

Matej Kohár, Beate Krickel

Compare and Contrast: How to Assess the Completeness of Mechanistic Explanation

Open Access via institutional repository of Technische Universität Berlin

Document type

Book chapter | Accepted version

(i. e. final author-created version that incorporates referee comments and is the version accepted for publication; also known as: Author's Accepted Manuscript (AAM), Final Draft, Postprint)

This version is available at

<https://doi.org/10.14279/depositonce-15257>

Citation details

Kohár, M.; Krickel, B. (2021) Compare and Contrast: How to Assess the Completeness of Mechanistic Explanation. In: Calzavarini, F.; Viola, M. (eds) Neural Mechanisms. (Studies in Brain and Mind, vol 17). Springer, Cham. pp. 395–424. https://doi.org/10.1007/978-3-030-54092-0_17.

Terms of use

This work is protected by copyright and/or related rights. You are free to use this work in any way permitted by the copyright and related rights legislation that applies to your usage. For other uses, you must obtain permission from the rights-holder(s).

Compare and Contrast: How to Assess the Completeness of Mechanistic Explanation

Matej Kohár and Beate Krickel

Institut für Philosophie, Literatur-, Wissenschafts- und Technikgeschichte,
Technische Universität Berlin, Berlin, Germany
Email: matej.kohar@gmail.com; beate.krickel@tu-berlin.de
Corresponding author: B. Krickel

Abstract Opponents of the new mechanistic account of scientific explanation argue that the new mechanists are committed to a ‘More Details Are Better’ claim: adding details about the mechanism always improves an explanation. Due to this commitment, the mechanistic account cannot be descriptively adequate as actual scientific explanations usually leave out details about the mechanism. In reply to this objection, defenders of the new mechanistic account have highlighted that only adding *relevant* mechanistic details improves an explanation and that relevance is to be determined relative to the phenomenon-to-be-explained. Craver and Kaplan (B J Philos Sci 71:287–319, 2020) provide a thorough reply along these lines specifying that the phenomena at issue are *contrasts*. In this paper, we will discuss Craver and Kaplan’s reply. We will argue that it needs to be modified in order to avoid three problems, i.e., what we will call the *Odd Ontology Problem*, the *Multiplication of Mechanisms Problem*, and the *Ontic Completeness Problem*. However, even this modification is confronted with two challenges: First, it remains unclear how explanatory relevance is to be determined for contrastive explananda within the mechanistic framework. Second, it remains to be shown as to how the new mechanistic account can avoid what we will call the ‘Vertical More Details are Better’ objection. We will provide answers to both challenges.

17.1 Introduction

It is widely agreed among the new mechanists that complete explanations are better than incomplete explanations, and that the closer an explanation is to being complete the better. However, there is an on-going discussion about what completeness of explanation amounts to (Baetu 2015; Miłkowski 2016; Craver and Kaplan 2020). Opponents of the mechanistic account of explanation object that the new mechanists

are committed to assuming that adding *any kind of details* about a mechanism will improve an explanation (Batterman and Rice 2014; Chirimuuta 2014; Levy 2014). As a consequence, the new mechanists are committed to the claim that, for example, mentioning quarks in the explanation of spatial memory will improve the explanation of the latter, listing all kinds of activities of all billions of neurons in the brain will improve explanations in neuroscience, and mentioning the exact location of the ion-channels or mentioning the exact diameter of the axon will improve the explanation of the action potential. This, according to the opponents, is obviously problematic, as actual scientific explanations do usually not provide all kinds of details. Real explanations are sketchy, they abstract away from details, and they do not necessarily mention quarks. And these explanations are good especially *because* they leave out details. Hence, the new mechanistic account fails as an adequate account of scientific explanation or is incomplete at best. This is what Craver and Kaplan label the '*More Details are Better*' (*MDB*) objection (Craver and Kaplan 2020).

Defenders of the new mechanistic approach, however, argue that they are not committed to such a 'More Details are Better' claim. They highlight that clearly only *relevant* details improve an explanation and that relevance is to be determined *relative to the phenomenon-to-be-explained* (Baetu 2015; Boone and Piccinini 2016; Miłkowski 2016; Craver and Kaplan 2020). Many authors focus on defending the view that explanations that abstract away from details can still be mechanistic (Boone and Piccinini 2016; Miłkowski 2016), or on how to empirically establish whether a given explanation is complete (Baetu 2015). In contrast, in a recent paper, Craver and Kaplan provide a detailed analysis of how the norm of completeness is to be understood in the context of the new mechanistic approach by elaborating on the *relevance relative to the phenomenon*-idea. In a nutshell, they argue that mechanistic explanation (or models) aim at explaining *contrasts*, such as the spiking of the action potential at -70 mV rather than -50 mV. Relevance has to be determined relative to these contrasts.

In this paper, we discuss Craver and Kaplan's (2020) reply to the MDB-objection. More specifically, the paper will proceed as follows: In Sect. 17.2, we present the MDB-objection and Craver and Kaplan's reply. In Sect. 17.3, we will highlight three problems for Craver and Kaplan's account that we will call the *Odd Ontology Problem*, the *Multiplication of Mechanisms Problem*, and the *Ontic Completeness Problem*. In Sect. 17.4, we will suggest modifications to Craver and Kaplan's reply that solve these problems. We will, in Sect. 17.4.1, introduce a distinction between ontic mechanisms, mechanism descriptions, and mechanistic explanatory tests that helps to solve the *Odd Ontology Problem* and the *Multiplication of Mechanisms Problem*. In Sect. 17.4.2, we will show that completeness is a predicate of mechanism descriptions and mechanistic explanatory texts rather than ontic mechanisms. We thereby solve the *Ontic Completeness Problem*. In Sect. 17.5, we will argue that even based on these modifications, the reply to the MDB-objection is confronted with two challenges: First, it remains unclear how explanatory relevance can be determined for contrastive explananda within the mechanistic framework (Sect. 17.5.1). Second, it remains to be shown how the new mechanistic account

can avoid what we will call the ‘Vertical More Details are Better’ objection (Sect. 17.5.2). We will provide answers to both challenges that essentially hinge on the idea that mechanistic explanations aim at identifying *crucial points of intervention*.

17.2 The MDB-Objection & Craver and Kaplan’s Reply

According to opponents of the new mechanistic approach, the mechanistic approach fails as it is unable to account for the fact that many explanations are good because they leave out details about the mechanism that is responsible for the phenomenon (Woodward 2013; Batterman and Rice 2014; Chirimuuta 2014; Levy 2014; Rice 2015). They hold that the mechanistic account entails or that mechanists even explicitly defend the claim that the more details an explanation mentions, the better this explanation is. This is the ‘More Details Are Better’ or ‘MDB’ objection (Craver and Kaplan 2020). More specifically, different authors seem to accuse the new mechanists of being committed to the following claims:

- a) Explanations should be complete.
- b) Explanations are complete if and only if they describe every detail of a mechanism.
- c) Adding details about a mechanism to an explanation always improves the explanation.

However, so the argument goes, often explanations are good because *they leave out* certain details of the mechanism (Woodward 2013; Batterman and Rice 2014; Chirimuuta 2014; Levy 2014; Rice 2015). For example, when cognitive neuroscientists explain cognitive capacities, they do not mention all 86 billion neurons of the brain and what they are doing (Chirimuuta 2014, 149). Network explanations are good because they mention the generic features of the network topology rather than more specific details about the implementing mechanism (Woodward 2013, 41). Optimality models provide explanations because they do *not* accurately represent the causes of a phenomenon, and they would provide “a worse (or perhaps no) explanation” at all if they described the causes in all detail (Rice 2015, 591). Completeness of mechanistic explanations, thus, must be interpreted in a way that is less demanding than b), or a) has to be rejected. In any case, c) has to be rejected as at least sometimes adding details about the mechanism does *not* improve the explanation.

Craver and Kaplan (2020) provide a detailed reply to this objection. They argue that their opponents are wrong in assuming that the new mechanists are committed to claims a), b), and c). According to Craver and Kaplan, the new mechanists are rather committed to what we summarize as claims a’), b’), c’). These claims are indeed compatible with the fact that some explanations are good even if they leave out details about the mechanism, or even if they are not complete.

- a') Completeness concerns stores of *explanatory knowledge* rather than models or explanations (Craver and Kaplan 2020, 310). The closer a store of explanatory knowledge is to being complete, the better. Although a complete model is in principle possible, actual models serve further non-explanatory purposes that such a complete model would fail to account for (such as understandability, computational tractability, unification) (Craver and Kaplan 2020, 308).
- b') A mechanistic explanation/model is complete if and only if it mentions all details of a mechanism that are *relevant for explaining a contrastive phenomenon P vs. P'* (all other things being equal) (Craver and Kaplan 2020, 300). If an explanation leaves out relevant details this will decrease its explanatory power and can only serve non-explanatory purposes (Craver and Kaplan 2020, sec. 7). However, an explanation does not have to be complete in this sense in order to have explanatory force (Craver and Kaplan 2020, 301).
- c') Adding details about a mechanism to a mechanistic explanation/model improves the explanation if and only if *the details are relevant for explaining the phenomenon P vs. P'* (Craver and Kaplan 2020, 303).

