Scientific Explanation: Mechanistic, Ontic, Epistemic

Abstract

Conceptions of mechanistic explanation, following Salmon, are orthodoxically rendered as ontic. However, the ontic conception is fallacious. Fortunately, the so-called 'New Mechanists' can achieve a far more credible conception of mechanistic explanation, which both generalizes to any scientific explanation that New Mechanists care to analyze mechanistically, and which leaves intact the rest of their research program.

1. Introduction

1.1 Mechanistic explanation

Arguably, the focal issues of logical empiricism no longer haunt contemporary philosophy of science. Previous emphases on set-theoretic formalization and axiomatization, intertheoretic reduction, monotonic derivation from laws, etc. have ceded way to concerns over actual scientific practice, such as experimental design and causal modeling (Bickle, 2003). This change in focus is particularly evident for those focused on disciplines like systems biology and neuroscience. And it is within this context that mechanistic conception of scientific explanation has proven gripping.

Mechanistic explanation is simply explanation in terms of mechanisms, conceived of as hierarchical systems. Given this point of departure, the research program of the so-called 'New Mechanists' has three main foci. One is to develop a comprehensively adequate conception of mechanistic explanation.[1] Another is to analyze the nature of hierarchical systems, particularly as exemplified by cases from different scientific disciplines, in terms of which any mechanistic explanation proceeds.[2] A third focus concerns strategies for discovery and experimentation.[3]

The most recent exemplification of this research program, in which all three foci converge, is Craver's (2007) landmark achievement in

---

[1]   See, e.g., Bechtel & Abrahamsen (2005); Bechtel & Wright (2008); Craver (2007); Cummins (2000); Delehanty (2005); Glennan (2002, 2005); Hegarty (1992, 2004); Hegarty et al. (2003); Waskan (2006); Wright (2002, 2007); Wright & Bechtel (2006). Railton (1978) and Salmon (1984) are often cited, perhaps mistakenly, as the progenitors of the contemporary conception of mechanistic explanation; too often overlooked is Harré's (1960) paper, which expresses a more developed conception that is more consonant with the research program of New Mechanists than either of these two.

[2]   See, e.g., Bechtel (1995, 2008); Bogen (2005, 2008); Craver (2001); Craver & Bechtel (2006, 2007); Glennan (1996); Machamer (2004); Machamer et al. (2000); Piccinini (2007, 2008); Steel (2004); Tabery (2004); Thagard (1998, 2003); Wimsatt (1994); Woodward (2002, 2003); Wright (2007).

[3]   See, e.g., Bechtel (2006); Bechtel & Richardson (1993); Craver (2002, 2003); Craver & Darden (2001); Darden (2006).

philosophy of neuroscience, entitled *Explaining the Brain*. The book's central project is to articulate a framework for understanding explanation in neuroscience. To that end, Craver assimilates his mechanistic conception of explanation to Salmon's (1984) ontic conception, which results in the thesis that neuroscientific explanations are ontic exhibitions of mechanisms. Such theses are highly representative of the general New Mechanist program.

My contention will be twofold. Firstly, I'll argue that the development of a mechanistic conception of scientific explanation is severely hampered if those explanations are rendered as ontic. Mechanistic explanations—whether in systems biology, neuroscience, or elsewhere—are not ontic exhibitions of mechanisms. The audacity of this claim derives from its heretical opposition to the now-orthodox way of thinking about the mechanistic and ontic conceptions, which categorizes them as synergetic if not even equivalent. (I'll call this view the 'orthodox view'.) Secondly, I'll argue that mechanistic conceptions of scientific explanation are plausible only if rendered as epistemic. The vapidity of this claim derives from the epistemic conception being the only game in town. Together, these two claims entail collapsing or otherwise rejecting the traditional epistemic/ontic distinction.

1.2 An inconsistent triad

The contrast between my contention and the orthodox view can be made salient by considering the following three claims.

(1)   *The best conception of scientific explanation is a mechanistic one.*
(2)   *Mechanistic conceptions of scientific explanation are ontic.*
(3)   *Ontic conceptions of scientific explanation are fallacious.*

Prima facie, (1)-(3) appear to form an inconsistent triad. Firstly, it appears that mechanistic conceptions can both serve as our best conception of scientific explanation and be rendered as ontic; but it would then seem false that ontic conceptions could be fallacious. Secondly, it appears that ontic conceptions can be fallacious and that the best conception can be a mechanistic one; but it would then seem false that mechanistic conceptions can be rendered as ontic. Thirdly, it appears that mechanistic conceptions can be rendered as ontic and ontic conceptions can be fallacious; but it would then seem false that the best conception could be a mechanistic one.

Various strategies are available for managing such inconsistencies. The most facile resolution would be to concur that at least one claim is false. And indeed, the consensus among New Mechanists is that only the first of these three aforementioned possibilities is endorsable: i.e., (1) and (2) are true, but (3) is false. One reason why is historical. The elaboration of both mechanistic and ontic conceptions of scientific explanation occurred primarily within the so-called 'Chicago School' of philosophy of science, following its migration to Pittsburgh, UC San Diego, Washington University in St. Louis, and elsewhere. A second reason, already alluded to, is more conceptual. Influenced by Salmon, New Mechanists have pursued an integration of these two conceptions, and—in a few cases—a

unification such that the classes of explanations they pick out more-or-less harbor all and only the same versions.[4]

Another strategy would be to deny that any of (1)-(3) are false, and instead offer them as the premises of a valid and sound argument whose conclusion is:

∴ (4)   *The best conceptions of scientific explanation are not very good.*

Whether or not this second strategy is unduly pessimistic, it is at least unpromising for anyone with a serious commitment to any one conception over any other. Rather than pursue an actual argument against it, or mention other strategies, let me instead just insist, with virtually all New Mechanists, that any such argument would be unsound even if valid. The twist offered here, however, is that New Mechanists have been mistaken about which of (1)-(3) is the false claim. It is (2), not (3). Mechanistic conceptions are the best conception of scientific explanation currently on offer. But the ontic conception is indeed fallacious, and has been mistaken for a viable competitor.

1.2 Agenda

The primary aim of this paper, then, is to disabuse us of the ontic conception of scientific explanation, and suggest that mechanistic explanation <u>must</u> be rendered epistemically. To that end, I briefly canvass Salmon's epistemic/ontic distinction in §2, and argue that Salmon failed to properly characterize the epistemic conception of scientific explanation. Resultingly, the motivation to abandon it for the ontic conception is based on a straw person argument. In §3, I characterize the orthodox view, focusing particularly on an argument from Craver sanctioning it. The argument involves a bifurcation of explanations into objective and subjective, and the bifurcation is in turn justified by the claim that the term *explanation* is linguistically ambiguous. In §4, I argue that the reasons for attributing ambiguity are few and weak; in particular, while *explanation* does pass one zeugmatic test—namely, syllepsis—it does so only by occurring in formulations that surreptitiously depend on the deployment of ellipsis. Once one recognizes that elliptical techniques are at work, it becomes plain that the only charitable and non-absurd interpretation of the orthodox view is a trivializing one, whereby explanation turns out out to be just what advocates of the epistemic conception say it is. In §5, I argue that the expression of the ontic conception bottoms out in what Craver calls 'filler terms'; explicating *exhibition*-talk non-elliptically fails to perform the labor demanded of it, and so fails to avoid being rendered epistemically. Finally, I conclude with a positive characterization of mechanistic explanation to replace what is offered by the orthodox view.

