged by most mechanists, and captured by the slogan that mechanistic hierarchies are merely ‘local’ (Craver 2007, chapter 5; Eronen 2015). The worry here is different, namely that distinct explanatory perspectives may generate cross-cutting mechanistic descriptions *of the very same system or phenomenon*. If perspectivalism about hierarchies allows cross-cutting of the same physical system, then one cannot infer from hierarchical organization in a particular model to the uniqueness of that organization in the world, thus threatening mechanism hierarchy realism.



Yet, in principle at least, function perspectivalism *does* allow cross-cutting hierarchies of the same physical system. Suppose different explanatory perspectives on some system assign cross-cutting causal relations; by the argument above, these different causal analyses induce cross-cutting hierarchies. The C-B principle shows how this single, cross-cut system is nevertheless logically coherent, by generating a translation between these two hierarchies. Suppose, for instance, that two different I-O analyses,  $S$  and  $T$ , of a single system  $X$  attribute causal relations that cross-cut each other — i.e. by the lights of  $S(X)$ ,  $a$  causally interacts with  $b$  and  $c$  is a lower level component of  $a$ , while by the lights of  $T(X)$ ,  $b$  causally interacts with  $c$ , which is itself a component of the higher level entity  $a$  (fig. 2). Taking inspiration from Craver and Bechtel, we can define the *apparent* causal relations in a mechanism  $\mathfrak{C}(\cdot)$  as the set of all intra-level causes, plus all combinations of cause and composition relations (in either direction). Then it is easy to see that  $\mathfrak{C}(S(X)) = \mathfrak{C}(T(X))$ , and consequently both hierarchies may plausibly have been derived from a single system through I-O analysis, i.e. function perspectivalism is consistent with cross-cutting hierarchies, and thus with anti-realism about hierarchical structure in the world, thus threatening MHR.



**Figure 2.** Components *a* and *b* causally interact for  $S(X)$ , and thus sit at the same mechanistic level, while *b* and *c* causally interact, and thus occupy the same level, for  $T(X)$ . Therefore the componential hierarchy of one perspective,  $S(X)$ , cross-cuts that of another,  $T(X)$ .


So, function perspectivalism implies level perspectivalism, and level perspectivalism is consistent with anti-realism about hierarchical organization. Nevertheless it may be that, contingently, all target systems in nature are hierarchically organized in a non-cross-cutting manner, such that all (empirically successful) perspectival attributions of function to real world systems will happen to be consistent with a single hierarchy, i.e. susceptible to nesting. All that we have shown here is that, once one accepts function perspectivalism, the ability to model a system mechanistically does not provide any principled reason for supposing that the model's hierarchical structure is really 'out there' in the world.

## 5. Cross-Cutting Levels in Practice

So, the would-be function perspectivalist/mechanism realist may not rely on a principled argument from model success to mechanistic hierarchy in the world, but they might hold out hope that nature is contingently hierarchical in the right way to license mechanism realism.

This hope is defeated if there are actual cases where different explanatory perspectives carve up a single system into mechanistic levels that cross-cut one another. We suspect such cross-cutting hierarchies are common; in this section we examine one example in detail. We take this example from neuroscience, often advanced as a paradigm domain for mechanistic explanation (see e.g. Craver 2007), and thus a salient testing ground for worries about mechanism realism.

A frequently rehearsed example of a mechanistically explained neural phenomenon is synaptic firing, i.e. the passage of an electrical impulse from one neuron to another via a chemical interaction at the synaptic gap (Craver 2007: 135-8; Machamer, Darden, and Craver 2000: 8-9; cf. Shepherd 1994; see fig. 3). The I-O function attribution in this case is transfer of electrical activation from one cell to another. Identifying this function isolates the pre-synaptic neuron, the post-synaptic neuron, and the molecular neurotransmitters released by the first and binding to the other as all interacting causally, and thus at the same level of a mechanistic hierarchy. The interaction between these components produces the overarching phenomenon of synaptic firing at a higher level of this hierarchy. Synaptic firing cannot be reduced to the behavior of one or another of the components at the lower level, as it requires more than just one neuron. Rather a ‘synapse’ is a gap bridging two neurons, signaling between which is enabled by the presence of neurotransmitters. Looking further ‘down’, we might expand our story about neuron–neuron interaction by examining the role of the proper subcomponents of each neuron in generating synaptic behavior, for instance, receptor molecules that bind the neurotransmitter, or ion channels that induce the actual flow of current across the cell wall. So far so good; this kind of description is precisely the sort of explanation that mechanistic decomposition is meant to be well-suited for.



**Figure 3.** A schematic decomposition of the mechanism responsible for synaptic firing, where black arrows indicate causal interactions, and dotted lines indicate the constitution relation between levels of the hierarchy.