According to Craver and Kaplan, claims a'), b'), and c') are compatible with the fact that good explanations often do not mention all details about a mechanism. Mentioning neurons in the brain is unlikely to provide a good mechanistic explanation as it is unlikely that all neurons and their activities will be relevant to explaining a given contrast about a cognitive capacity. However, if the activities of all neurons in the brain were explanatorily relevant for a given contrast, adding more details about the neurons would improve the explanation. This does not imply that the less detailed explanation has no explanatory force at all. Similarly, many network explanations highlight relevant details about the underlying mechanism (e.g. the structural, causal, or functional connectivity within a mechanism; see Craver (2016)), and are, thus, explanatory. Adding details about the mechanism will improve the explanation if these details are relevant for a given contrast-to-be-explained. Importantly, adding details to an already explanatory powerful network explanation does not imply that the original explanation is thereby rejected. Mechanistic explanations are essentially multilevel (Craver and Kaplan 2020, 304).

For the purpose of this paper, we will focus on Craver and Kaplan's reply to claims b) and c). We agree that it is odd to read into the mechanistic account the idea that adding *any kind* of detail would improve a given explanation. Early on, the new mechanists (including Craver (2007a)) have made clear that mechanistic details that are to be mentioned in an explanation are individuated *relative to the phenomenon-to-be-explained*, and they have pointed out that mechanistic details need to be *relevant* to the phenomenon at hand. Especially Craver has provided a detailed account of what relevance amounts to in the mechanistic context—constitutive explanation being the primary focus. In his so-called *mutual manipulability account of constitutive relevance*, Craver makes use of Woodwardian interventionism to specify the notion of relevance:

(Mutual Manipulability Account of Constitutive Relevance)

A mechanistic detail (an acting entity, X's ϕ -ing) is constitutively relevant for a given phenomenon (an acting entity, S's ψ -ing) if and only if:

- (i) X's ϕ -ing is a spatiotemporal part of S's ψ -ing,
- (ii) there is an ideal intervention on X's ϕ -ing by means of which one can change S's ψ -ing, and
- (iii) there is an ideal intervention on S's ψ -ing by means of which one can change X's ϕ -ing. (Craver 2007a, 153)¹

Based on these assumptions about the relevance of mechanistic detail for a given phenomenon, Craver and Kaplan introduce the following two-step account as a reply to the MDB objection and as a strategy to argue for b') and c'): First, they define what they call 'Salmon-Completeness' or 'SC' to spell out in which sense ontic mechanisms are complete:

Salmon-Completeness (SC): The Salmon-complete constitutive mechanism for [the phenomenon] P versus P' is the set of all and only the factors constitutively relevant for P versus P'. (Craver and Kaplan 2020, 300)

Then, based on this notion of SC, Craver and Kaplan define under which conditions adding details to an explanation (they speak of "models" instead of "explanations") improves its explanatory power:

More Relevant Details Are Better (MRDB): If [explanatory] model M contains more explanatorily relevant [i.e., constitutively relevant] details than M* about the SC mechanism for P versus P', then M has more explanatory force than M* for P versus P', all things equal. (Craver and Kaplan 2020, 303).²

As already mentioned, mechanists have always stressed the importance of *relevance* of mechanistic detail *for a given phenomenon*. In addition, Craver and Kaplan argue that phenomena are contrasts of the form P vs. P' (such as 'Socrates died vs. remained alive', 'Socrates died quickly vs. at some other rate') (Craver and Kaplan 2020, 296). They stress that this is not an entirely new idea and not foreign to the mechanistic thinking: Craver has explained already in his Craver 2007a, b book that phenomena are "multifaceted" (Craver 2007a, 125) and that one needs to explain various contrasts in order to fully explain a phenomenon. Furthermore, the idea that mechanistic explananda are contrastive follows naturally from the application of Woodwardian interventionism to make sense of explanatory relevance. However, in the present context, this addition is still new. The reason is that the notion of constitutive relevance has, so far, only been defined for non-contrastive phenomena, i.e., only for acting entities/Ss' ψ -ings.

The general spirit of Craver and Kaplan's reply to the MDB-objection is convincing. Clearly, only adding detail that is *explanatory relevant for the phenomenon-to-be-explained* increases the power of an explanation. However, there are three

¹For the sake of argument, we ignore the challenges for the mutual manipulability account. For a discussion of these challenges see (Romero 2015; Baumgartner and Gebharder 2016; Baumgartner and Casini 2017; Kästner 2017; Baumgartner et al. 2018; Krickel 2018b).

²Craver and Kaplan use the label 'MDB_r' (with an index). For our purposes, it is more convenient to use the label 'MRDB'.

problems with the specifics of Craver and Kaplan's account and two general challenges. In the next two sections, we introduce the three problems and offer a solution. This solution will allow for a modification of Craver and Kaplan's reply without departing much from their approach. However, two challenges remain, which will be discussed in Sects. 17.5 and 17.6.

17.3 Three Problems: Odd Ontology, Multiplication of Mechanisms, Ontic Completeness

We will call the first problem for Craver and Kaplan's reply to the MDB objection the *Multiplication of Mechanisms Problem*. Given that mechanisms are individuated relative to the phenomena they are supposed to explain, the introduction of the contrastive phenomenon brings with it a multiplication not only of phenomena but also of mechanisms. On the original mutual manipulability account (see previous section), there was only one mechanism for each phenomenon such as the action potential, muscle contraction, or a rat's navigating through a maze (S's ψ -ing). Based on the contrastive interpretation of phenomena, Craver and Kaplan are committed to the view that there are multiple mechanisms—one for each contrast. For example, the action potential has many different features that may figure in such a contrast. It has a certain speed, voltage, etc., its refraction period has a certain length, and presumably other features. For each of these features, we can formulate an unbounded number of contrastive phenomena. For instance, in connection with voltage, there would be the contrastive phenomenon "action potential with voltage 70mV rather than 30mV", another one "action potential . . . rather than 35mV" and so on. None of these contrastive phenomena is inherently more or less worthy of explaining. Furthermore, each of these contrastive phenomena would individuate a mechanism that is responsible for it. This generalizes to each of the unbounded number of contrasts one can formulate with respect to action potential or any other phenomenon taken as an acting entity. But on the grounds of parsimony, such multiplication of phenomena and mechanisms should be avoided.

The second problem is the *Odd Ontology Problem*. It can be formulated as a dilemma: If mechanisms and phenomena in Craver and Kaplan's account are supposed to be ontic, then they either cannot be contrasts, or Craver and Kaplan are committed to an odd ontology. The original mutual manipulability account had a straightforward ontic reading: X's ϕ -ing and S's ψ -ing are both acting entities that are real things in the world (Machamer et al. 2000). The part-whole relation between the mechanistic component and the phenomenon is a mind-independent ontic relation between these two acting entities. Since ideal interventions need only be "logically possible" (Woodward 2003, 128), also the mutual manipulability condition is a mind-independent relation between the phenomenon and the mechanistic component. Craver and Kaplan commit to an *ontic conception of scientific explanation*. According to this conception, explanations are objective things in the

world (Craver and Kaplan 2020, sec. 5). This implies that mechanisms, which are the explanations according to the ontic conception, as well as their explanantia, i.e., phenomena, are real, ontic things. Based on the contrastive interpretation of the phenomenon, however, it becomes unclear in which sense phenomena and the relation between a mechanism and a phenomenon can be ontic. Clearly, contrasts of the form P vs. P' are not real entities. P' does not actually occur, and contrasts in this context are not entities at all but comparisons that scientists make. If at all, contrastive phenomena can be *descriptions* of things that are compared. But how does an ontic mechanism cause or constitute a contrast if the latter is a description?

The third problem, we call the *Ontic Completeness Problem*. It consists in the fact that Craver and Kaplan talk about ‘ontic completeness’ and define Salmon-Completeness (see Sect. 17.2) in terms of ontic mechanisms. It is a category mistake to speak of ontic mechanisms as complete or incomplete. Ontic mechanisms are the way they are, and it does not make sense to say that an ontic mechanism is complete or incomplete as if we had to build mechanisms similar to an IKEA cupboard where one screw is always missing. In explanatory contexts, mechanisms are already in existence. They are neither complete, nor incomplete. They just are. This reasoning is based on a thesis that may be called the *They Just Are Principle* (inspired by Craver 2007a, b, p. 27): ontic, mind-independent things on their own do not have normative or evaluative properties, they are neither good, nor bad, neither complete, nor incomplete. It is either our *descriptions* of the things in the world that are complete or incomplete; or it is a feature of a set of ontic things *relative* to a normative description (such as in the IKEA case: There should be six screws in the box!). Based on the They Just Are Principle, another way to formulate the Ontic Completeness Problem is in terms of the following argument:

1. Craver and Kaplan take explanation to be ontic: the mechanism that explains the phenomenon and the phenomenon itself are mind-independent things in the world.
 2. There is no norm about which things *should* be components of a mechanism.³
- (K1) Hence, it does not make sense to say that a mechanism that explains a phenomenon is complete or incomplete (from 1, 2, and the They Just are Principle).
3. Salmon-completeness is defined for (constitutive) mechanisms.
- (K2) Hence, according to Craver and Kaplan it is possible to say about a mechanism that it is complete or incomplete (from 3).

Thus, Craver and Kaplan run into a contradiction (between K1 and K2). However, the They Just Are Principle already suggests a way to modify Craver and Kaplan’s account such that the *Ontic Completeness Problem* can be avoided: define com-

³Note that the mutual manipulability account is not a description about what *should* be a component of a mechanism for a given phenomenon but rather a recipe for determining what *is* a component of a mechanism for a given phenomenon.

pleteness as a feature of *descriptions* of mechanisms rather than for mechanisms themselves. This is what we will do in the next section.