From this agenda, it's hopefully clear that my criticisms of the ontic conception, which generally take the form of a reminder about what explanations are, should be interpreted as a self-critical conceptual <u>advance</u> for the New Mechanists, and an emendation for Craver's project in particular. After all, it hardly follows from the prescription that New Mechanists ought to relinquish the ontic conception, and a fortiori the

───────────────

[4]   See, e.g., Craver (2007: 21, 39); Delehanty (2005); Tabery (2004); cf. Machamer et al. (2000: 7).

orthodox view, that they need abscond from any other aspect of their research program.

2. An inheritance

2.1 Salmon's epistemic/ontic distinction

Throughout his career, Salmon consistently delineated 'two grand traditions of scientific explanation' (1984: 15-20, 84 ff.; 1998: 69, 126, 320). He named the first tradition, characterized by its focus on logic, the *epistemic conception*, and named the second tradition, characterized by its focus on causality, the *ontic conception*.[5, 6] Arguably, Salmon's best summarization of the distinction between these two conceptions contrasted his own claim that to explain a phenomenon is to show how it fits into a causal nexus (1984: 19) with Hempel's (1965: 488) claim that it is to show that it fits into a nomic nexus.[7]

Salmon divided the epistemic conception intro three versions, the most renown of which was his so-called 'inferential version', whereupon explanation involves making inferences—prototypically deductive—in first-order logic from explanans-statements to an explanandum-statement. Virtually all philosophers of science have followed Salmon in not only treating the epistemic conception to be little more than his inferential version, but also in treating his inferential version to just be the covering law model tout court (given the further constraint that one explanans-statement be a law of nature).

According to that model, there are two main features of scientific explanations: (i) logical necessity and (ii) nomic expectability (Salmon 1984: 16, 84 ff.). Firstly, an explanation E for some phenomenon ϕ is an argument with at least one premise stating a law and at least one stating initial boundary conditions, such that the truth of those premises logically necessitate the truth of conclusion codifying ϕ. Secondly, singular events or instances of laws are therefore subsumed or "covered" under more general laws, such that one should rationally be able to expect the

---

[5]   For the epistemic conception as Salmon conceived it, see Hempel (1965); for the ontic conception, see Salmon (1984) and Scriven (1975). Although it's generally been sidelined, Salmon also distinguished these two conceptions from the 'modal conception'; for a potential bridge between the ontic and modal conceptions, see Woodward (2002, 2003). For criticism of his historical scholarship, see Fetzer's (1992) argument that Salmon's 'history of scientific explanation' was improperly insulated from his personal autobiography.

[6]   Salmon occasionally framed the focus of the ontic conception in terms of mechanisms instead of causality (e.g., 1984: 132, 239; 1998: 324). As is well known, however, Salmon never defined *mechanism*; nor did he ever characterize the concept MECHANISM (see Glennan, 2002). The likely reason why is that he had none in mind.

[7]   In subdisciplines such as biochemistry, behavioral neuroscience, neuroeconomics, social psychology, psychopharmacology, etc., there are far too few nexuses that can be nomically specified, and so, far too few fittings of phenomena into them, and so, far too few explanations pace Hempel. But since those disciplines are rife with putative explanations, Hempel's claim presumably mischaracterizes the nature of explanations in those disciplines.

explanandum to be the case given her knowledge of the general laws and initial conditions.[8]

For Salmon and other advocates of the ontic conception, scientific explanations are not arguments. In particular, they are not inferences from sets of statements to statements, much less any other kind of representations. Rather, scientific explanations are exhibitions of pattern-fitting. 'Scientific explanation, according to the ontic conception, consists in exhibiting the phenomena-to-be-explained as occupying their places in the patterns and regularities which structure the world', as Salmon put it (1984: 239), and again, 'to explain an event is to exhibit it as occupying its place in the discernible patterns of the world' (1984: 18, 121).

As is well known, what makes the ontic conception gripping is that the patterns within which explananda are exhibited as fitting into are prototypically causal: 'Causal relations lie at the foundations of these patterns and regularities; consequently, the ontic conception has been elaborated as a causal conception of scientific explanation' (1984: 239), and again, 'an explanation of an event involves exhibiting that event as it is embedded in its causal network and/or displaying its internal causal structure' (1998: 325; see also 1984: 19). Indeed, at times, Salmon seemed to suggest that the concept CAUSE is the non-negotiable core of the ontic conception: 'It is only by providing adequate treatments of causal concepts, I believe, that we can successfully implement the ontic […] conception of scientific explanation' (1984: 122).

Given these 'two grand traditions', it is unsurprising that New Mechanists have framed their view of explanation in terms of Salmon's ontic conception. The perceived conceptual proximity of the ontic and mechanistic conceptions owes to their common focus on causality, rather than logic. For the sake of a slogan, their common focus might be expressed by misappropriating that of the early 20th century phenomenologists': *back to the* (*causes of*) *things themselves!*

2.2 The epistemic conception of scientific explanation

The above slogan, misappropriation aside, belies a deep misrepresentation of what debates over scientific explanation are about. No conception of scientific explanation is so rude as to posit a mutually exclusive focus on either causality or logic. Even Hempel's favorite case of D-N explanation, for instance—i.e., Semmelweis's explanation that puerperal sepsis was being caused by the introduction of decomposing cadaverous tissue into patients' bloodstream via open surgical wounds—was one in which causation figured centrally. For Hempel, statements about the singular causes of particular events presuppose general laws of nature, and these presuppositions are what "drive" scientific explanations (1966: 53). However, one need not agree with Hempel in the least to recognize both (i) the pseudo-dichotomy between the logical versus causal structure of explanations, and—more importantly—(ii) the implausibility of Salmon's reduction of the epistemic conception to the covering law model.

---

[8]   By these two features, the covering law model subsumes several others: the Deductive Nomological (D-N), Deductive Statistical (D-S), Inductive Statistical (I-S), Deductive Nomological Probabilistic (D-N-P) models, and, to a lesser extent, the Statistical Relevance (S-R) and Expected Information (E-I) models. Yet, each version is slung with well-known, and in some cases fatal, problems; New Mechanists are therefore right to look askance at the covering law model.

The epistemic conception is the view that a scientific explanation E is a finite class of operations on representations. Although maximally schematic, this characterization is neither trivial nor universally accepted (advocates of the modal and ontic conceptions, for instance, reject it). Its main utility resides in organizing its various versions that elaborate the specific nature of the operations and representations involved.

Prototypically, the operations are inferences and other "movements" of thought for wielding cognitive models—often of mechanistic activities—whose aim to enrich our knowledge of an explanandum ϕ.[9] Abounding is a large class of such prototypical versions beyond the covering law model, including one version positing processes of translation and mental manipulation of physically isomorphic scale models (Kraik, 1943), another focusing on deductive inferences on sets of formal possibility tableaux (Johnson-Laird, 1983), and yet another on exductions from image-specific intrinsic cognitive models (Waskan, 2006).[10] Beside these prototypical versions, recent work on mechanistic explanation in particular has focused on the cognitive operations involved in mechanical reasoning with diagrams, images, and other non-linguaformal representations—particularly, the inferential structure involved in visuospatial animation and simulation (Hegarty, 1992, 2004; Hegarty et al., 2003).[11]

These prototypical versions of the epistemic conception focus on operations and representations at the personal level. Worth mentioning are two that focus on the subpersonal level. For Churchland (1989), scientific explanation is a kind of reconfiguration of synaptic weights in neurocomputational state-space, in which the representations are vectors and the operations on them are vector-to-vector transformations that assimilate novel representations of phenomena to attractors. For Thagard (1998a,b, 2003), scientific explanation involves making reasonable and coherent non-deductive inferences, which take the form of solving constraint-satisfaction problems; solving such problems involves implementing heuristics at Marr's algorithmic level for computing a coherence function.