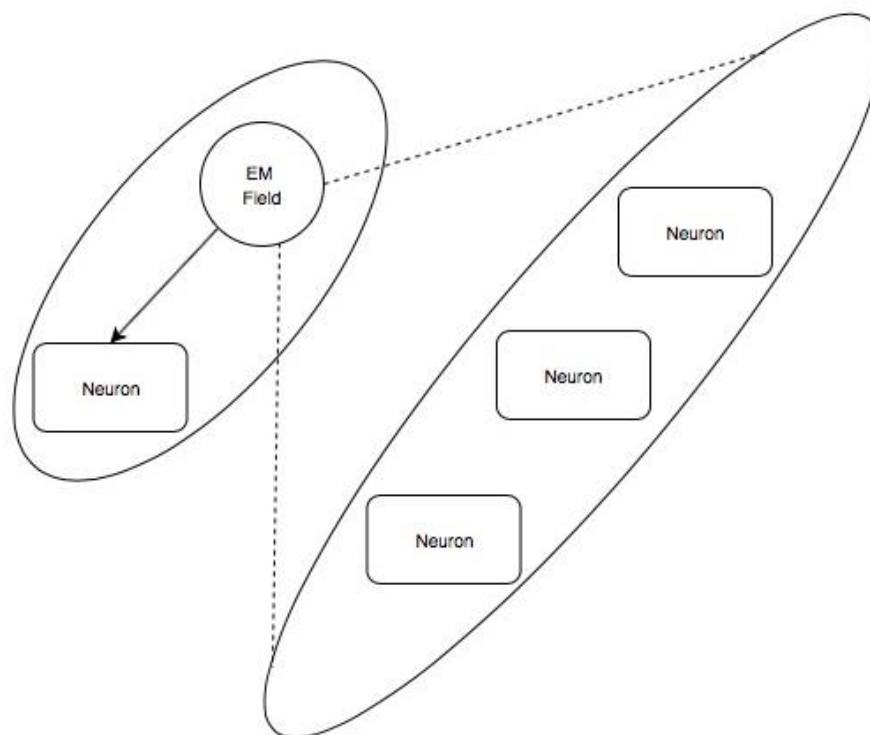
Now consider another, less well-known neural phenomenon: ephaptic coupling.

Ephaptic coupling is the process by which an electromagnetic field can have an effect on the firing rate of an individual neuron — in a typical case, a field generated itself by the firing of another group of neurons (see e.g. Anastassiou *et al.* 2011). The possibility of ephaptic coupling was experimentally demonstrated in the 1940’s (Arvanitaki 1942), and has since then been studied both *in vitro* and *in vivo*. A point of continued contention concerns whether ephaptic coupling plays any significant role in normal brain function. It is relatively uncontroversial that it can play a role in pathological conditions — for instance, ephaptic coupling amplifies random activity, allowing disruptive neural “cross-talk” between demyelinated neurons in trigeminal neuralgia (Love and Coakham 2001). There is also some evidence for a non-pathological role; for instance, it has been demonstrated *in vivo* that ephaptic coupling affects functionally relevant features of neurons, such as spike timing and threshold of neural response (Anastassiou *et al.* 2011); furthermore, recent work has offered evidence that ephaptic coupling functions to regulate periodic neural activity (Chiang *et al.* 2019).

From the standpoint of the function perspectivalist, it does not really matter whether ephaptic coupling does or does not perform a critical function in the brain. We may take the

perspective that it does, and thereby provide a mechanistic explanation such as that developed by Chiang *et al.* (2019). Low frequency periodic activity in hippocampus is believed to play a role in memory consolidation (Dickson 2010). Chiang *et al.* (2019) observed the propagation of low frequency periodic neural activity of this sort in hippocampal tissue that had been severed in order to ensure synaptic coupling played no role. They then developed a computer model demonstrating how this propagation may be explained through ephaptic coupling induced by an endogenous electromagnetic field (Chiang *et al.* 2019: 261). Their study seems to be a clear case of mechanistic explanation, decomposing a system into components and their characteristic interactions, and yet it generates a causal hierarchy that cross-cuts the standard synaptic transmission mechanism outlined above. In order to see why, we will need to investigate the ephaptic coupling mechanism in a little more detail.