17.4 Solving the Problems: A Threefold Distinction & Two Notions of Completeness

Before we can provide solutions for the problems presented in the previous section, we have to do a little bit of stage setting. For the purposes of this paper, we will make the following assumptions that most new mechanists accept (including Craver and Kaplan):

Mechanism Characterization Mechanisms are entities and activities organized such that they are responsible for a phenomenon. (Machamer et al. 2000; Craver 2007a; Illari and Williamson 2012; Glennan 2017)

Etiological vs. Constitutive Mechanisms consist of those and only those acting entities that are either causally or constitutively relevant for a phenomenon. (Craver 2007a)

Constitutive Relevance Constitutive relevance is spelled out in terms of two necessary conditions: (i) spatiotemporal parthood, (ii) mutual manipulability. (Craver 2007a, b)

Levels of Mechanisms Mechanisms and mechanistic explanations come in hierarchies that are determined by relations of constitutive relevance and that are local to the phenomenon-to-be-explained. (Craver 2007a; Craver and Bechtel 2007)

Phenomena Phenomena are acting entities and to explain a phenomenon means to explain various contrasts. (Craver 2007a; Kaiser and Krickel 2017)

Singularism/Nominalism Mechanisms, entities, and activities are concrete particulars. Types are descriptions/models summarizing details about concrete particulars. (Glennan 2017; Krickel 2018a)

Purpose of Explanation The core function of explanation is to show how a phenomenon is situated in the causal structure of the world (Craver 2007a, 200), chiefly for the purpose of *intervening* into the phenomena (Craver 2007a, 93). Abstraction, in the sense of *ignoring explanatorily relevant details*, has only non-explanatory virtues (Craver and Kaplan 2020, sec. 7).

Unique Endeavour Explaining a phenomenon is a unique scientific endeavor that is distinct from prediction and description (Craver and Kaplan 2020, sec. 3).

Most new mechanists (including Craver and Kaplan) would also accept the following commitment:

Explanatory Relevance Explanatory relevance is constitutive relevance (in the case of constitutive explanation) or causal relevance (in the case of etiological explanation).

However, the equivalence between constitutive and explanatory relevance cannot hold assuming the contrastive view of explananda and the view that constitutive

relevance holds between ontic things. Explanatory relevance and constitutive relevance are identical only if explanatory relevance is a relation between ontic things as well. If explananda are contrasts, however, we run into the Odd Ontology Problem if we assume that they can stand in constitutive relevance relations. If we resolve the Odd Ontology Problem by viewing contrastive explananda as descriptions, then they cannot enter into relations of constitutive relevance. One way to solve this problem is to reject that explanatory relevance and constitutive relevance are identical and allow for explanatory relevance to relate descriptions of ontic things that stand in constitutive relevance relationships. Our account will ultimately decouple explanatory relevance from constitutive relevance in that way.

It is important to note that unlike Craver and Kaplan, not all new mechanists accept the ontic conception of explanation mentioned in the previous section. Neither will we, as already indicated in the previous paragraph. Indeed, as we will argue in the next section, it is the conflation between ontic mechanisms and explanations that leads into the two problems mentioned in the previous section. Roughly, our solution will be to introduce a threefold distinction that helps us to keep apart the ontic and the epistemic aspects of explanation (Sect. 17.4.1). This threefold distinction will illuminate the role and meaning of completeness within the mechanistic framework (Sect. 17.4.2).

17.4.1 Ontic Mechanisms, Mechanism Descriptions, Explanatory Texts

In order to be able to unambiguously talk about ontic and epistemic issues, about matters of description vs. matters of explanations, we will distinguish between three elements:

1. **Ontic phenomena and mechanisms:** Ontic phenomena and ontic mechanisms are concrete particulars (see *Singularism/Nominalism* above). Ontic phenomena are acting entities such as a neuron firing, an axon terminal releasing neurotransmitter, a muscle contracting, a mouse navigating the Morris Water maze (Kaiser and Krickel 2017). Ontic phenomena are objects of explanatory endeavors, and targets of investigations. The mechanistic ontology is committed to the view that ontic phenomena are constituted by equally real ontic mechanisms (Illari and Williamson 2011), which are composed of acting entities with a spatiotemporal organization particular to each phenomenon. However, explanatory practices are only mediately concerned with ontic phenomena and mechanisms. Instead, explanation consists in constructing two types of texts: mechanism descriptions and mechanistic explanatory texts.
2. **Mechanism descriptions:** Mechanism descriptions are texts or other knowledge items that can be found in textbooks, journals, or other scientific media. The ontic mechanism is the truthmaker of the mechanism description. Mechanism descriptions are not guided by any particular explanatory interest but aim at

neutral description of the mechanism that later (via explanatory texts — see below) can be used for various explanations. Ideally, mechanism descriptions mention *all acting entities that are constitutively relevant* for a given ontic phenomenon *with maximal detail*. For example, ideally, the description of the mechanism responsible for a neuron's firing will mention, say, how many ions and ion-channels are involved, where they are located, what size they have, etc. for every point in time of the occurrence of the mechanism. The description of a single mechanism may span a number of publications, with only a part of the whole description exhibited in one place. In this they are close to Craver and Kaplan's "stores of explanatory knowledge". It is important to note that this is to be understood as a regulative ideal (see Railton (1980) for a similar view). Much scientific work consists in refining mechanism descriptions and filling in any gaps in them, although in practice, all mechanism descriptions actually available in the scientific community are incomplete.

3. **Mechanistic explanatory texts:** Mechanistic explanatory texts are vehicles of explanation, i.e., they are the explanantia. Each mechanistic explanatory text is an answer to a particular why-question. Why-questions, in our account, following Dretske (1972) and the spirit of Craver and Kaplan (2020) require explaining a particular contrast, whether explicitly, or implicitly stated. Mechanistic explanatory texts contain information from mechanism description relevant for explaining a particular contrastive explanandum. Note that it is only in mechanistic explanatory texts that contrasts play a role. Neither ontic phenomena, nor mechanism descriptions are in any way concerned with contrasts. Although mechanistic explanatory texts depend on mechanism descriptions, in practice even incomplete mechanistic descriptions can furnish the researcher with enough information to construct numerous mechanistic explanatory texts concerning various contrasts. Additionally, research that aims at answering particular why-questions, i.e. at constructing particular mechanistic explanatory texts can lead to the discovery of hitherto unknown ontic constituents, thus enriching the overall mechanism description. The question remains, however, what information goes into the mechanistic explanatory text and whether these texts always improve with the addition of further details. This will be taken up in Sect. 17.5.1.

As we will see in Sect. 17.4.2, the distinction between mechanism descriptions and mechanistic explanatory texts allows us to formulate different norms of completeness for descriptions and explanatory texts. Kaplan and Craver's talk of mechanistic models which at the same time describe a mechanism for a phenomenon and provide the explanatorily relevant factors for a contrast precludes one from acknowledging that depending on the purpose of the model a different completeness norm is appropriate. Therefore, what Craver and Kaplan call "mechanistic models" can on a case-by-case basis be classified as either mechanism descriptions or mechanistic explanatory texts.

One advantage of this threefold distinction is that it allows us to maintain the idea that mechanistic explanations have to pick out the ontic relations between a mechanism and a phenomenon—in contrast to a strict epistemic view that assumes

that explanation has purely pragmatic or psychological functions. Furthermore, it is only by drawing a distinction between mechanism descriptions and explanatory texts that one can account for the idea that to explain a phenomenon is more than merely describing the mechanism that is responsible for it. Only *specific* ways of describing the ontic mechanism are explanatory. For example, assume that a certain volume of water has a temperature of $-17\text{ }^{\circ}\text{C}$ and is therefore frozen (under otherwise normal conditions) (Craver and Kaplan 2020, 304–305). If we want to explain as to why the water is frozen in contrast to not frozen, the explanation will mention that the temperature is below $0\text{ }^{\circ}\text{C}$ rather than simply that the temperature is $-17\text{ }^{\circ}\text{C}$. Both answers provide correct descriptions of the ontic mechanism and are thus explanatory according to the ontic conception of scientific explanation. However, only the fact that the temperature is below $0\text{ }^{\circ}\text{C}$ is explanatorily relevant as it does not matter whether the temperature is -17 , -16 , . . . , or $-1\text{ }^{\circ}\text{C}$ for its being frozen. Take another example: assume that we want to explain why neurotransmitter are released rather than not released. As a matter of fact, the release was caused by the rise in intracellular Na^{+} concentration. As both descriptions of the neurotransmitter release mechanism in terms of ‘rise in intracellular Na^{+} concentration’ as well as in terms of ‘rise in the membrane voltage’ are true and causally relevant, both aspects should count as equally explanatory according to the ontic conception of scientific explanation. However, what is explanatory relevant is the rise in the membrane voltage that goes along with it and that could have been induced by any other cations (Craver 2007a, 205). This shows that not all ways of describing an ontic mechanism are explanatory. Only describing a mechanism by providing what we call a “mechanistic explanatory text” is explanatory. We will say more on how explanatory texts are generated in Sect. 17.5.1.

A further advantage of our threefold distinction is that it allows us to solve the Multiplication of Mechanisms and the Odd Ontology problem. We can uphold the idea that *explananda are contrastive* without being committed to the view that ontic phenomena are contrastive. We can thereby avoid being committed to an odd ontology without having to deny that phenomena are ontic. Explananda are Why-questions or requests for explanation of the form ‘Why P rather than P’?. Explanantia are explanatory texts, i.e., representations of ontic mechanisms that identify the elements of a mechanism description that are relevant to a given explanatory contrast P vs. P’. Furthermore, our account does not multiply mechanisms. Mechanisms are not individuated with regard to explanatory contrasts. Rather, they are individuated relative to ontic phenomena, which are acting entities, such as neurons that fire, rats that navigate mazes, or muscles that contract (Craver 2007a; Kaiser and Krickel 2017). If at all, we multiply explanatory texts. It makes sense to say that the number of contrasts that can be formulated for a given ontic phenomenon is at least possibly infinite.