Two more popular versions of the epistemic conception are the erotetical and unificationist. Erotetic versions treat the representations to be essentially why-questions ($¿r_i$, $¿r_j$, $¿r_k$, …) and the operations on them to be essentially that of answering, such that E has the form $answer(¿r_i \wedge ¿r_i \wedge$

---

[9] Note that the sense of (epistemic) 'operations' in the context of explanation here is not that of (ontological) 'operations' in the context of mechanisms. The latter is a technical term designating what comprises mechanistic activities. Note also that a consequence of this view is that failing to aim at increasing knowledge of some ϕ_i precludes being a scientific explanation E.

[10] Although mental space theory has not yet been brought to bear on scientific explanation, it is potentially another such prototypical version. The representations are base, input, and blended mental spaces—'partial assemblies of cognitive potential'—constructed in online thought and speech via operations such as compression, mapping, identification, blending, integration, etc. (see, e.g., Fauconnier, 2007).

[11] The modifiers *mechanistic* and *mechanical* each profile hierarchical systems, but differ semantically according to construed degree of abstraction. Hence, *mechanical* is used for conceiving of more concrete—often "push/pull"—human-scale objects (e.g., as in so-called 'naïve physics); *mechanistic* modifies nouns and full nominals that profile a much wider range of systems.

$ιr_k$ …).[12] The unificationist version treats E to be an axiomatization of a set of accepted beliefs K (Kitcher, 1981), whereupon scientists apply sets of argument patterns (#$_1$, …, #$_n$) so as to derive a maximally unified explanatory store, $E(K)$. Also worth mentioning are information-theoretic versions, which characterize the representations to be essentially probabilistically defined units of information, and the operations on them to be algebraic.[13]

Each of these and sundry other versions share two features. Firstly, they're all versions of the epistemic conception, in which scientific explanation involves operations on representations. Secondly, they're all competitors to the covering law model, and, in some cases, to Salmon's characterization of the epistemic conception whole cloth. The project of further explicating and adjudicating between these versions is unnecessary and hardly something that could be undertaken here anyway; but all that is needed is the (eventual) recognition that one can remain silent on questions about the epistemic conception (e.g., what versions fare best, how to analyze the extension of 'operation', how to delimit representations, what the norms governing such explanations are, what the problems of each version are, etc.), and still acknowledge that the covering law model—while most infamous—merely stands as one (flawed) version among a litany of others.

Consequently, the epistemic conception is much more expansive than is usually acknowledged because what philosophers of science have acknowledged is Salmon's characterization. Salmon, however, artificially reduced the space of possible versions of the epistemic conception in order to motivate the ontic conception. Of course, demonstrating that his characterization was a straw person is not yet to demonstrate that Salmon's ontic conception is also fallacious, and so insufficient to fully demonstrate that the distinction should be collapsed or otherwise rejected. In the rest of the paper, I address myself to arguing thus.

3. The orthodox view

The orthodox view is the view that assimilates the mechanistic conception to the ontic conception. On its further ontic rendering, this mechanistic conception omits all non-ontic entities—differential equations, diagrams, models, names, reasons, truth and reference, etc. Explanations are exhibitions rather than arguments, and what is exhibited are the mechanisms that explain (alt., 'the fittings of mechanisms within causal patterns' that explain). Such explanations are, inter alia, physical/causal relationships between explanantia and explananda (Salmon, 1984), such that an explanans just is a mechanism and an explanandum just is the

---

[12]   Scientific explanations frequently take the form of answers to *how*-question or *how-possibly*-questions, in which case erotetic analyses are not necessary. And they are insufficient to the extent that scientific explanations require understanding what produces the explanandum and how. As Railton famously noted, 'explanations […] may legitimately be regarded as unsatisfactory unless we can back them up with an account of the mechanism(s) at work' (1978: 208).

[13]   Greeno (see also Fetzer, 1992) is sometimes credited with the canonical formulation, in which information transmission is a measure of explanatory goodness; but, since goodness metrics are metrics of goodness and not accounts of explanation per se, and since the information-theoretic version is not the central one, I won't pursue it here.

phenomenon causally produced by its activities.[14] Its component parts, operations, and organization are themselves what do the explaining, and they do so in virtue of simply <u>being</u> the explanans.

Exemplifying the orthodox view are the indefinitely many remarks that the mechanisms responsible for a given phenomenon explain it. For instance, consider Salmon's statement that 'to understand the world and what goes on in it, we must expose its inner workings. To the extent that causal mechanisms operate, they explain how the world works' (1984: 133). Similar exemplifications are particularly ubiquitous among New Mechanists. For instance, consider Bechtel's claim that 'the behavior of a system at a given level of organization in nature is explained by processes at lower levels and their interactions in the confines of that system' (1995: 163).[15] Or consider Machamer et al.'s claim that 'it is not the regularities that explain but the activities that sustain the regularities' (2000: 22), which is echoed by Delehanty, who wrote: '[…] it is clear what [mechanisms] are supposed to do for us: they provide causal explanations of the phenomenon of interest' (20005: 721). Other exemplifications are focused on specific cases: e.g., in discussing the mechanism of neurotransmitter release, Craver avers, 'the explanandum is the release of one or more quanta of neurotransmitters in the synaptic cleft. The explanans is the mechanism linking the influx of Ca^{2+} into the axon terminal' (2007: 22). Suffice it to say, then, 'orthodox view' quite clearly names the orthodox view.

Although the New Mechanist literature is rife with such exemplifications, self-consciously explicit and extended defenses of the orthodox view are less apparent. Craver's is both the richest and most recent, and so is worth quoting in full:

> [Salmon] defended an ontic view, according to which explanations are objective features of the world. This idea can be brought out by considering an ambiguity in the term, *explanation*. Sometimes explanations are *texts*—descriptions, models, or representations of any sort that are used to convey information from one person to another. [...] Other times, the term *explanation* refers to an objective portion of the causal structure of the world, to the set of factors that bring about or sustain a phenomenon (call them *objective explanations*). What explains the accident? The ice on the road, the whiskey, the argument, the tears, and the severed brake cables. There are mechanisms (the objective explanations) and there are their descriptions (explanatory texts). Objective explanations are not texts; they are full-bodied things. They are facts, not representations. They are the kinds of things that are discovered and described. There is no question of objective explanations being 'right' or 'wrong', or 'good' or 'bad'. They just are. Objective explanations, the causes and mechanisms in the world, are the correct starting point in thinking about the criteria for evaluating explanatory texts in neuroscience. […] Good mechanistic explanatory texts (including prototypes) are good in part because they correctly represent objective explanations. Complete explanatory texts are complete because they represent all and only the relevant portions of the causal structure of the world. Explanatory texts can be accurate enough and complete enough, depending on the pragmatic context in

---

[14]   This canonical characterization of the relationship between a mechanism and a phenomenon ϕ to be explained should be eshewed. In my view, when all is settled, a (good) mechanistic explanation will identify ϕ with the activity of an actual mechanism, which is just to say that the relationship between that activity and ϕ is identity (much as the relationship between a mechanism and its constituents is constitution). Consequently, that relationship is neither realization, implementation, grounding, 'underlying', 'causally producing', or 'being responsible for' or any number of other metaphysical relations—contra virtually all New Mechanists (especially those influenced by Machamer et al., 2000 (e.g., Bechtel & Abrahamsen, 2005; Craver, 2001; Darden, 2006; Delehanty, 2005; Glennan, 2002; Piccinini, 2007; Tabery, 2004); cf. Craver & Bechtel, 2007).