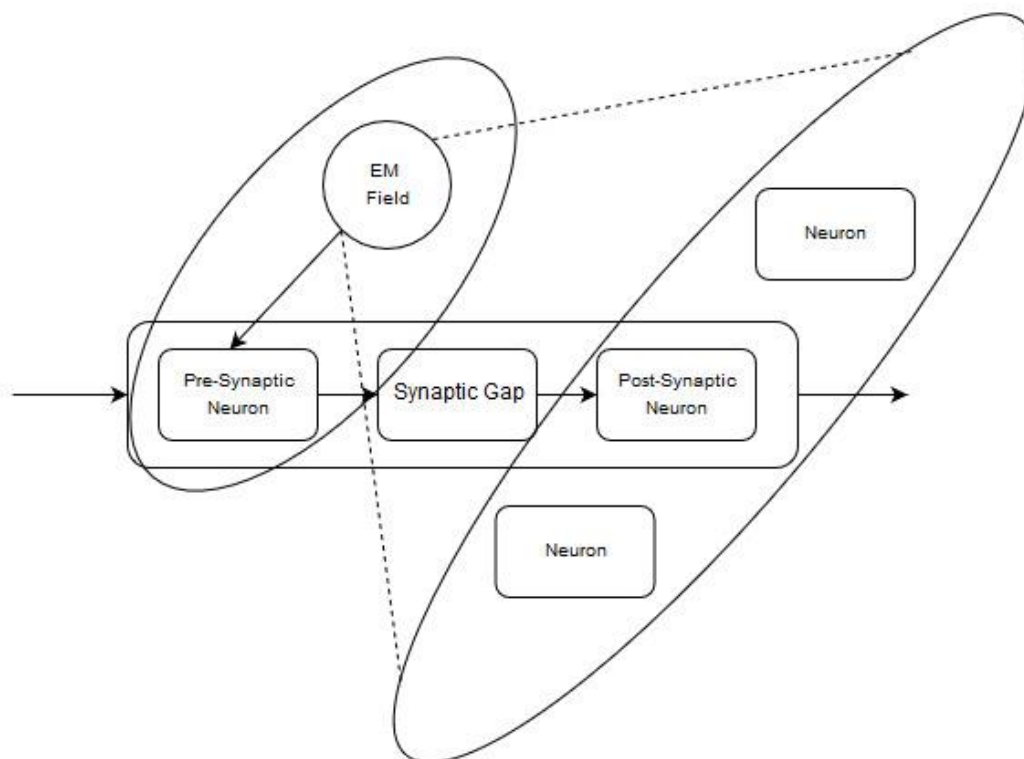
An endogenous electromagnetic (EM) field is generated by the coordinated behavior of a group of neurons, and is not reducible to the behavior of individual neurons, nor a mere aggregate of that behavior. Rather it is a holistic phenomenon, depending for its constitutive properties on the neurons' coordinated interaction. In other words, the EM field occurs at a higher level of the mechanistic hierarchy than the generating neurons themselves. However, the causal interaction at issue occurs between this EM field and a single, target neuron, the one entrained by it. In other words, because ephaptic coupling occurs between an EM field and a neuron, both that neuron and the field should appear on the same level of the corresponding mechanistic hierarchy (see fig. 4).



**Figure 4.** A schematic decomposition of the mechanism responsible for ephaptic coupling. As in the previous diagram, the black arrow indicates a causal interaction between an electromagnetic field and a neuron, while the dotted lines indicate that the field is constituted by the activity of some other set of neurons. Each large oval designates distinct level of the hierarchy, presented at a slanted angle for reasons that will soon become clear.

In principle, both of these mechanisms (synaptic firing and ephaptic coupling) may be active at the same time, involving some of the same components. Even if the functional relevance of ephaptic coupling is not yet well-understood, it is still nonetheless a genuine phenomenon with a plausible mechanistic explanation. Thus, we have here a single system in which two mechanisms cross-cut one another, with a set of components that are situated at one level in the first mechanism appearing at two different levels in the second (see fig. 5). Relative to the perspective of synaptic firing, the pre- and post-synaptic neurons are situated at the same mechanistic level, as demonstrated by the fact that they interact causally with one another; relative to the ephaptic perspective, however, these very same neurons might be

situated at distinct levels, as the electromagnetic field acting on the pre-synaptic neuron (for example) may be constituted by the firing of some other set of neurons that includes the post-synaptic neuron. To complicate matters further, it could even be the case that one and the same neuron is both involved in generating an electromagnetic field and in having its firing rate influenced by that field, thereby simultaneously occupying two mechanistically distinct ‘levels’. Thus, there is no ‘objective’ fact about whether these neurons are on the same or different levels, as their position in the hierarchy can only be determined relative to an explanandum phenomenon (either synaptic firing or ephaptic coupling), and the perspectival function attribution it induces.



**Figure 5.** The two mechanisms cross-cutting one another. Here the post-synaptic neuron is part of the group of neurons whose activity constitutes the electromagnetic field, which in turn exerts a causal influence on the pre-synaptic neuron. Entities at a single level of the synaptic mechanism appear at two different levels of the ephaptic mechanism.