Finally, our threefold distinction between ontic mechanisms, mechanism descriptions and mechanistic explanatory texts helps us to solve the Ontic Completeness Problem. Ontic mechanisms are the way they are, and it does not make sense to say that an ontic mechanism is complete or incomplete. What can be complete or incomplete, though, is knowledge about or representations of ontic mechanisms.

These representations can be either mechanism descriptions or explanatory texts. However, the norms of completeness for mechanism descriptions, on the one hand, and for explanatory texts, on the other, are different. The different norms of completeness will be developed in the next section.

17.4.2 *Descriptive Vs. Explanatory Completeness*

As already indicated in our characterizations of mechanism descriptions and explanatory texts, the primary difference between the norms of completeness for mechanism descriptions and mechanistic explanatory texts is that mechanism descriptions obey *norms for descriptive completeness*, while mechanistic explanatory texts obey *norms for explanatory completeness*.

As was stated in Sect. 17.4.1, mechanism descriptions should describe every detail of the ontic mechanism in order to be descriptively complete. We will, thus, formulate a notion of descriptive completeness for mechanism descriptions:

(Descriptive Completeness) A mechanism description is descriptively complete if and only if it contains all and only the constitutively relevant details about the ontic mechanism for phenomenon P.

Here, ‘constitutive relevance’ can be straightforwardly interpreted in line with the original mutual manipulability account as presented in Sect. 17.2. Based on this notion of ontic completeness, a ‘More Relevant Details Are Better’-claim for mechanism descriptions can be derived:

More Descriptively Relevant Details are Better (MDRDB): If a mechanism description D contains more constitutively relevant details than D* about the ontic mechanism for phenomenon P, then D has more descriptive power than D* for P, all things being equal.

As we saw in Sect. 17.4.1, we endorse this MDRDB claim as a regulative ideal for mechanism descriptions. However, this does not imply that mechanistic *explanations* are always better with additional detail, since mechanism descriptions are not vehicles of explanation on our account.

Completeness of explanatory texts is more closely tied to the explanatory interests expressed by the request for explanation. Furthermore, as specified in the characterization of our threefold distinction, explanatory texts are formed based on mechanism descriptions rather than ontic mechanisms directly. Therefore, we will call the norm of completeness for explanatory texts ‘explanatory completeness’:

(Explanatory Completeness) An explanatory text is explanatorily complete if and only if it mentions all and only the explanatorily relevant details for P vs. P’ contained in the mechanism descriptions for P and P’.

As we saw above, in order to make sense of explanatory completeness, the notion of explanatory relevance cannot straightforwardly be interpreted along the lines of constitutive relevance. We provide a positive account of explanatory relevance in

Sect. 17.5.1. Still, we can formulate a ‘More Relevant Details Are Better’-claim for explanatory texts based on the norm of explanatory completeness:

More Explanatorily Relevant Details are Better (MERDB): If an explanatory text T contains more explanatorily relevant details for P vs. P' than T* from the mechanism descriptions for P and P', then T has more explanatory power than T* for P vs. P', all things being equal.

Drawing the distinction between descriptive and explanatory completeness is important for the following reason: remember the example of the freezing water; the ontic mechanism, as a matter of fact, involves a temperature of -17°C . However, the exact value is not explanatorily relevant for why the water is frozen. Rather, the temperature's being below 0°C is explanatorily relevant. Still, as a matter of fact, the property of having a temperature of below 0°C is realized by the actual temperature of -17°C . An ontic mechanism cannot have the property of involving a temperature of below 0°C without instantiating a specific value. This indicates that the norms for evaluating the completeness of mechanism descriptions differ from the norms of completeness of explanatory texts. Take another example. Imagine a neuron's firing that is brought about by the transmission of electric current along an axon where exactly 14,560 ions are involved. Among the components of the ontic mechanism, thus, are 14,560 ions. However, for the explanation of why a specific neuron fired rather than not, the exact number of ions is irrelevant. What matters is that some number of ions above some threshold >0 is involved. The difference is captured by the contrastive formulation of the phenomenon: it only matters that *some rather than none* ions are involved in order to explain *why the neuron fired rather than not*.

Opponents of the mechanistic account might have confused descriptive and explanatory completeness, or in other words, they may have mistaken mechanism descriptions for mechanistic explanations. If one, as Craver and Kaplan do (though they use the term ‘model’), insists that explanations have a contrastive explanandum, it becomes clear that descriptive completeness is not an ideal of mechanistic explanation (though it is an ideal of mechanism description as formulated above).

However, given that Craver and Kaplan define Salmon-Completeness already with respect to a contrastive phenomenon, they blur the distinction between explanatory completeness and descriptive completeness. Thereby they are committed to the view that ontic mechanisms have properties such as ‘a temperature below 0°C ’ without any determinant of this property being among their components. The mechanism cannot be said to have a temperature of -17°C as this property is explanatorily irrelevant for why the water is frozen. Similarly, they are committed to the claim that the neurotransmitter release mechanism is composed of just a rise of the membrane voltage without any corresponding rise in Na^+ . However, determinables (such as ‘a temperature below 0°C ’ or ‘a rise in the membrane voltage’) have to be realized by a determinant (such as ‘a temperature of -17°C ’, or ‘rise in intracellular Na^+ concentration’) in order to exist. It would be odd to assume that the ontic mechanism consists of the determinables but does not contain the determinants. By introducing the distinction between ontic mechanisms, mechanism

descriptions, and explanatory texts, we can account for the reality of determinables: ontic mechanisms contain the determinants (such as a temperature of -17°C) and these determinants are mentioned in the mechanism description; explanatory texts however may mention determinables (such as a temperature of below 0°C) that are explanatorily relevant for a given explanatory contrast.

In a nutshell: the general spirit of Craver and Kaplan's reply to the MDB objection, i.e., that mechanistic explanations are only improved by adding details if the details are explanatorily relevant to a given contrastive phenomenon, is correct. However, in order to avoid the Multiplication of Mechanisms Problem, the Odd Ontology Problem, and the Ontic Completeness Problem, their reply has to be modified. We introduce a distinction between ontic mechanisms, their descriptions, and the explanatory texts that are generated based on the descriptions. However, two challenges remain—as we will show in the next section.

17.5 Two Challenges: Constitutive Relevance & Vertical Completeness

The two remaining challenges are not only challenges for Craver and Kaplan's reply to the MDB-objection but for the mechanistic account in general. The first challenge stems from the fact that, on the one hand, Craver and Kaplan and many other mechanists want to think of the explanantia of mechanistic explanation in terms of contrasts (i.e., what we call explanatory texts). On the other hand, they hold that explanatory relevance is constitutive relevance. However, constitutive relevance is spelled out in terms of mutual manipulability between *ontic phenomena and their spatiotemporal parts* and not in terms of a contrastive account of phenomena. There is at least a gap here: How can we determine what is to be part of a mechanistic explanatory text and how can this be combined with constitutive relevance? We will discuss and answer this question in Sect. 17.5.1.

The second challenge stems from the fact that Craver and Kaplan only address *one version* of the MDB-objection. We will show that there are two different versions of this objection—the *vertical* and the *horizontal* version. So far, there is no successful answer to the vertical version of the objection. We will discuss this objection and a possible reply in Sect. 17.5.2.

17.5.1 Explanatory Relevance, Contrasts, and Constitutive Relevance

As we saw in Sect. 17.4.1, mechanists typically equate explanatory relevance with constitutive relevance. However, this option is not open to any account of mechanistic explanation which views explananda as contrastive. Craver and Kaplan

persist in viewing explanatory relevance and constitutive relevance as equivalent, but the mutual manipulability account of constitutive relevance has never been explicitly combined with a contrastive account of the phenomenon. Indeed, as presented earlier, the mutual manipulability account of constitutive relevance defines under which conditions an acting entity (X 's ϕ -ing) is constitutively relevant for another acting entity (S 's ψ -ing). In their paper, Craver and Kaplan do not explain how the combination of the mutual manipulability account and a contrastive phenomenon is supposed to work.

In our view, explanatory relevance is distinct from constitutive relevance, though constitutive relevance constrains explanatory relevance in the sense that all explanatorily relevant details concern mechanism constituents, i.e., entities and activities that are constitutively relevant to the explanandum phenomenon. However, not all constitutively relevant detail is explanatorily relevant. Therefore, we must provide a further criterion for identifying explanatorily relevant details.

Giving a satisfactory account of explanatory relevance requires spelling out exactly the relation between mechanism descriptions and mechanistic explanatory texts. As we saw in Sect. 17.4.1, mechanistic explanatory texts answer contrastive why-questions. The why-question, or *explanation request*, determines the class of contrast phenomena P' , with which the actual phenomenon is compared. For example, one might ask "Why is the water in the glass frozen?". This explanation request is incomplete, because it does not specify a contrast class explicitly. However, implicitly, we could read the contrast in as "Why is the water in the glass frozen, rather than thawed?". This picks out the phenomenon P : *water in the glass is frozen*, and the contrast class P' : *water in the glass is thawed*. Other why questions could be posed which pick out different contrast classes. For instance, one could ask: "Why is the water in the glass frozen rather than evaporating?". This why-question picks out the same phenomenon P , but a different contrast class P' : *water in the glass is evaporating*.