[15]   Bechtel (2005, 2008: 18) has recently endorsed a non-ontic, epistemic view of mechanistic explanation.

which the explanation is requested and given. Objective explanations are not variable in this way. (2007: 27)

In this passage, Craver begins with the unobjectionable distinction between mechanistic representations and represented mechanisms, and then claims that the term *explanation* can and does ambiguously refer to both, mutatis mutandis. These claims are then pressed into service as the basis for an argument sanctioning the orthodox view:

(5)   *Terms for explanation are ambiguous.*
(6)   *Explanations can be bifurcated into explanatory texts and objective explanations.*
(7)   *Mechanistic explanations are a species of objective explanations.*
∴ (8)   *The orthodox view is warranted.*

Of course, Craver and I both concur that mechanistic explanatory texts (subjective explanations?) aim to increase knowledge about mechanisms, and that the successful ones are those whose representations are operated on in ways that track their particulars—i.e., that 'accurately characterize the causal structure of the world', as he puts it (2007: 27 fn. 3). Salmon would have surely agreed, having himself supposed that scientific explanation more generally aims at understanding, and understanding results from knowing how things work (1984: 240). Obviously, 'knowledge of how things work' is an epistemic matter if anything is, which is just to say that analysis of explanatory texts properly requires an epistemic conception. Doubtlessly, analysis of the nature of those workings goes hand-in-hand with constructing causal and mechanistic explanations and is crucial for assessing them; but that's an analysis of causes and mechanisms, not explanation. According to the orthodox view, this last distinction doesn't follow, as some mechanistic explanations are ontic exhibitions of mechanisms.

4.   Why the ontic conception is fallacious

4.1 An ambiguity?

Premise (5) is crucial for Craver's argument. But why think it's true? Presumably, constructions involving *explanation* pass linguistic tests for syntactic ambiguity, but their doing so is utterly irrelevant to the issue at hand since Craver's appeal was to semantic ambiguity. And when we turn to semantic ambiguity, we find that the reasons for attributing it to the term *explanation* are few and weak. For instance, it's not the case that *explanation* has a univocal phonological structure, but different orthographic and semantic structures, as homophones do (e.g., *basal/basil*, *break/brake*, *seller/cellar*); nor is it the case that it has a single orthographic form but different phonological and semantic poles, as homographs do (e.g., *desert/desert*, *live/live*, *bass/bass*); nor is it the case that *explanation* is homonymous, having the same orthographic and phonetic structure but constituting fully separable and unrelated lexical items (e.g., *bark/bark*, *duck/duck*). So, why think that premise (5), upon which the argument turns, is true?

Worse, the term *explanation* fails other batteries of linguistic tests for semantic ambiguity. For instance, terms like, e.g., *ring* (circular mark, circular piece of jewelry, group of people operating together clandestinely,

boxing arena, etc.) or *coffee* (brownish hue in color-space, social gathering, flavor of ice cream, type of bean, hot liquid beverage, etc.) have multiple but related lexical senses, and are ambiguous in this way. The ambiguity of these polysemous terms can be brought out in tests involving mixed inferences. For example, the following syllogism,

(9)   *Her bike wheels are now completely true.*
(10)   *What is now completely true cannot be false.*
∴ (11)   **Her bike wheels cannot be false.*

involves an equivocation of its major term, wherein the lexical sense of *true* in (9) has the meaning of being corrected so as to spin along a vertical plane, but has the meaning of being maximally accurate in (10). However, unlike these and myriad other terms, *explanation* does not have polysemous senses of this sort, and so does not pass these sorts of mixed inference tests. Confirmation can be found in most dictionaries.

Another battery of tests involve various types of zeugma, in which concurrent constructions are governed by a single term (usually classed grammatically as a noun or verb). Of these tests, the most important for our purposes is syllepsis, which is exemplified by cases in which the constructions are semantically orthogonal:

(12a)   **The mistake and the asylum patient were committed.*
(b)   **The mass firing happened at both neuronal and professional levels.*

In (12a), there are two incongruent trajectors of the finite clause, suggesting that two different senses of the verb *committed* are being invoked; in (12b), the nominal subject *firing* simultaneously pertains to activation patterns of neurons and termination of employment. Interestingly however, *explain* and *explanation* do pass tests of syllepsis:

(c)   **She explained why the ValuJet crash in the Florida Everglades occurred, and so did the oxygen canisters.*

In (12c), the verb profiles a particular type of process with the same landmark, but apparently incongruent trajectors: the agent designated by the pronoun *she* and the entities designated by the count noun *canisters*. And yet, doesn't passing just one test for semantic ambiguity establish the bifurcation into subjective and objective explanations proposed by Craver? No. *Explanation* passes tests for syllepsis only in virtue of occurring in constructions whose formulation surreptitiously depends on the deployment of elliptical techniques. And once we understand the role of ellipsis in formulating such constructions—discussed in the following subsection—it becomes evident that the putative "ambiguity" of *explanation* is easily explained away. (To put the point somewhat paradoxically, passing this particular test of syllepsis is what allows us to recognize why *explanation* is unambiguous.)

4.2 Syllepsis via ellipsis

The archetypal explainer is a person. And so to say that mechanisms explain, as the New Mechanist orthodoxy does, is to speak metaphorically. It is apiece with claiming

(13a)   *My car refuses to start.*
(b)   *Her mesocorticolimbic dopamine system will eventually pay a heavy price for the transition from drug use to abuse.*

The metaphors in (13) are metaphors involving personification. Various further analyses of each are available,[16] but regardless, the point to fixate on is that mechanisms give (alt., 'are'?) explanations in precisely the same sense that the cars have poor etiquette, are moody or susceptible to learned helplessness, that DA_1 systems beg, borrow, and steal, and so forth. Relatedly, the archetypal products produced and exhibited by an explainer are representations. Typically, these representations are reasons; other possibilities include the cognitive models that explainers construct and control, the sentences they utter or figures they draw, the journal articles they publish, and so forth. And so to say that mechanisms are exhibited by explanations, as the New Mechanist orthodoxy does, is to speak <u>metonymically</u>. It is apiece with saying that

(14a)   *Radiology isn't answering the phone.*
(b)   *That quote from Hesiod is perfect for your purposes.*

The metonymy of *Radiology* in (14a) is based on a whole-to-part shift in its profile, such that the whole department or unit metonymically is conceptualized as standing for the individuals who comprise it. The metonymy in (14b) is equally clear; when one quotes Hesiod, what one quotes (alt., 'exhibits') is not a token archaic Greek vocalization of the deceased poet named by the singular term *Hesiod*, but a portion of certain representations associated with him, e.g., (translated/transliterated) linguistic expressions from a certain edition of *Works and Days*. The entity is referred to, metonymically, as a convenient stand-in for what is actually being conceptualized—namely, the text itself.