The phenomena of ephaptic coupling and synaptic firing provide a clear case of mechanistic decomposition with cross-cutting levels, where two token entities (the pre- and post-synaptic neurons) within one physical system can be described at either the same level or different levels depending on the phenomenon of interest. This is not a case of comparing levels across spatially distinct mechanisms, such as when Craver argues that it does not make sense to ask whether “a Na<sup>+</sup> ion in the hippocampal LTP mechanism is at the same level as a Na<sup>+</sup> ion in the retinal mechanism of phototransduction” (Eronen 2015: 44; cf. Craver 2007: 191). Rather the issue is that there can be multiple different decompositions *of one and the same system*, each of which serves a distinct explanatory purpose, but whose mechanistic hierarchies cross-cut one another, with an entity at one level in one hierarchy at a higher (or lower) level in another. Thus, this example shows that hierarchical organization is not objective, but rather perspectival, and so, insofar as we accept the C-B principle, and take mechanisms as constitutively hierarchical, the presence of mechanisms in the world is also not a matter of objective fact, but of perspective.

Finally, it is worth reminding ourselves why this is a problem for the function perspectivalist in particular, rather than for mechanists in general. The crucial point is that, for any mechanist who endorses realism about functions, there is space to differentiate between these two putative mechanistic explanations by distinguishing real from apparent functions. A selectionist might claim (for instance) that synaptic firing was selected to perform the function of passing information between neurons. Then, the mechanistic hierarchy induced by the (correct) attribution of this function is actually out there in the world, and can participate in ontic explanations. In contrast, they might argue that ephaptic coupling is also a real phenomenon, but not one that has been selected for, and thus should not be accorded the same status as the (true) mechanistic organization of the brain. The perspectivalist, however, rejects the idea that any kind of function attribution might be more



‘proper’ than any other, and thus cannot make this move. For her the relevant function attribution (synaptic firing or ephaptic coupling) will depend on our explanatory aims and interests. Since function is necessary to determine causal interactions, and causal interactions are necessary to determine mechanistic hierarchy, then this hierarchy too will depend on our explanatory aims and interests. Logically, this result follows even when perspectival hierarchies do not cross-cut; what this section demonstrates is that cross-cutting does actually occur, and so even the optimistic hope of finding contingent hierarchy in the world ultimately fails. Consequently, the function perspectivalist must also endorse hierarchy perspectivalism, and thus abandon MHR.

## 6. What Realism is Left?

Where does all this leave the would-be perspectival realist, i.e. someone who wants to endorse perspectivalism about functions but still remain a realist about mechanisms in some sense? If hierarchical organization is constitutive of mechanistic structure, and function perspectivalism rules out realism about hierarchies, then the function perspectivalist cannot consistently be a *mechanism hierarchy* realist. Rather, she can at best be realist about parts of her model that together do not add up to a full-fledged mechanistic hierarchy. However, there is still quite a bit of substantive structure one may be realist about that does not depend on the reality of the mechanistic hierarchy, and is thus still available to the function perspectivalist.

One aspect of the realist picture that is untouched by the above discussion is the parsing of the world into entities and activities. We’ve given no reason to suppose that function perspectivalism implies perspectivalism about entities or their interactions. But the mechanism realist requires more than just entities and activities, she requires those entities and activities be organized *in a specific way*. The existence of entities and activities alone is consistent with a thoroughgoing reductionist worldview, on which only the most fundamental

entities and relations really exist. Such a view is only ‘mechanistic’ in a vacuous sense, as it exhibits no structure other than brute causality (Craver and Tabery, 2015, for instance, explicitly rule out appeals to fundamental laws or causes as mechanistic, since they admit no further decomposition). Function perspectivalism need not mandate thoroughgoing reductionism, and it permits some degree of objective organization in the world, just not as much organization as is required for realism about mechanistic hierarchies.

To see exactly how much realism the function perspectivalist is left with, it will help to describe the constituent features of a mechanism in a somewhat different way than is usually done. In particular, we need to distinguish constraints on the mechanistic view of *composition*, i.e. the relationship between entities understood at different degrees of granularity, from constraints on the mechanistic view of *interaction*, i.e. the types of (typically causal) relationships that may occur between entities. In an abstract, formal sense, both composition and interaction may be thought of as *relations* between entities, characterized by certain properties. A mechanism may be understood as a structure in which the interplay between these two relations (compositional and causal) is severely constrained, with causal interactions occurring only between entities at the same level of composition, and levels of composition being defined in terms of the possibility of causal interactions. The possibility of cross-cutting hierarchies within a single system shows that this condition of mutual constraint must be rejected by the function perspectivalist, as she must countenance cases like that we considered in the previous section, where two interacting entities may sit simultaneously at the same and different ‘levels’, depending on which explanatory perspective one adopts. The function perspectivalist may nevertheless still endorse a picture of the world on which both composition and interaction are separately constrained, and thus on which the world is quite rigidly structured — it is just that this structure is not mechanistic

per se (or at least, not prior to our adoption of a mechanistic explanatory perspective). To see this, we need to first introduce a few formal details.