Based on the mechanistic idea that explanation is about finding crucial points of intervention (see *Purpose of Explanation* in Sect. 17.4), explanatory texts should identify these crucial points of intervention. Crucial points of intervention with respect to the contrast P vs. P' are those which allow one to change P into P' with *minimal effort*. A preliminary characterization of the contents of mechanistic explanatory texts, thus, is the following:

(*Contents of METs_{preliminary}*) A mechanistic explanatory text T explaining a contrast " P vs. P' " has the form "because C rather than C' ", where C is a set of constituents of the description of the actual mechanism M_{actual} for P and C' is a set of constituents of a description of a possible mechanism M_{possible} , where the following holds:

- (i) M_{possible} is a member of a set S of possible mechanisms each sufficient to bring about P' ,
- (ii) C and C' contain all and only constituents that differ in the description of M_{actual} and the description of M_{possible} and that are also differences between M_{actual} and all other members of S ,
- (iii) M_{possible} could in principle be created by means of *intervening into* M_{actual} with *minimal effort* compared to the effort that would be necessary for the creation of each other member of S .

Condition (i) in the definition above should be self-explanatory. The mechanistic explanatory text explains why P occurred rather than P', by exhibiting the constituents which were crucial to the occurrence of P and not P'. Therefore, we must compare the mechanism for P with some mechanism for P', in order to isolate these crucial constituents. There are usually many ways in which a phenomenon from the contrast class P' could be constituted. Consider the example of the action potential. If we want to explain why the neuron fired as opposed to maintaining resting membrane potential, we must contend with the fact that there are innumerable differences within the class of neurons maintaining resting potential (e.g., in the number of ion channels open, etc.). We must select one mechanism for P' from the set of possible mechanisms to compare to the actual mechanism.

Condition (ii) further specifies that once we have selected one particular possible mechanism for comparison, we are interested only in those differences between this contrast mechanism and the actual mechanism for the phenomenon that apply to all the other mechanisms sufficient to bring about the contrast phenomenon. This is what enables us to say that the correct explanation for "Why is the water in the glass frozen rather than thawed?" is that the temperature of the water is below 0 °C, or more exactly, that the temperature of the water is -17 °C rather than *above* 0 °C. The difference between -17 °C in the actual mechanism and above zero degrees holds for all instances of the contrast class (all other things being equal).

Condition (iii) identifies which of the many possible mechanisms sufficient for bringing about the contrast phenomenon should serve as the basis of the comparison. 'Minimal effort' can be defined as a function of (a) the *amount* of required interventions, and (b) the *similarity* of the required interventions (the more similar the required interventions are, the least effort). Hence, we have to know how to count and compare interventions in order to identify the mechanism M_{possible} that is to figure in our comparison.

Woodward (2003, 98) defines interventions with the help of three variables $\{I, X, Y\}$, where I is the intervention variable. When I takes a particular value, it sets the value of X ; X is the putative cause; and Y the putative effect. Since we are here concerned with constitutive rather than causal explanation, we do not take X and Y to be putative causes and effects. Instead X is a variable representing the mechanism constituent, and Y is a variable representing the phenomenon P (or temporal parts thereof, see Krickel (2018b)). According to this view, two interventions are identical, if and only if their defining I -variables, X -variables, and Y -variables are the same. Since we are interested in mechanism descriptions that describe mechanisms that are sufficient to bring about the contrast phenomenon P', we can assume that Y is the same for all relevant interventions (Y represents the phenomenon and can take the values that represent P and P'). Hence, in order to determine which mechanism description of the counterfactual mechanisms that bring about P' are most similar to the mechanism description of the actual mechanism, we have to determine which counterfactual mechanism requires the least effortful interventions individuated by their I s and their X s, giving us a set of required interventions $\{\{I_1, X_1\}, \dots, \{I_n, X_n\}\}$ for each possible mechanism for P'.

In order to determine which of these sets of interventions involves the ‘least effort’, we have to know how to count *I*s and *X*s and how to determine their similarity.

There is a practical problem for the counting and comparing of the intervention variables I_1 - I_n . In many cases, scientists know *what* would have to be changed in order to build a counterfactual mechanism M^* from an actual mechanism M . However, they do not know *how* this change could be brought about. For example, scientists do often know what component of a pathological mechanism is responsible for the symptoms and therefore know that changing this component would lead to an improvement of the symptoms. But they do not know how to change this component. Much effort in medical research goes into inventing better drugs to be able to change mechanisms in the right way. In the present context, the consequence is that, in practice, formulating the explanatory text for a given contrastive request for explanation often is conditional on the fact that we do not know what the intervention variable represents and how similar or different it will be compared to other interventions. We therefore need to decouple the measure of minimal changes from the count of intervention variables. However, the measure we ultimately choose should respect the interventionist insight that the number of intervention variables matters. This can be achieved if we make our measure sensitive to similarities and differences between *X*s. If the targets of interventions are similar in specific ways, it is likely that they can be intervened on with just a single intervention variable *I*.

The practical problem does not arise for the counting and comparison of the *X*s. Counting and comparing *X*s means to count and compare mechanistic components. These mechanistic components, in our framework, are described in the mechanism descriptions. Hence, in the end in order to determine which interventions require the least effort, we have to count and compare the differences between the *descriptions* of all (nomologically) possible mechanisms for P' . This results in the following characterization of the contents of mechanistic explanatory texts:

(*Contents of METs*) A mechanistic explanatory text T explaining a contrast “ P vs. P' ” has the form “because C rather than C' ”, where C is a set of constituents of the description of the actual mechanism for P M_{actual} and C' is a set of constituents of a description of a possible mechanism M_{possible} , where the following holds:

- (i) M_{possible} is a member of a set S of possible mechanisms each sufficient to bring about P' ,
- (ii) C and C' contain all and only constituents that differ in the description of M_{actual} and the description of M_{possible} and that are also differences between M_{actual} and all other members of S ,
- (iii) the description of M_{possible} is more similar to the description of M_{actual} than the description of any other member of S .

Thus, we have to be able to determine when a mechanism description D is more similar to another mechanism description D' than some further mechanism descriptions D^* , D^{**} etc.

Let us look at two more informal examples of the kind of comparison that figures in constructing METs. For example, suppose we are trying to construct an explanatory text answering the question: “Why did the car go straight, as opposed

to turning right?'. The mechanism description for the ontic mechanism in which the car is going straight at speed 90 km/h with rattling bumpers includes a number of constituents describing the activity of the spark plugs. However, the mechanism description of the most similar mechanism M_{possible} , which would underlie the car's turning right includes the very same constituents describing the activity of spark plugs. Therefore, when answering the question regarding going straight in contrast to turning right, this information will not be included in the explanatory text. Spark plug activity is not different across the two cases. To make a car turn right, rather than go straight, one should intervene on the wheels, not on the spark plugs.

In practice, the problem of constructing the correct MET will be compounded by the fact that there may be numerous ways of exhibiting the contrast phenomenon. For instance, let's look at explaining why the car goes straight rather than standing still. Will the explanatory text mention spark-plug activity? Perhaps surprisingly, the answer is still no. Although the paradigm case in which the car stands still is one where the engine does not run, and spark plugs do not spark, there is another class of situations in which cars stand still, i.e. when they are idling with the engine running in neutral gear, or when brakes have been applied. In these cases, spark-plugs do spark in the same way as when the car goes straight. The mechanism description for the idling case, or the braking case will be closer to the description of the actual mechanism, because all the (many) engine parts will work in the same way, and thus receive the same description, as in the actual mechanism. In fact, the contrast class might be too heterogeneous to admit any set of differences satisfying point (ii) of the definition. This would suggest that the contrast must be explained piecemeal.

The matter of comparing mechanism descriptions is complicated by the fact that mechanism descriptions can be given in various forms, such as spoken word, written text, diagram, etc. and two mechanism descriptions can contain the same information about the same mechanism, even though they superficially differ. In order to resolve this issue, we stipulate that mechanism descriptions can be transformed into a canonical form:

(Canonical Form of MDs): A mechanism description in its canonical form is a set of 4-tuples $\langle E, A, S, T \rangle$, where E stands for some entity, A, for the activity this entity is performing, S for the (relative) spatial region in which this activity is performed, and T for the time during which the activity is performed. A single mechanism description will consist of many such 4-tuples strung together.

In the rest of Sect. 17.5, we will need to distinguish between 'constituents', i.e., 4-tuples in a mechanism description and 'elements', which are any of the 4 parts which make up a constituent. Note that constituents in a mechanism description describe ontic constituents. When we refer to constituents of ontic mechanisms, this will always be specified in full. Mechanism descriptions in sentential or diagrammatic form can be, at least in principle, converted to this canonical form.

Two further questions arise with respect to mechanistic descriptions: the question of grain, and the question of sameness. The question of grain asks how detailed mechanism descriptions are. In practice, the answer varies, because different particular mechanism descriptions will be exhibited with varying detail. However,

in Sects. 17.3 and 17.4 we saw that mechanism descriptions follow the norms for ontic completeness. This means that we can at least specify the conditions under which a mechanism description is better than another one characterizing the same ontic mechanism. The answer is in line with the MRDB claims formulated in the previous section: the more detailed a mechanism description the better. The best mechanism description describes all ontic constituents, and it describes all of them to maximum detail. A scientific community which has more fine-grained mechanism descriptions at its disposal is better off than a scientific community with only coarse-grained mechanism descriptions. This is because the former scientific community can explain more contrasts than the latter one. Further note that in practice scientific communities have descriptions at various levels of grain available to them, and they can construct coarser descriptions if need be by substituting less determinate denotations for entities, activities, places and times. Thus, the scientific community with the more fine-grained description will always also be in possession of the coarse-grained description.

Secondly, when are two mechanism descriptions equivalent? For mechanism descriptions in non-canonical forms, the answer is simple — two such descriptions are equivalent if and only if they can be transformed into the same canonical form without adding or leaving out any empirical content. Two mechanism descriptions in the canonical form are the same, if they contain the same constituents.