The ontic conception fundamentally relies on these elliptical techniques, and the analysis of *explanation* would not pass the test for syllepsis without them; moreover, the implication of ellipsis is what explains away the term's putative ambiguity. For instance, no one would interpret Thagard, who wrote that 'biochemical pathways explain by showing how changes within a cell take place as the result of the chemical activities of the molecules that constitute the cell' (2003: 238), as having claimed—literally—that biochemical pathways explain or show, much less that those pathways step out for lunch, pass the sodium chloride, stop to ask for directions, or sometimes go on holiday.

Of course, competent hearers easily navigate the metonymical statements and metaphors of advocates of the ontic conception, and they do so either by explicitly applying principles of charity, or by implicitly recognizing that elliptical techniques are being employed, or both. That competent hearers do, and should, interpret them as speaking elliptically is normal and unsurprising, as the techniques for doing so are frequently relied upon to facilitate expository or communicative purposes. Indeed, what would be surprising would be any interpretation of advocates of the ontic conception and orthodox view as having spoken either literally or non-elliptically, which is just to say that any problems are caused by taking those views seriously.

The fallacy embodied in the ontic conception should now be in full view. Biochemical pathways, changes in Ca^{2+} concentrations, phasic DA_1

_____

[16]   For instance, the conceptual source-to-target target mappings are constructed by blending the frames for personal activities (e.g., interlocution, negotiation and bargaining, economic transaction) that structure the generic and input spaces with other input spaces for non-personal systems (see, e.g., Fauconnier, 2007).

bursting or allostatic mesocorticolimbic dysfunction, oxygen generators, etc. are simply inapposite candidates for doing any explaining. Component parts and component operations constitute mechanisms, not explanations; and they do not—cannot—constitute a possible explanans, on pain of absurdity. Rather, the relevant parts, operations, and organization of a mechanism must always and everywhere, inter alia, be codified into structural or functional <u>representations</u> of some sort. And this is, in part, just what advocates of the epistemic conception claim. Part and parcel of articulating the epistemic conception is therefore giving just this reminder when advocates of the ontic conception forget themselves.

To better understand how passing the test for syllepsis is being explained away, reconsider the available scientific explanation about the historical event mentioned in (12c). Rendered epistemically, the explanation is a cohesive set of ratiocinative procedures applicable to models of the constituent sequence of processes, the entities participating in them, the causal structure connecting those processes and entities, and so forth.[17] Crucially, there is no alternative story for advocates of the ontic conception to tell; for there is no way to exhibit how the oxygen canisters fit within a causal pattern, and thus no way to construe the two conjoined sentences in (12c) literally or non-elliptically. The obvious reason why is that being extant a necessary condition on the very possibility of (ontically) exhibiting ϕ as part of the causal structure of the world, yet the canisters were incinerated upon impact more than a decade ago. Consequently, the ontic referent of the expression *oxygen canisters* itself could not possibly explain anything, contra Craver and other advocates, for there is no such referent.

Moreover, advocates of the ontic conception and orthodox view also face a dilemma. Consider constructions parallel to (12c), such as

(12d)   *The decreased expression of AMPA/kainate receptors caused the weak EPSPs, and so did the explanation.*

Presumably (12d) passes the test for linguistic ambiguity if, and only if, (12c) does. On one hand, if (12d) does pass, then it's a sylleptic construction precisely because the term *explanation* therein does not profile a Craverian objective explanation; but this effectively suggests that the verb *explain* can be informally but unambiguously characterized as a (deeply normative) narrative practice that cognizers engage in to understand the world; the non-cognizant world does not itself so engage. The plural mass noun *explanation* unambiguously refers to the products of that narrative practice—what I am calling explanations and what Craver calls 'explanatory texts'. On the other hand, if advocates of the ontic conception and orthodox view want to maintain their view, then they should claim that (12d) does not pass precisely because *explanation* is to be understood in the sense of Craverian objective explanations. But then, the motivation for claiming that (12c) is undermined; those advocates

---

[17]   Claims that the oxygen canisters (cf., 'the ice on the road, the whiskey, the argument, the tears, and the severed brake cables') explained the crash are easily paraphrased away: the matrix of representations evoked by the expression *oxygen canisters* is necessary for understanding the output of those ratiocinative procedures on those models. Subsequently, the onus of providing an intelligible story otherwise is actually on advocates of the ontic conception. Thanks to Gualtiero Piccinini for pointing this out to me. Note also that the relevant expression in (12c), *oxygen canister*, is itself metonymical—a whole standing in for a particular active zone (e.g., activated firing pins, leaky valve seals, certain quantities of oxygen, etc.).

would be committed to claiming that it's a nonsylleptic construction—i.e., the ontic verb *cause* isn't actually ambiguous—because the two entities profiled by the two clausal trajectors, i.e., the decrease in receptor population density and the explanation, respectively, enter into causal-mechanical relationships. This is, I submit, an extremely counterintuitive consequence of their view. Explanations don't cause, e.g., changes in Ca^{2+} concentrations, decreases in receptor population density, protein synthesis, etc.; nor are they entities or activities that comprise mechanisms, pace Craver and other New Mechanists. Worse, it simply begs the question of Craverian objective explanations at hand. This turning of the tables suggests that the orthodoxy is committed to an (unnecessarily) idiosyncratic conception of mechanistic explanation.

My deflationary reading of Craver's argument for the ontic conception and orthodox view has been that any interpretation of them, which takes their advocates at their word, is uncharitable; for such interpretations leaves those advocates committed to absurdities about what scientific explanations are. In contrast, the most charitable interpretation of the ontic conception is the one that reveals it as counterfeit; for its *explanation-as-exhibition*-talk is merely elliptical, and what it is elliptical for just is the epistemic conception.

4.3 Facticity

Part of the reason for Craver's bifurcation into subjective and objective explanations, and thus his defense of the ontic conception, is historical. He, like most other New Mechanists, have simply been cued by Salmon's terminology:

> The term *explanation* may be construed in either of two ways. It may be taken as the combination of the explanans and explanandum—the explanans having its place in the causal pattern or causal nexus in which it is objectively embedded. This term may also be used to refer to the combination of explanans-statements and explanandum-statement—the linguistic entity that is used to present the objective facts. (1984: 13, 274).

Given these two construals, Salmon continued, 'as I prefer to construe these terms, the explanandum is a fact—it is the fact-to-be-explained. […] To explain a fact is to furnish an explanans; these general and particular facts account for the explanandum. […] The explanans is a complex of objective facts. […]' (1984: 274).