Let  $<$  represent the composition relation, i.e. read  $a < b$  as ‘ $a$  is a component of  $b$ ’. As mechanists understand it,  $<$  is something like the proper parthood relation studied by mereologists. Proper parthood is typically thought to be both *irreflexive* ( $not\ a < a$ ) and *antisymmetric* ( $a < b$  and  $a \neq b$  implies  $not\ b < a$ ), i.e.  $<$  is *asymmetric*: if  $a$  is a component of  $b$ , then  $b$  is not a component of  $a$  (Koslicki, 2008; McDaniel 2009). This is an assumption about composition that corresponds to the strict hierarchy of layers of entities posited by mechanists. Affirming that  $<$  is asymmetric constitutes a substantive position, rejecting, for instance, arguments that proper parthood might be susceptible to mereological loops, and thus not asymmetric (Tillman and Fowler 2012). Furthermore, mereologists typically consider proper parthood to be *transitive* ( $a < b$  and  $b < c$  implies  $a < c$ ), yet composition as understood in the strict hierarchical sense by mechanists is plausibly not transitive, as (intuitively) if an entity is a component of a sub-mechanism, it is not (necessarily) also a component of the mechanism of which that sub-mechanism is a part (Craver 2015: 14). This substantive view of composition is not directly undermined by function perspectivalism, so the would-be perspectival realist may consistently continue to assert it. Even in the case of a cross-cutting hierarchy, like that which we discussed in the previous section, there need be no entity that is a part of itself, nor any pair of entities that are parts of one another. It is only the strict ordering of entities at one (causal) level and not another that this example rules out, not any of the more fundamental or intuitive constraints on composition.

Next, let  $\sqsubset$  represent the interaction relation, i.e. ‘ $a$  (causally) interacts with  $b$ ’. Like  $<$ , one might consider  $\sqsubset$  to be *irreflexive*:  $a$  does not interact with itself. Unlike  $<$ ,  $\sqsubset$  might not be anti-symmetric, though: if  $a$  and  $b$  are entities, and  $a \sqsubset b$ , we might also wish to allow that  $b \sqsubset a$ . The exact nature of  $\sqsubset$  is less specified in the mechanism literature than that

of  $<$ ; nevertheless,  $\sqsubset$  could be fleshed out in a variety of ways, drawing for instance on the literature on the logical structure of causality. One issue up for debate, for instance, is whether interaction should be thought of as transitive — a property endorsed by some in the logic of causation literature (Koons 1999) and denied by others (Scheffler 1993). The point is just that, as for  $<$ , there is room for the would-be perspectival realist to commit herself to a substantive position concerning the interactive structure of the world, without in any way contradicting her perspectivalism about functions (and thus, mechanistic hierarchies). She can be a realist about causes, and make substantive claims about which patterns of causation are permitted and which not; she just can't (consistently) maintain that causal interactions are restricted to occurring between entities at the same level of a (compositional) hierarchy.

What this line of reasoning reveals is that the ontological commitments of mechanism realism are actually quite substantive and specific, and that there is a wide variety of realist positions on the compositional and causal-interactive structures of the world that fall short of full-blooded mechanism realism, and are thus still available to the function perspectivalist. What mechanism realism adds to asymmetric, intransitive composition and irreflexive, (potentially) symmetric interaction is a restriction on how composition and interaction are permitted to combine, i.e. the mutual constraint that we mentioned previously. To state this restriction formally, we need a few more definitions. Write  $<_{\text{T}}$  for the transitive closure of  $<$ , and  $\sqsubset_{\text{ST}}$  for the symmetric, transitive closure of  $\sqsubset$ . The 'family' of  $a$ ,  $F_a$ , is the set of all entities  $x$  such that  $a <_{\text{T}} x$  or  $x <_{\text{T}} a$ . Crucially, note that, since composition is irreflexive,  $a \notin F_a$ . The "peer group" of  $a$ ,  $[a]$ , is the set of all  $x$  such that  $a \sqsubset_{\text{ST}} x$ . Since  $\sqsubset_{\text{ST}}$  is symmetric,  $a \in [a]$ . Intuitively, the family of  $a$  are all entities other than  $a$  related to it by some chain of composition relations, while the peer group is the set of all entities  $a$  is related to by some chain of interactions, i.e. all entities at the same 'level' as  $a$ . Then, the fundamental constraint on composition and interaction endorsed by mechanism realism may be stated as:

$$\forall x(F_x \cap [x] = \emptyset)$$

This constraint rules out cross-level interactions. To see why, suppose  $a$  and  $b$  are on different levels of a mechanistic hierarchy, yet  $a \sqsubset b$ . If  $a$  and  $b$  are on different levels, then (without loss of generality) there is some  $c$  such that  $b \sqsubset_{ST} c$  and  $c \in F_a$  (by the definition of level). But then, by the transitivity of  $\sqsubset_{ST}$ ,  $a \sqsubset_{ST} c$ , and so  $c \in [a]$  as well, and the constraint is violated. In other words, no member of the family of an entity may (causally) interact with it, and no member of the peer group of an entity may be in its family, or, in more familiar terms, causal interactions may only occur between entities at the same level, and entities must be at the same level if they causally interact. It is important to clarify here that  $x$  ranges only over the entities in some narrowly defined spatio-temporal region, rather than over all entities in the world — this is the force of the mechanist’s assertion that mechanistic hierarchies are purely local. Nevertheless, within this narrow region, the claims of mechanism realism are substantive, and are violated by the possibility of cross-cutting levels entailed by function perspectivalism.

Mechanists often state their position as if it is near trivial, feeling the need to defend it against the worry that everything constitutes a mechanism (Craver and Tabery 2015). Furthermore, they feel comfortable asserting that causal interactions or compositional relations that violate the strict mechanistic hierarchy are suspect, or draped in a “shroud of mystery” (Craver and Bechtel 2006, 562). Our purpose in stating the hierarchical commitments of the mechanism hierarchy realist in this new way is to shift intuitions on this issue. What we hope to have accomplished in this section is make it clear that, (1) the assertion of mechanism hierarchy realism is substantive, and severely restricts the permitted interplay between mereology and causality; and furthermore, (2) there are a wide range of

substantive metaphysical positions one might take on the causal and mereological structure of the world that are richer than full reductionism, yet not so restrictive as mechanism hierarchy realism. It is worth emphasizing that this is a result the mechanist should embrace: mechanism hierarchy realism is significantly more interesting as a claim about either explanation or the world if it rules out many alternatives, and thereby makes a substantive, positive claim.

Let us conclude by revisiting the implications of this analysis for the function perspectivalist. The function perspectivalist is an anti-realist about proper functions, i.e. she argues that there is no fact of the matter about what the function of a mechanism is, and that function attributions are made from an explanatory perspective for the purpose of constructing a mechanistic explanation. In section 4, we demonstrated that, because function attribution is necessary to determine the relevant causal interactions for a mechanistic explanation, and because it is these causal interactions that determine which entities are at the same level, then function perspectivalism implies levels perspectivalism, and is logically consistent with the existence of multiple explanations (of the same system) that posit cross-cutting hierarchical structures. In section 5, we offered an example from neuroscience of a system for which two cross-cutting mechanistic explanations are available; this illustrated that cross-cutting hierarchies are not merely a logical possibility, they do actually occur in empirically robust mechanistic explanations. So, the function perspectivalist has no license to conclude from a successful mechanistic explanation that the mechanistic structure it describes exists, objectively, in the world. What we have shown in this final section is that a mechanistic explanation nonetheless plausibly involves substantive (compositional and causal-interactive) structure that is left untouched by function perspectivalism. Although she would, of course, require some further, realist argument to justify attributing this structure to the world, the function perspectivalist could make this attribution without contradicting her

anti-realism about proper functions (and thus also mechanistic levels and hierarchies). The upshot of all this is that, although there is some objective structure in the world that may be inferred from a successful mechanistic explanation, it is potentially less substantial than a mechanism hierarchy realist would have you believe, and therefore compatible with a full-blown perspectivalism about mechanistic explanation.

## **7. Conclusion**

We have argued that function perspectivalism is inconsistent with mechanism hierarchy realism, because proper functions are required in order to determine the relevant causal interactions that define mechanistic levels and hierarchies, and without these there is no interesting sense in which we can say that mechanisms are ‘really’ out there in the world. The interesting sense of mechanism depends on there being strict mutual constraints between composition relations and causal interactions, such that the latter only occur between entities at the same level of a hierarchy. If function attribution is perspectival, then it is possible for cases to arise where two mechanistic hierarchies ‘cross-cut’ one another, such that an entity cannot unambiguously be located on a single level. We illustrated this concern with an example drawn from neuroscience (the paradigmatic home of mechanistic explanation), where the phenomena of synaptic firing and ephaptic coupling give rise to actual cross-cutting mechanistic hierarchies. Without function realism it is impossible to determine which mechanistic hierarchy receives ontic priority, and thus mechanism realism must be abandoned if one wants to maintain function perspectivalism. There is, however, a thinner sense of realism that is still available to the perspectivalist, consisting of real causal interactions and real compositional relations, but falling short of the strict hierarchical structure required for mechanism realism.