Can the comparison between mechanism descriptions be formalized such that a general recipe for how to compare mechanisms can be made available? Our proposal is that the minimal set of differences between mechanism descriptions M and M^* can be computed based on adapting the concept of generalized *edit distance*. This measure is frequently used in computer science to reason about string matching and indirectly about graph matching. Adopting this framework is licensed by the fact that mechanism descriptions can be transformed into our canonical form.

In computer science edit distance of string s from s^* is based on the number of steps required to transform s into s^* . Each step consists of applying one of a set of permitted edit operations to one character of s . Different applications sanction different sets of permitted edit operations. Additionally, a cost, or weight, is associated with each permitted edit operation. The edit distance is the sum of these costs (Cohen et al. (2003); see also papers cited therein). The most well-known version of edit distance for strings is the so-called Levenshtein distance (Levenshtein 1966). Levenshtein distance permits the following operations: character insertion, character deletion and character replacement, all of which are of equal cost 1. The Levenshtein distance from the string ‘dogged’ to ‘froggy’ is 4 – 1 addition, 2 replacements, and 1 deletion.⁴ Other related measures use a more restrictive set of edit operations (e.g. disallowing direct replacement; Wagner and Fischer (1974)), or on the contrary a more permissive set of edit operations (e.g. allowing direct transposition, etc.; Damerau (1964)). For some applications, weights different from 1 are used, so that some operations are more costly to perform than others (Monge

⁴‘dogged’ → ‘fdogged’ → ‘frogged’ → ‘froggyd’ → ‘froggy’

and Elkan 1997). Although the edit-distance framework was originally devised for imprecise string-matching, similar measures have now been used for comparing graphs, such as semantic networks (Bunke and Shearer 1998).

A version of edit distance can be straightforwardly applied to mechanism descriptions in their canonical forms. Instead of performing edit operations on characters in a string, we can define edit operations on constituents in mechanism descriptions. Two mechanism descriptions M and M^* are the same, if the edit distance from M to M^* is 0.⁵ Alternative ways of exhibiting contrast phenomena described by contrast mechanism descriptions M^* , M^{**} , M^{***} etc. can be ranked according to their edit distance from the actual mechanism description M . The one with the lowest edit distance from M is the appropriate contrast.

At this point we are left to specify the appropriate set of edit operations for mechanism descriptions. Firstly, there are straightforward equivalents for insertion and deletion. Constituent-addition and constituent-deletion are equivalent to character-insertion and character-deletion respectively. There is also an operation roughly equivalent to character-replacement. This is element-replacement, which consists in replacing one of the 4 elements in a constituent with a different element of the same category. Thus, one is permitted to change $\langle E_i, A_i, S_i, T_i \rangle$ to any of the following: $\langle E_i^*, A_i, S_i, T_i \rangle$, $\langle E_i, A_i^*, S_i, T_i \rangle$, $\langle E_i, A_i, S_i^*, T_i \rangle$ and $\langle E_i, A_i, S_i, T_i^* \rangle$.

Constituent-insertion and constituent-deletion are both weighted 1. Element-replacement, on the other hand, is weighted 0.5. This is to ensure that the cost for changing $\langle E_i, A_i, S_i, T_i \rangle$ into $\langle E_i^*, A_i^*, S_i^*, T_i^* \rangle$ is higher than the cost of changing fewer elements in a constituent. The distance from $\langle E_i, A_i, S_i, T_i \rangle$ to $\langle E_i^*, A_i^*, S_i^*, T_i^* \rangle$ is 1.5, by 3 element-replacements. The distance from $\langle E_i, A_i, S_i, T_i \rangle$ to $\langle E_i^*, A_i^*, S_i^*, T_i^* \rangle$ is 2 — either 4 element-replacements at 0.5 each, or 1 constituent-deletion and 1 constituent-insertion at 1 each.

Apart from equivalents for the standard string edit operations, we introduce an edit operation unique to comparing mechanism descriptions, called *mass element-replacement*. Mass element-replacement is our attempt at discounting such systematic changes to multiple constituents, for which a single intervention variable is likely to be responsible. Such systematic changes should be discounted because in formulating mechanistic explanatory texts we are interested in finding crucial points of intervention, where one can intervene with minimal effort.

In mass element-replacement, applying the same change to a group of relevantly similar constituents has the same cost as a simple element-replacement: 0.5. By relevant similarity, we mean that:

- a) The entity elements E of these constituents can be subsumed by the same type description. Thus, we can apply a change to, e.g., all constituents whose entity element is an electron.

⁵This criterion is equivalent to the one on p. 21 above. The edit distance from M to M^* is 0 iff M and M^* have the same constituents.

- b) The activity element A of these constituents can be subsumed under the same type description. Thus, we can apply a change to, e.g., all constituents whose activity element is a fission.
- c) The space element S of these constituents falls in a specific range, say a sphere, with a defined centre and radius.
- d) The time elements T of these constituents are synchronous or fall into a determined interval.

The range of constituents targeted by a mass element-replacement at once can be narrowed down by specifying that they be similar according to two or more of these similarity criteria. For example, we can specify that we want to change constituents involving electrons, but only those within 20 centimeters of an electric coil. The change applied to a group of constituents specified in this way need not concern the element on which their similarity depends. We can specify a group of constituents by noting that they involve electrons, and systematically change their locations in space, or slow them down, for instance.

The notion of applying the *same change* to a group of constituents also requires elucidation. Mass element-replacements are element-replacements on every constituent in the specified group. However, only certain types of replacement should be discounted. Specifically, we propose that:

- a) The activity elements of a group of relevantly similar constituents can be replaced at once with cost 0.5, if all the activity elements A are replaced by elements A* which can be subsumed under the same type descriptions. Replacing all fissions in the mechanism description with fusions, for example, is an edit operation with cost 0.5.
- b) The space elements of a group of relevantly similar constituents can be replaced at once with cost 0.5 by scaling (changing the size of constituents by a constant ratio), translation (moving the constituents in a uniform fashion, e.g., all 30 cm to the right) and rotations (tilting the constituents over).
- c) The time elements of a group of relevantly similar constituents can be replaced at once with a cost of 0.5 by scaling (changing the duration of constituents by the same ratio).

The mechanism description edit distance with these edit operations (constituent-insertion, constituent-deletion, element-replacement, and mass element-replacement) as specified here, is meant to guide judgments about the minimal differences between mechanism descriptions in a way that parallels the results one would obtain by counting interventions, without requiring us to know the intervention variables but only based on the target variables X_i - X_n . In particular, the rules for mass-element-replacement are founded on the intuition that similar things can be changed by a single intervention in a systematic way.

At the same time, however, this proposal is provisional, and subject to amendment as the framework is further developed. In this paper, we include it to demonstrate the possibility of developing sophisticated semi-formal modes of reasoning about mechanistic explanation.

17.5.2 *The Vertical Version of the MDB-Objection*

As explained above, the second challenge for Craver and Kaplan's reply consists in the fact that they only address the *horizontal version of the MDB-objection* but do not provide a satisfying reply to the *vertical version*. In this section, we explain the difference between these two versions and why Craver and Kaplan fail to address one of them.

Ontic mechanisms form hierarchies in such a way that the same acting entity can be a phenomenon that is constituted by a mechanism, but at the same time a constituent in another higher-level phenomenon (Craver 2007a, chap. 5). Therefore, mechanistic hierarchies can be said to have a horizontal and a vertical dimension. The *horizontal dimension* of mechanism hierarchies is the one along which constituents are related by non-constitutive causal relations and by relations of temporal precedence. It is called 'horizontal', because it corresponds to the horizontal axis of the Craver diagram (Craver 2007a, 121). The *vertical dimension* of mechanism hierarchies is the one along which constituents are related by part-whole relations. This corresponds to the vertical axis of the Craver diagram.

Based on this distinction, the MDB-objection can be read as a claim about the *horizontal* dimension of hierarchies of mechanisms or as a claim about the *vertical* dimension of hierarchies of mechanisms. For example, opponents accuse the new mechanistic account of claiming that adding more *horizontal details* to an explanation by, say, listing the exact positions of ion-channels improves an explanation. And they object that the new mechanistic account implies that adding *vertical detail*, by adding, for example, details about quarks always improves an explanation. As we will show, Craver and Kaplan's reply to the MDB-objection accounts only for horizontal completeness but fails to account for vertical completeness. To see this, more accurate definitions of horizontal and vertical completeness are required.

Horizontal completeness can be defined for mechanism descriptions as well as explanatory texts:

(Horizontal Completeness_{description}) A mechanism description is horizontally complete if and only if the description mentions at least one set of constitutively relevant acting entities of the ontic mechanism for phenomenon P that is minimally sufficient for bringing about P.

(Horizontal Completeness_{text}) An explanatory text is horizontally complete if and only if the text mentions all explanatorily relevant factors for P vs P' from the minimally sufficient sets of acting entities mentioned in the horizontally complete mechanism description for phenomenon P and the horizontally complete mechanism description for P'.

The acting entities contained in a horizontally complete set of mechanism components will either constitute the mechanism at different points in time (i.e., they form horizontal chains), or they will occur in different spatial location. This is implied by the requirement that the set of acting entities described in horizontal completeness_{description} must be *minimally sufficient*. In the given context, the term 'minimally sufficient' is supposed to imply that the set of acting entities does not contain redundant members, i.e., acting entities that, given the other members of the

set, do not make a difference to the higher-level phenomenon. This implies that the acting entities that are members of this set will not spatiotemporally overlap.