Count nouns like *fact*, *faktum*, *facto*, etc. are genuinely linguistically ambiguous (Johnston, 2004) across a range of natural languages. Following the foreclosure of Middle English (late 15th-early 16th c.), *fact*-talk initially served to designate actions, occurrences, and events—i.e., Salmon and Craver's 'causal structure of the world'—but later underwent processes of grammaticalization and semantic bleaching to yield new, insubstantive senses in use with adverbial and gerundial clauses. These 'bleached' senses were evident by the early 20th century (recall Frege's equation of facts with true propositions in *Der Gedanke*), and were the principle fundament of the infamous 1950 Austin-Strawson-Cousin debate, in which Strawson leveraged the complaint that facts are merely 'psuedomaterial correlates' of true sentences. Justification for Strawson's complaint might appeal to the inability of *fact*-talk (e.g., *the fact that p*) to take predicates of occasion, duration, or spatiotemporal extension:

(13)   *The fact that norepinephrine is involved in flight-or-fight responses occurred in 2003.
(14a)   *The fact that there are resting potentials has lasted for less than 200 years.
(b)   *The fact that memory systems are localizable is still awaiting disconfirmation.
(15a)   *The fact that retinal cells exist is longer than the fact that pyramidal cells are at a lower level than hippocampi.
(b)   *The fact $A^2 + B^2 = C^2$ is older than $E = MC^2$

Unlike actions, mechanistic activities, events, or occurrences, facts seem not to possess temporal or spatial properties. That being the case, Salmon's preference for identifying extant mechanisms and their activities tout court with facts, the ambiguity in *fact*-talk notwithstanding, requires justification. That justification is not too hard to come by, and Salmon can hardly be blamed for having idiomatic preferences for treating facts in their substantive Austinian senses of singular events and actions rather than their insubstantive Fregean sense of true propositions.

However, it is quite another thing to induce confusion born of regimenting those preferences philosophically. Salmon imposed his preference for an Austinian conception of facts onto conceptions of scientific explanation, claiming that the objective relationship between the 'explanatory facts and the fact-to-be-explained' obtains in cases of explanation (e.g., *the onset of a storm objectively explains the falling barometric pressure*), and fails to obtain in cases of non-explanation (e.g., *hexing salt does not explain its dissolution*) (1984: 13). Minimally, the problem with doing so is the non-sequitur committed by overextending a case of genuine linguistic ambiguity (i.e., *fact*-talk) to a pseudocase (i.e., *explanation*-talk).[18]

On the view offered here, the extant in-/activities of mechanisms are just that—the mere existence of activities and inactivities. Explanations of them are not produced until an explainer contributes her or his epistemic and/or cognitive labor. Salmon's confusion has been parlayed by New Mechanists, wherein the adoption of the orthodox view mislocates <u>the space</u> of explanation. Instead of understanding explanation as taking place within the space of reasons, New Mechanists have put explanation out in the space of mechanisms. This is tantamount to a kind of category mistake.

Here, it might be helpful to distinguish Craverian objective explanations, which are 'portions of the causal structure of the world', from the objectivity of explanations, which represent those portions, potentially aright. Pace Craver, objective explanations are extra-representational. New Mechanists are, of course, right to focus on that causal structure, but wrong to identify causal structure with norm-governed ratiocinative practices such as explanation, or the products thereof. Salmon's claim that objective relationships hold between genuine explanans and explananda can therefore be interpreted as nothing more than that genuine explanatory relationships are objective. But objectivity—like truth or maximal accuracy—is primarily a property of representations. And the bifurcation of explanation into subjective and objective does not

---

[18]   Craver is not the only New Mechanist overextending ambiguity attributions from *fact*- to *explanation*-talk à la Salmon's claims that (Austinian) facts explain other (either Austinian or Fregean) facts: 'whenever $1 is inserted and the button marked *coke* is pressed, a coke appears in the slot at the bottom of the machine. The fact that the statement is true is explained by facts about how the Coke machine works' (Glennan, 2005: 446).

follow from the subjectivity or objectivity of explanations in the epistemic sense. Consequently, it appears that *exhibition-of-mechanism*-claims are merely claims that mechanistic explanations are operations on objective representations of mechanisms.

5. Filler Terms

Plainly, mechanisms are "integral" to mechanistic explanations, and so somehow "figure" in them. According to advocates of the epistemic conception, they do so only in the trivial sense that those explanations would be without representational content were there no mechanisms or their activities about which to explain. Presumably, this trivial and figurative sense is not what advocates of the orthodox view intended in staking out either the ontic conception or the orthodox view. Is there <u>any</u> other way to make sense of the *explanation-as-exhibition*-talk so central to the ontic conception?

Craver (2007: 113) rightly notes that one distinguishing aspect of mechanism sketches, unlike models of mechanisms, is their frequent devolutionary reliance on 'filler terms', which are used to gesture toward the organizational features, the activities and their component operations, and the component parts of mechanisms without detailing what is being designated. According to Craver, common examples in neuroscientific discourse include: *trigger*, *produce*, *encode*, *inhibit*, *mediate*, *represent*, *store*, *control*, etc. Yet, the very use of filler terms, however, is problematic: 'filler terms are barriers to progress when they veil failures of understanding. If the term *encode* is used to stand for some process-we-know-not-what, and if the provisional status of that term is forgotten, then one has only an illusion of understanding' (2007: 113).

Craver's point is here well-received, but does not go far enough; and we should apply it self-critically, no less to mechanistic explanation than to mechanisms. The claim that scientific explanations are exhibitions of mechanisms is, if explicated at all, explicated via appeals to certain cognates—*demonstrate*, *present*, *reveal*, *display*, *lay bare*, *show*, *indicate*, etc.—which are equally filler terms and equally often used in ways that can have just the pernicious consequences alluded to. Consider, e.g., Salmon's statement that, according to the orthodox view, scientific explanation 'involves laying bare the mechanisms […] that bring about the fact-to-be-explained' (1998: 328), or Craver & Bechtel's claim that, in explaining, scientists are 'presenting a mechanism as an explanation' (2006: 473). As such claims evince, it is usually no clearer what is being meant when these filler terms are invoked in lieu of, or to explicate, *exhibition*-talk. This is particularly the case since neither are they strictly synonymous with each other or with *exhibit*, nor are they or the metaphors they invoke themselves ever analyzed in any detail.

Moreover, all and any such explications are compatible with the epistemic conception. Consider the frequent attempts to explicate *exhibition* as showing, presentation, indication, etc., as in 'mechanistic explanation in neurobiology and molecular biology involves showing or demonstrating that the phenomenon to be explained is […] the result of productive activities' (Machamer et al., 2000: 22), or in Salmon's contrast that explaining a phenomenon is to show <u>how</u> it fits into a <u>causal</u> nexus

15

versus showing <u>that</u> it fits into a <u>nomic</u> nexus. Prototypically, such neurobiological showings involve producing representations of types (not tokens) of mechanisms from which inferences can be made about how the target ϕ can, ultimately, be identified with its activity.[19] For example, to explain what an action potential is and how it works, instructors often show an image presenting a type, namely, the action potential; to explain heart disease, one typically does not show or present a token heart and follow the evolution of its pathology over time; to explain what a posterior parietal motor command is and how it works, one abstracts away from token motor commands and invokes a representation of a type, namely, posterior parietal motor command; and so on. Hence, when Machamer et al. claim, in discussing the action potential, that 'it is through these activities of these entities that we understand how depolarization occurs' (2000: 13), how should we understand them? According to the epistemic conception, this is perfectly acceptable as a manner of speech, but unacceptable as a piece of serious philosophical analysis; for what our understanding literally proceeds 'through' are the representations of those entities and activities and the (epistemic) operations thereon.