Depending on one's personal tastes and prior commitments, this result could also be taken as a *modus tollens* against function perspectivalism. If so, however, one will need to provide a realist account of proper functions, i.e. an explanation for how a mechanism comes to have a function independent of our explanatory interests, in order to then preserve the realist account of levels, hierarchies, and thus mechanisms themselves. Some such realist notion of proper function will also be needed if one is to consistently endorse an ontic account of explanation, claiming that explanations are provided not by our *descriptions* of mechanisms, but rather by the real mechanisms themselves, out there in the world. The function perspectivalist, on the other hand, may have to give up on ontic explanation, as without any real mechanisms there is nothing other than our descriptions (or models) of mechanisms that could fulfil the role of 'explanation' in 'mechanistic explanation'. This outcome is especially problematic for Craver, who has previously endorsed both function perspectivalism and ontic explanation. We will ourselves part way here, one of us taking the road '*modus ponens*', and seeking the right level of sub-mechanistic structure to endorse as a perspectival realist; the other taking the road '*modus tollens*', and resuming the search for the right naturalization of proper functions.



## References

Anastassiou, Costas, Rodrigo Perin, Henry Markram, and Christof Koch. 2011. "Ephaptic coupling of cortical neurons." *Nature Neuroscience* 14:217-23.

Anderson, Michael. 2015. "Mining the Brain for a New Taxonomy of the Mind." *Philosophy Compass* 10/1:68-77.

Arvanitaki, A. 1942. "Effects invoked in an axon by the activity of a contiguous one." *Journal of Neurophysiology* 5/2:89-108.

Baumgartner, Michael and Lorenzo Casini. 2017. "An Abductive Theory of Constitution." *Philosophy of Science* 84:214-33.

Bechtel, William. 2009. "Looking down, around, and up: Mechanistic explanation in psychology." *Philosophical Psychology* 22/5:543-564.

———. 2016. "Using Computational Models to Discover and Understand Mechanisms." *Studies in History and Philosophy of Science A* 56:113-121.

Bechtel, William and Adele Abrahamsen. 2005. "Explanation: a mechanist alternative." *Studies in History and Philosophy of Science C* 36/2:421-41.

Boorse, Christopher. 1976. "Wright on Functions." *The Philosophical Review* 85/1:70-86.

Chiang, Chia-Chu, Rajat Shivacharan, Xile Wei, Luis Gonzalez-Reyes, and Dominique Durand. 2019. "Slow periodic activity in the longitudinal hippocampal slice can self-propagate non-synaptically by a mechanism consistent with ephaptic coupling." *The Journal of Physiology* 597.1:249-69.

Craver, Carl. 2001. "Role Functions, Mechanisms and Hierarchy" *Philosophy of Science* 68:31-55.

———. 2007. *Explaining the Brain*. Oxford: OUP.

———. 2013. "Functions and Mechanisms: A Perspectivalist View." *Functions: Selection and Mechanisms*, ed. Philippe Huneman, 133-58. Dordrecht: Springer Netherlands.

———. 2014. "The Ontic Account of Scientific Explanation." In *Explanation in the Special Sciences*, eds. Marie I. Kaiser, Oliver R. Scholz, Daniel Plenge and Andreas Hüttemann, 27-52. Springer Verlag.

———. 2015. "Levels." In *Open MIND* 8, eds. Thomas Metzinger and Jennifer Windt. Frankfurt am Main: MIND Group.

Craver, Carl and William Bechtel. 2007. "Top-down causation without top-down causes." *Biology and Philosophy* 2:547-63.

Craver, Carl and James Tabery. 2015. "Mechanisms in Science." In *The Stanford Encyclopedia of Philosophy*, ed. Edward Zalta.

Cummins, Robert. 1975. "Functional Analysis." *Journal of Philosophy* 72:741-64.

Darden, Lindley. 2002. "Rethinking Mechanistic Explanation." *Philosophy of Science* 69(S3): 342-53.

———. 2007. "Mechanisms and Models." In *The Cambridge Companion to the Philosophy of Biology*, eds. David Hull and Michael Ruse. Cambridge, UK: CUP.

Dewhurst, Joe. 2016. "Gualtiero Piccinini: *Physical Computation*." *Philosophical Psychology* 29/5:795-7.

Dickson, Clayton. 2010. "Ups and downs in the hippocampus." *Behavioral Brain Research* 214:35-41.

Eronen, Markus. 2015. "Levels of Organization: A Deflationary Account." *Biology & Philosophy* 30/1:39-58.