The two notions of horizontal completeness are usually not co-extensional. A horizontally complete mechanism description will usually mention more acting entities and more details about them than an explanatory text based thereon. On the assumption that the physical realm is causally closed, and each physical event has a physical effect (we ignore quantum events for the sake of argument), an ontic mechanism will at each point in time of its occurrence be composed of at least one acting entity. A mechanism description, ideally, mentions all of them. However, horizontal completeness of explanatory texts is compatible with there being gaps in the text. For the explanation of, say, why the action potential peaks at +40 mV rather than +50 mV it may not be explanatorily relevant what happened one millisecond after the stimulus onset. Hence, the explanatory text will be silent about what happened one millisecond after stimulus onset and leave a gap.

Horizontal completeness in the sense defined above is a goal of mechanism description as well as of explanation. A mechanism description that is not horizontally complete misses some acting entities that are crucial for the occurrence of the ontic phenomenon; in other words, it misses some parts of what might be called the ‘constitutive basis’ of a phenomenon. Similarly, an explanatory text that is not horizontally complete does not fully explain why the phenomenon P occurred rather than P’. Thus, the closer a description or text is to horizontal completeness the better. Two MRDB-claims can be formulated (where ‘horizontal’ details are those that bring us closer to horizontal completeness):

Horizontal descriptive MRDB-claim: If a mechanism description D contains more horizontal details than D* about the ontic mechanism for phenomenon P, then D has more descriptive power than D* for P, all things being equal.

Horizontal Explanatory MRDB-claim: If a mechanism text T contains more horizontal details from the mechanism description that are explanatorily relevant for P vs. P’ than T*, then T has more explanatory power than T* for P vs. P’, all things being equal.

Does Craver and Kaplan’s reply to the MDB objection apply to both descriptive and explanatory horizontal completeness? As Craver and Kaplan’s account makes use of the contrastive formulation of the phenomenon it only captures the horizontal *explanatory* MRDB-claim (on the assumption that it is modified in line with our threefold distinction). The original mutual manipulability account, however, that defined constitutive relevance relative to an ontic phenomenon P can account for the horizontal descriptive MRDB-claim only. Hence, even though the mutual manipulability account as well as Craver’s and Kaplan’s solution to the MDB objection each on its own fail to account for the horizontal descriptive *and* the horizontal explanatory MRDB-claim, taken together they capture both.

However, as will become clear, this combinatory strategy does not work for vertical completeness. Again, vertical completeness norms can be defined for mechanism descriptions as well as for explanatory texts:

(*Vertical Completeness_{description}*) A mechanism description is vertically complete if and only if the description is (descriptively) horizontally complete at each mechanistic level.

(*Vertical Completeness_{text}*) An explanatory text is vertically complete if and only if the text is (explanatorily) horizontally complete at each mechanistic level.

Vertical completeness of mechanism descriptions is a descriptive matter. A mechanism description is vertically complete iff it describes the whole mechanism. This implies that a mechanism description is vertically complete if and only if it goes down to the fundamental level (if there is one). On the assumption of physicalism, everything supervenes on the fundamental physical level. As a consequence, every mechanistic hierarchy bottoms-out at the fundamental physical level. A vertically complete mechanism description will thus mention all acting entities that are constitutively relevant for a given phenomenon P at each level until down to the fundamental level. This gives us a further MRDB-claim:

Vertical Descriptive MRDB-claim: If a mechanism description D contains more vertical details than D* about the ontic mechanism for phenomenon P, then D has more descriptive power than D* for P, all things being equal.

In contrast to mechanism descriptions, explanatory texts do not always go down to the fundamental level. As a matter of fact, mechanistic explanations that can be found in the life sciences and other special sciences do not always mention, say, fundamental particles. Therefore, mechanistic explanatory text should be considered vertically complete even if they do not go down to the fundamental level. How to account for this? Craver and Kaplan briefly discuss this question in a footnote. They explain that their account

is consistent with the possibility that higher-level causes can ‘screen off’ lower-level components. Once the higher-level component is fixed, and the relevant background assumptions of the request for explanation are made explicit, differences among the lower-level components no longer make an additional difference to the explanandum phenomenon. In Woodward’s [(Woodward 2018)] terms, the lower-level parts can be irrelevant conditional on the behaviour of higher-level components. For example, once we know the sodium current across the membrane, the precise locations of the individual sodium ions are irrelevant. Likewise, it does not matter which of the thousands upon thousands of sodium channels do and do not open. This is why the total current equation can explain current in terms of conductance changes, while bracketing future knowledge of precisely how these changes are brought about. These low-level differences make no relevant difference once the higher-level behaviour is fixed. One consequence of this is that a complete explanation for a properly specified explanandum phenomenon need not (and typically does not) end in quarks (for further discussion, see Craver [2007a, b]). Which of the various multilevel relevance relationships happens to be screened off depends on the contrastive specification of the explanandum. (Craver and Kaplan 2020, n. 16)

This footnote suggests that Craver and Kaplan take Woodward’s 2018-account of *conditional irrelevance* to be a potential answer to the question of how to account for the fact that explanatory texts do usually not go down to the fundamental level. According to Woodward, a set of variables Y_k is *irrelevant for a variable E conditional on some additional variables X_i* iff (i) changes in the variables X_i are causally relevant to E, (ii) changes in the variables Y_k are causally relevant to E, and (iii) given the values of X_i are fixed, changes in Y_k make no difference to E (Woodward 2018). Applied to the present context, ‘causal relevance’ mentioned in (i) and (ii), has to be replaced by ‘constitutive relevance’. On the assumption that E is the phenomenon variable at level L_0 ; the X_i variables represent the mechanistic components at level

L₋₁, and the Y_k variables represent the mechanistic components at L₋₂, we end up with the following account of *conditional irrelevance of lower mechanistic levels*:

(*Conditional Irrelevance of Lower Mechanistic Levels*) A set of variables Y_k representing mechanistic components at level L_{-n} ($n > 1$) is irrelevant for a phenomenon E at level L₀ conditional on variables X_i representing mechanistic components at level L_{-n + m} ($n > m > 0$) iff:

- a) changes in the variables X_i are constitutively relevant to E ,
- b) changes in the variables Y_k are constitutively relevant to E , and
- c) given the values of X_i are fixed, changes in Y_k make no difference to E .

Conditions a) and b) are clearly satisfied for all components at all mechanistic levels (otherwise they would not be at lower levels as ‘being at a lower level’ is defined in terms of ‘being constitutively relevant’). The problem is that c) will be necessarily satisfied as well. The reason is that variables Y_k are also constitutively relevant for X_i . This follows from the definition of mechanistic levels: Y_k are at a lower level than X_i iff the former are components of the mechanism for the latter, which is the case iff the former are constitutively relevant for the latter (Craver 2007a, 189). As a consequence, each change in Y_k makes a difference to E only *via* a change in X_i . If a change in Y_k did not induce a change in X_i but still changed E , this would imply that the change in Y_k is *not* constitutively relevant for X_i . Thus, it would not be at a lower level than X_i . As a consequence, if Woodward’s account of conditional irrelevance was applied in the present context, necessarily, all levels lower than L₋₁ would turn out to be irrelevant for the phenomenon at L₀ conditional on the first lower level L₋₁. In other words, all lower levels (except for the first lower level) turn out to be always explanatorily irrelevant. Note that each lower level would be explanatorily relevant to the level directly above, but never to any higher level. Entities and activities at level L₋₂, for example, would be relevant to phenomena on L₋₁, but never to the original explanandum at L₀. Explanations would always stop at the first lower level. However, this may be too restrictive.⁶ We should allow for

⁶Note that the question we are interested in differs from the question that Woodward answers with his account of conditional irrelevance. Our question is ‘When is an explanation improved by going down the mechanistic hierarchy?’ Woodward’s question is ‘When is a higher-level explanation better than or as good as a fundamental level explanation?’ Woodward’s perspective differs from ours in the sense that in his context it is commonly assumed that (i) there are different explanations at different levels (whereas we assume that there is one explanation that can extend over multiple levels), and (ii) that lowest-level explanations are by default the preferred ones (due to considerations of causal closure and exclusion). Based on these considerations, the question arises whether higher-level explanations can at least sometimes be better or at least as good as lowest-level explanations. Here, Woodward provides a convincing answer: a given higher-level explanation is at least as good as the lowest-level explanation if the lowest-level explanation is irrelevant for the explanandum conditional on the higher-level explanation. In the mechanistic picture, however, explanation is a top-down matter: while the first lower-level is clearly explanatorily relevant for the phenomenon (say, the activity of the hippocampus is clearly explanatorily relevant for spatial memory), the lowest level is clearly not (say, the interactions between quarks is clearly irrelevant for the explanation of spatial memory). The question, then, is where in the mechanistic hierarchy explanatory relevance stops.

the possibility that *at least sometimes* going further down the mechanistic hierarchy improves an explanation.

The fact that Woodward's notion of conditional irrelevance makes lower levels always irrelevant shows that there is a further problem for Craver and Kaplan's reply to the MDB objection: either their account is too restrictive if they adopt Woodward's notion, or it is too permissive if they do not provide an alternative way of determining where explanatory texts bottom-out. As a consequence, they cannot account for vertical explanatory completeness, i.e., the vertical completeness norms for explanatory texts. Therefore, they are still confronted with what may be called the 'Vertical MDB-objection':

(Vertical MDB-objection) According to the new mechanists, an explanation of a higher-level phenomenon is always improved by adding more constitutively relevant lower-level details. However, explanations of higher-level phenomena do not usually go down to the fundamental level. Hence, the new mechanistic account of explanation fails.