That scientists explain by representing, inter alia, types of mechanisms rather than by showing or presenting tokens suggests that this explicatory maneuver would fail to perform the labor demanded of it, even if neither *exhibition* nor its cognates <u>were</u> filler terms (Wright & Bechtel, 2006). Firstly, presentations or mere showings of token mechanisms would inhibit the operations of abstraction and idealization that are often so crucial to making explanations effective as scientific currency. Secondly, there are obvious limitations (e.g., molar, human-scale, spectrally observable) on what sort of mechanisms can be and cannot be presented, or laid bare, or displayed. No social psychologist explains, or could explain, the mechanisms of group violence by presenting or merely showing them; no neuroscientist explains the phenomena of E-LTP or saltatory conduction by merely putting it on display. The overarching reason is that presentations, or displays, revealings, etc. of mechanisms, in their ontic, non-epistemic sense, are neither necessary nor sufficient for mechanistic explanation. For example, neuroscientists can mechanistically explain the neuronal basis of forgetting by simply telling a good story, or by writing a report; no actual mechanism need present itself to do so. And presentations of mechanisms are insufficient to do any heavy lifting in an explanation of how their activities are produced or organized, since presentation would only be relevant if it were made against the background of a large corpus of antecedently possessed conceptual knowledge and scientific understanding. For example, if the phenomena of vesicular dehiscence or GABAergic inhibition or kindling in epileptic seizures are cognitively abstruse to someone, would presenting those phenomena to them, even were it necessary for explanation, be sufficient? No, plainly. A great deal of para-presentational commentary would also need to be deployed; for explanation aims at understanding, as it is often said. (Of course, what that amounts to is a difficult but further question.)

Explicating *exhibition* as indication fails for similar reasons, as do combinations of presentation and indication. The plural mass nominal *indications* designates indicative acts: demonstrative pointings, deictic acts of referring, peri-linguistic acts of gesturing, etc. at a given mechanism's component parts and their operations and organization, which, together, constitute its activity. Such acts are shot through with representing and reasoning, as are attempts to understand what is being indicated or shown or demonstrated. After all, what one takes to be the explanans is entirely a matter of what ones finds cognitively salient. Of course, in the service of explanation, such acts can be highly complementary, but no set of such acts communicate explanations of complex phenomena in any ontic sense; no such demonstrative gesturing or pointing could unequivocally specify what it was that had been indicated, since myriad structure, function, and organization would be consistent with any such indication. Not only could one inquire, ad nauseum, about what had been indicated, but such inquiries might not be optional: the problem of indeterminancy of reference is hard enough with linguistic representations, but orders of magnitude more so with gesticulation, deictic gestures, demonstrative pointing, etc.

New Mechanists' legitimate point about the questionable use of filler terms to designate the component operations and parts of the activities of mechanisms should be self-critically applied to the the orthodox view of mechanistic explanations as exhibitions (alt., 'showings', 'fittings', 'demonstrations', 'presentations', etc.) of mechanisms. Again, the conditions in the world may provide the material or the ground of explanation, and our behaviors may be of help in communicating an explanation, but conditions in the world do not themselves explain, and no amount of or ontic exhibition or non-epistemic presentation will make that the case. And, a fortiori, no amount of complexification or simplification of presentation or indication is relevant given that nothing is explained by ontic exhibition. Explication of *explanation-as-exhibition-as-*[cognate] merely repackages the epistemic conception, and in a less perspicuous manner at that.

6. Conclusion

New Mechanists' theoretical platform has inherited a fundamental tension from Salmon. At once, they've desired an account of explanation that sloughs off the logical, cognitive, semantic, grammatical, and inferential concerns that misled logical empiricists, while simultaneously helping themselves to these and other aspects of the epistemic conception that they abandon. Ushering such a conception out the front door whilst sneaking it in the back has subsequently led to subtle and serious confusions about the nature of explanation, and has marred the virtues of an otherwise powerful and sophisticated framework in the philosophy of science (Wright & Bechtel, 2006).

I've argued for an deflationary interpretation of ontic conception of scientific explanation, as well as of the orthodox view that assimilates the mechanistic conception to it. Accordingly, the orthodox rendering of mechanistic explanation as ontic should be understood elliptically, i.e., merely just emblematic stand-ins for a richer (epistemic) account of how

to construe the relationship between explanatia and explananda. For example, many of the versions of the epistemic conception in §2.2 are amenable to the following working definition of mechanistic explanation.

> (*ME*) A mechanistic explanation is a class of ratiocinative procedures on a set of models $\{m_i, m_j, \ldots, m_n\}$ of a repetitive mechanistic activity, Σ ϕ-ing, that yields knowledge of how to identity the explanandum, ϕ, with the ϕ-ing of Σ.

As a working definition, some adjustments are required. For instance, implicit in (*ME*) is the thought that knowledge of how to identity the explanandum, ϕ, with the ϕ-ing of Σ, depends on how those models $\{m_i, m_j, \ldots, m_n\}$ represent the internal structure and function of Σ; assumed in (*ME*), then, is that those models will specify in detail the component parts $(\pi_1, \ldots, \pi_n)$ of Σ, the component operations $(\psi_1, \ldots, \psi_n)$ that are coordinated so as to compose the repetitive higher-level activity of ϕ-ing, and the internal dynamics and organization. Moreover, 'ratiocinative procedures' inclusively allows mechanistic explanations to implicate a variety of operations and types of reasoning (e.g., simulation, abduction, mental animation), and this too needs further specification. In any case, New Mechanists should be able to develop a more sophisticated conception of mechanistic explanation around (*ME*). Some of that work is already in place; Glennan (2002: S347), for instance, offers an explication of what a mechanistic model m_i is, and Woodward (2002: S375) offers a necessary condition on being such a model.[20] But the point to emphasize here is that mechanistic explanantia, in neuroscience and elsewhere, are (epistemic) operations on models $m_i, m_j, \ldots, m_n$ and other idealized representations of certain kinds of hierarchical systems, and the good ones systematically increase understanding of how those representations might ultimately be mapped onto representations of the target explanandum.[21]

Subsequently, a virtue of this proposal is that a central supposition of the New Mechanists remains untouched: scientific understanding is centrally focused on understanding what mechanisms do and how they do it. Another virtue is that common ways of analyzing explanations are accommodated. For instance, as Craver rightly notes, a consequence of the orthodox view is that certain predicates are inapplicable to mechanistic explanations; e.g., they are neither *good* nor *bad*, *accurate* nor *infelicitous*, *relevant*, *incomplete*, *substandard*, *clever*, *elegant*, *simple*, etc. If the epistemic conception is the only game in town, as I've argued, this inauspicious consequence doesn't occur. Similarly, claims that there are multiple competing objective explanations for any given target phenomenon, or that scientists sometimes discover which of several possible explanations is the best explanation—which are otherwise nonsensical on the orthodox view—are perfectly sensible (Waskan, 2006). The upshot of my proposal, however, is that it requires recognizing that

---

[20] One worry is that (*ME*) fails to make explicit the norms of good mechanistic explanations (Craver, personal communication). That is true: no norms are explicit in that definition. However, it hardly follows that they are inarticulable. Moreover, advocates of the epistemic conception can appeal to all the same norms appealed to by advocates of the ontic conception, pari passu. For instance, Craver (2007: 26) articulates five such constraints; in the parlance of the epistemic conception, good mechanistic explanations will satisfy each of these, and possibly several others.

[21] Note that I do _not_ claim that all neuroscientific explanations, much less all scientific explanations more generally, are mechanistic. Nor do I claim that non-mechanistic conceptions of explanation are implausible.

the ontic conception is threadbare as a philosophical analysis of scientific explanation.