Garson, Justin. 2013. "The Functional Sense of Mechanism." *Philosophy of Science* 80/3:317-33.

———. 2017. "A Generalized Selected Effects Theory of Function." *Philosophy of Science* 84/3:523-43.

Glennan, Stuart. 2002. "Rethinking Mechanistic Explanation." *Philosophy of Science* 69(S3):S342-S353.

———. 2005. "Modeling mechanisms." *Studies in History and Philosophy of Science C* 36/2:443-64.

———. 2017. *The New Mechanical Philosophy*. Oxford: OUP.

Hardcastle, Valerie. 1999. "Understanding functions: A pragmatic approach." In *When biology meets philosophy*, ed. Valerie Hardcastle, 27-46. Cambridge, MA: MIT Press.

Illari, Phyllis. 2013. "Mechanistic Explanation: Integrating the Ontic and Epistemic." *Erkenntnis* 78/2:237-55.

Illari, Phyllis, and Jon Williamson. 2012. "What is a mechanism?" *European Journal for the Philosophy of Science* 2/1:119-35.

Kaiser, Marie and Beate Krickel. 2017. "The Metaphysics of Constitutive Mechanistic Phenomena." *British Journal for the Philosophy of Science* 68:745-79.

Kästner, Lena. 2018. "Integrating mechanistic explanations through epistemic perspectives." *Studies in the History and Philosophy of Science* 68:68-79.

Koons, Robert. 1999. "Situation mereology and the logic of causation." *Topoi* 18/2:167-74.

Koslicki, Kathrin. 2008. *The Structure of Objects*. Oxford: OUP.

Love, Seth and Hugh Coakham 2001. "Trigeminal neuralgia: pathology and pathogenesis." *Brain* 124/12:2347-60.

Machamer, Peter, Lindley Darden, and Carl Craver. 2000. "Thinking about Mechanisms." *Philosophy of Science* 67(1):1-25.

Maley, Corey, and Gualtiero Piccinini. 2017. "A Unified Mechanistic Account of Teleological Functions for Psychology and Neuroscience." In *Explanation and Integration in Mind and Brain Science*, ed. David Kaplan, 236-56. Oxford: OUP.

McDaniel, Kris. 2009. "Extended Simples and Qualitative Heterogeneity." *The Philosophical Quarterly* 59:325-31.

Millikan, Ruth. 1989. "In defense of proper functions." *Philosophy of Science* 56:288-302.

Mossio, Matteo, Cristian Saborido, and Alvaro Moreno. 2009. "An organizational account of biological functions." *The British Journal for the Philosophy of Science* 60/4: 813-41.

Neander, Karen. 1991. "Functions as selected effects." *Philosophy of Science* 58/2:168-84.

Piccinini, Gualtiero, and Carl Craver. 2011. "Integrating psychology and neuroscience: Functional analyses as mechanism sketches." *Synthese* 183(3):283-311.

Potochnik, Angela. 2017. *Idealization and the Aims of Science*. Chicago, IL: University of Chicago Press.

Povich, Mark, and Carl Craver. 2017. "Mechanistic Levels, Reduction, and Emergence." In Glennan & Illari (eds.), *The Routledge Handbook of Mechanisms and Mechanical Philosophy*. Routledge.

Putnam, Hilary, and Paul Oppenheim. 1958. "Unity of Science as a Working Hypothesis." In Feigl *et al* (eds.), *Minnesota Studies in the Philosophy of Science*, vol. 2. Minneapolis, MN: Minnesota University Press.

Raerinne, Jani. 2011. "Causal and mechanistic explanations in ecology." *Acta Biotheoretica* 59(3-4): 251-71.

Romero, Felipe. 2015. "Why there isn't inter-level causation in mechanisms." *Synthese* 192:3731–3755.

Rosenbleuth, Arturo, Norbert Wiener, and Julian Bigelow. "Behavior, Purpose and Teleology." *Philosophy of Science* 10/1: 18-24.

Scheffler, Uwe. 1993. "On The Logic Of Event Causation." *Logic and Logical Philosophy* 1:129-55.

Shagrir, Oron, and William Bechtel. 2017. "Marr's Computational Level and Delineating Phenomena." In *Explanation and Integration in Mind and Brain Science*, ed. David Kaplan, 190-214. Oxford: OUP.

Shepherd, Gordon M. 1994. *Neurobiology*, 3<sup>rd</sup> edn. Oxford: OUP.

Tillman, Chris and Gregory Fowler. 2012. "Propositions and Parthood: The Universe and Anti-Symmetry." *Australasian Journal of Philosophy* 90/3:525-39.

Wilhelm, Isaac. 2019. "The Ontology of Mechanisms." *Philosophy of Science* 116(11): 615-36.