In sum, New Mechanists have rightly rallied around Salmon's metaphor that 'the underlying causal mechanisms hold the key to our understanding the world' (1984: 260) but wrongly taken too seriously its metaphorical entailments. By collapsing or rejecting the epistemic/ontic distinction, and thus self-critically rethinking the orthodox view, a far more credible account of scientific explanation can be achieved. Some problems will remain, of course, for us to solve or dissolve as our general research program develops; but the sooner we recognize that there are no viable non-epistemic conceptions of scientific explanation, the sooner we can get on with showcasing our work on mechanistic explanation in all areas of scientific inquiry.

Acknowledgements

References

Bechtel, W.P. (1995). Biological and social constraints on cognitive processes: The need for dynamical interactions between levels of organization. *Canadian Journal of Philosophy* (*suppl*) *20*: 133-164.

Bechtel, W.P. & Abrahamsen, A. (2005). Explanation: A mechanistic alternative. *Studies in History and Philosophy of the Biological and Biomedical Sciences*, *36*: 421-441.

Bechtel, W.P. (2006). *Discovering Cell Mechanisms: The Creation of Modern Cell Biology*. Cambridge: Cambridge University Press.

Bechtel, W.P. (2008). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. New York: Routledge.

Bechtel, W.P. & Richardson, R. (1993). *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton: Princeton University Press.

Bechtel, W.P. & Wright, C.D. (2008). What is psychological explanation? In P. Calvo & J. Symons (eds.), *Routledge Companion to the Philosophy of Psychology*. New York: Routledge.

Bickle, J. (2003). *Philosophy and Neuroscience: A Ruthlessly Reductive Approach*. Dordrecht: Kluwer.

Bogen, J. (2005). Regularities and causality: Generalizations and causal explanations. *Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*: 397–420.

Bogen, J. (2008). Causally productive activities. *Studies in the History and Philosophy of Science*, *39*: 112–123.

Churchland, P.M. (1989). *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge, MA: MIT Press.

Craver, C.F. (2001). Role functions, mechanisms, and hierarchy. *Philosophy of Science*, *68*: 53-74.

Craver, C.F. (2002). Interlevel experiments and multilevel mechanisms in the neuroscience of memory. *Philosophy of Science*, *69*: S83-S97.

Craver, C.F. (2003). The making of a memory mechanism. *Journal of the History of Biology*, *36*: 153-195.

Craver, C.F. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.

Craver, C.F. & Bechtel, W.P. (2006). Mechanism. In S. Sarkar & J. Pfeifer (eds.), *Philosophy of Science: An Encyclopedia* (469-478). New York: Routledge.

Craver, C.F. & Bechtel, W.P. (2007). Top-down causation without top-down causes. *Biology and Philosophy*, *22*: 547-563.

Craver, C.F. & Darden, L. (2001). Discovering mechanisms in neurobiology: The case of spatial memory. In P.K. Machamer, R. Grush, & P. McLaughlin (eds.), *Theory and Method in the Neurosciences* (112-137). Pittsburgh: University of Pittsburgh Press.

Cummins. R. (2000). "How does it work?" versus "what are the laws?": Two conceptions of psychological explanation. In F. Keil & R. Wilson (Eds), *Explanation and Cognition* (117-144). Cambridge, MA: MIT Press.

Darden, L. (2006). *Reasoning in Biological Discoveries: Mechanisms, Interfield Relations, and Anomaly Resolution*. New York: Cambridge University Press.

Delehanty, M. (2005). Emergent properties and the context objection to reduction. *Biology and Philosophy*, *20*: 715-734.

Fauconnier, G. (2007). Mental spaces. In D. Geeraerts & H. Cuyckens (eds.), *The Handbook of Cognitive Linguistics* (351-376). Oxford: Oxford University Press.

Fetzer, J.H. (1992). What's wrong with Salmon's history: The third decade. *Philosophy of Science*, *59*: 246-262.

Glennan, S.S. (1996). Mechanisms and the nature of causation. *Erkenntnis*, *44*: 49-71.

Glennan, S.S. (2002). Rethinking mechanistic explanation. *Philosophy of Science*, *69*: S342-S353.

Glennan, S.S. (2005). Modeling mechanisms. *Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*: 443-464.

Harré, R. (1960). Metaphor, models, and mechanism. *Proceedings of the Aristotelian Society, 60*: 101-122.

Hegarty, M. (1992). Mental animation: Inferring motion from static diagrams of mechanical systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*: 1084-1102.

Hegarty, M. (2004). Mechanical reasoning as mental simulation. *Trends in Cognitive Science*, *8*: 280-285.

Hegarty, M., Kriz, S., & Cate, C. (2003). The roles of mental animations and external animations in understanding mechanical systems. *Cognition & Instruction*, *21*: 325-360.

Hempel, C.G. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.

Johnston, D.K. (2004). The natural history of fact. *Australasian Journal of Philosophy*, *82*: 275-291.

Kitcher, P.S. (1981). Explanatory unification. *Philosophy of Science*, *48*: 507-531.

Machamer, P. (2004). Activities and causation: The metaphysics and epistemology of mechanisms. *International Studies in the Philosophy of Science*, *187*: 27-39.

Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, *67*: 1-25.

Piccinini, G. (2007). Computing mechanisms. *Philosophy of Science*, *74*: 501-526.

Piccinini, G. (2008). Computers. *Pacific Philosophical Quarterly*, *89*: 32-73.

Railton, P. (1978). A deductive-nomological model of probabilistic explanation. *Philosophy of Science*, *45*: 206-226.

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

Salmon, W.C. (1998). *Causality and Explanation*. Oxford: Oxford University Press.

Scriven, M. (1975). Causation as explanation. *Noûs*, *9*: 3-16.

Steel, D. (2004). Social mechanisms and causal inference. *Philosophy of the Social Sciences*, 34: 55-78.

Tabery, J. (2004). Synthesizing activities and interactions in the concept of a mechanism. *Philosophy of Science*, *71*: 1-15.

Thagard, P. (1998a). Explaining disease: Causes, correlations, and mechanisms. *Minds and Machines*, *8*: 61-78.

Thagard, P. & Verbeurgt, K. (1998b). Coherence as constraint satisfaction. *Cognitive Science*, *22*: 1-24.

Thagard, P. (2003). Pathways to biomedical discovery. *Philosophy of Science*, *70*: 235-254.

Wimsatt, W.C. (1994). The ontology of complex systems: Levels, perspectives and causal thickets. *Canadian Journal of Philosophy* (*suppl*), *20*: 207-274.

Woodward (2002). What is a mechanism? A counterfactual account. *Philosophy of Science* (*suppl*) *69*: S366-S377.

Woodward, J.F. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

Waskan, J. (2006). *Models and Cognition: Prediction and Explanation in Everyday Life and in Science*. Cambridge, MA: MIT Press.

Wright, C.D. (2002). Reformulating mechanistic explanation. Unpublished manuscript.

Wright, C.D. (2007). Is psychological explanation going extinct? In H. Looren de Jong & M.K.D. Schouten (eds.), *The Matter of the Mind: Philosophical Essays on Psychology, Neuroscience and Reduction* (249-274). Oxford: Blackwell.

Wright, C.D. & Bechtel, W.P. (2006). Mechanisms and psychological explanation. In P. Thagard (ed.), *Handbook of Philosophy of Psychology and Cognitive Science* (39-79). New York: Elsevier.