Avoiding the vertical MDB-objection requires providing a criterion based on which one can decide how much lower-level detail is necessary to fully explain a phenomenon but just enough to not add explanatory irrelevant details. We have already seen that one possible solution fails: Woodward's account of conditional irrelevance seems to be too restrictive, as it implies that adding lower-level details never improves an explanation once the details about the first lower-level are fixed.

A second potential solution can be found in Machamer et al. (2000). They argue that explanations "bottom out" at levels that "are accepted as relatively fundamental or taken to be unproblematic for the purposes of a given scientist, research group, or field" (2000, 13). Even though this may be a good pragmatic answer to what it means for an explanation to be complete *relative to what someone finds interesting*, it does not provide an objective criterion of when an explanation is vertically complete (by "objective" we mean "not dependent on the interests of any individual" and "not pragmatic"). For example, a sociologist may be happy with explaining a social phenomenon just at the macro-level, and she may therefore declare it to be complete given her explanatory interests. But that does not mean that the explanation is complete in any objective way. The sociologist may happily admit that the explanation indeed is not (objectively) complete, but that the division of labour in science allows her to leave the further bits of the explanation to other researchers.

Such a pragmatic, non-objective criterion is problematic as a reply to the vertical MDB-objection. Craver and Kaplan and other mechanists accepting the ontic conception of explanation aim at developing an account of *objective* explanation (Craver and Kaplan 2020, 300). In order to be able to stick to this aim, the vertical MDB-objection should not be addressed by adding pragmatic constraints to the account. Indeed, if it turned out that the vertical MDB-objection can only be rejected if one were to integrate pragmatic considerations, this would constitute a rejection of the ontic conception of explanation, and in the end, admitting that the opponents were right: mechanists who aim at providing objective criteria for explanation in

fact would turn out to be committed to what the vertical MDB-objection ascribes to them.

Luckily, we think that an objective criterion can indeed be provided. This criterion can be inferred from the purpose of explanation of reaching “understanding of where, and sometimes how, to intervene and change the world for good or for ill” (Craver 2007a, 93) (see Sect. 17.4). Based on this, one can infer that an explanation is better than some other explanation if it identifies more crucial points of intervention, i.e., if it identifies where to intervene such that the intended phenomenon is produced in the most economic fashion (with minimal effort). On our account, mechanistic explanatory texts exhibit the differences between the description of the mechanism for the actual phenomenon and the description of the mechanism for the contrast phenomenon that is most similar to the description of the actual mechanism compared to the descriptions of all other possible mechanisms for the contrast phenomenon. The vertical completeness issue is resolved by extending this principle to the choice of the appropriate bottoming-out level. That is, in choosing the appropriate level to stop the explanation, we are attempting to find the crucial points of intervention. Crucial points of intervention are those, where we can find the most systematic and least disruptive way to transform the description of the actual mechanism to a description of the contrast phenomenon.

The choice of the appropriate bottoming-out level is assessed in an equivalent way to the choice between two or more competing ways to exhibit the contrast phenomenon from Sect. 17.5.1. The only difference is that instead of computing edit distances for complete mechanism descriptions of possible mechanisms, one finds edit distances from the mechanism description for P to the mechanism description for P' for each competing level. The appropriate bottoming-out level is the one with the smallest edit distance.

(Bottoming-Out in METs) Mechanistic explanatory texts bottom out at that level, for which the edit distance from P to P' is minimal in comparison to other available levels. Where there is a tie, lower bottoming-out levels are preferred.

The criterion proposed here has three interesting features. Firstly, in most cases it only gives us defeasible justification for the belief that our explanation of any particular contrast is vertically complete. This is because it is always possible that on a lower, so far unexplored level of mechanism, the differences between the actual mechanism description and the mechanism for the contrast phenomenon will be more systematic, thus allowing a shorter transformation procedure from the one to the other. Even if we find a level where the transformation procedure only has one step, it is possible, though unlikely, that a lower level will be found at which the transformation procedure also has just one step. In this situation, we think it uncontroversial that one should prefer the deeper explanation. In practice, though, such situations would be exceedingly rare.

Secondly, the criterion does not intrinsically favour either lower, or higher-level explanation. Rather, the appropriate level at which the explanation is complete is contingent on the result of empirical investigations. Further, vertical completeness of explanation may differ across phenomena, and across contrasts related to the

same phenomenon. This means that the criterion we propose is non-arbitrary, but intimately tied to explanatory practice.

Lastly, this criterion helps explain why mechanistic models in individual special sciences tend to bottom out at levels containing similar entities and activities specific for each discipline or sub-discipline. It can be hypothesized that the entities and activities at these levels contain crucial points of intervention for the contrastive explananda of interest to the sub-disciplines in question. For instance, even though the mechanism underlying certain depressive episodes is highly complex, it appears that serotonin mediated synapses in a number of brain circuits play a crucial role. On higher levels of mechanistic description, the contrast between a depressed episode and normal functioning must be accounted for by citing a number of disparate differences in many brain regions. But on a lower level, this contrast is accounted for by a higher number of highly systematic differences having to do with neurotransmitter concentrations. Other seemingly complex contrasts on higher levels of mechanism in the brain may turn out to be due to systematic differences in neurotransmitter concentration, secretion or inhibition. The research into these kinds of differences constitutes the discipline of psychopharmacology.

17.6 Conclusion

The aim of this paper was to find a satisfactory solution to the MDB-objection. We showed that the most promising extant account due to Craver and Kaplan (2020) introduces new problems, namely the Odd Ontology Problem, the Multiplication of Mechanisms Problem, and the Ontic Completeness Problem. Furthermore, that account is still incomplete, as it leaves open how explanatory relevance with respect to contrasts is to be determined. And even worse, it is still vulnerable to a version of the MDB-objection, i.e. the vertical MDB-Objection.

Our account builds on the foundational idea by Craver and Kaplan, and it resolves all five of these issues. The Odd Ontology and the Multiplication of Mechanisms problems are avoided because our threefold distinction between ontic phenomena, mechanism descriptions and mechanistic explanatory texts only introduces contrasts as a feature of mechanistic explanatory texts. Ontic phenomena are not contrastive, and there are no ontic mechanisms for every conceivable contrast. The Ontic Completeness problem is solved, because instead of formulating completeness norms for ontic mechanisms (Craver and Kaplan's SC) we provide completeness norms for both mechanistic descriptions and mechanistic explanatory texts.

Additionally, we provide criteria for explanatory relevance of mechanistic details relative to contrastive explananda. In our account, this means to determine the contents of mechanistic explanatory texts, which enables us to determine which constituents from the mechanism description should be cited to account for any particular contrast. Since in our account contrasts are not ontic, we can keep the original account of constitutive relevance as an account of the dependency relation between ontic mechanisms and ontic phenomena.

Finally, in contrast to Craver and Kaplan's account, our account avoids the vertical version of the MDB-objection. According to our proposal, mechanistic explanatory texts bottom out at those levels of the mechanistic hierarchy where the edit distance from P to P' cannot be decreased by going a level down. This level will contain the crucial points of intervention for turning phenomenon P into contrast phenomenon P'.

References

- Baetu, T. M. (2015). The completeness of mechanistic explanations. *Philosophy of Science*, 82, 775–786. <https://doi.org/10.1086/683279>.
- Batterman, R. W., & Rice, C. C. (2014). Minimal model explanations. *Philosophy of Science*, 81, 349–376. <https://doi.org/10.1086/676677>.
- Baumgartner, M., & Casini, L. (2017). An abductive theory of constitution. *Philosophy of Science*, 84, 214–233. <https://doi.org/10.1086/690716>.
- Baumgartner, M., & Gebharter, A. (2016). Constitutive relevance, mutual manipulability, and fat-handedness. *The British Journal for the Philosophy of Science*, 67, 731–756. <https://doi.org/10.1093/bjps/axv003>.
- Baumgartner, M., Casini, L., & Krickel, B. (2018). Horizontal Surgicality and mechanistic constitution. *Erkenntnis*. <https://doi.org/10.1007/s10670-018-0033-5>.
- Boone, W., & Piccinini, G. (2016). Mechanistic abstraction. *Philosophy of Science*, 83, 686–697. <https://doi.org/10.1086/687855>.
- Bunke, H., & Shearer, K. (1998). A graph distance metric based on the maximal common subgraph. *Pattern Recognition Letters*, 19, 255–259. North-Holland. [https://doi.org/10.1016/S0167-8655\(97\)00179-7](https://doi.org/10.1016/S0167-8655(97)00179-7).
- Chirimuuta, M. (2014). Minimal models and canonical neural computations: The distinctness of computational explanation in neuroscience. *Synthese*, 191, 127–153. Springer Netherlands. <https://doi.org/10.1007/s11229-013-0369-y>.
- Cohen, W. W., Ravikumar, P., & Fienberg, S. E. (2003). A comparison of string metrics for matching names and records. In *ACM international conference on Knowledge Discovery and Data Mining (KDD), Workshop on data cleaning, record linkage, and object consolidation, 2003*, 3:73–78.
- Craver, C. F. (2007a). *Explaining the brain: Mechanisms and the mosaic Unity of neuroscience*. New York: Oxford University Press.
- Craver, C. F. (2007b). Constitutive explanatory relevance. *Journal of Philosophical Research*, 32, 1–20. https://doi.org/10.5840/jpr_2007_4.
- Craver, C. F. (2016). The explanatory power of network models. *Philosophy of Science*, 83, 698–709. <https://doi.org/10.1086/687856>.
- Craver, C. F., & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology and Philosophy*, 22, 547–563. <https://doi.org/10.1007/s10539-006-9028-8>.
- Craver, C. F., & Kaplan, D. M. (2020). Are more details better? On the norms of completeness for mechanistic explanations. *The British Journal for the Philosophy of Science*, 71, 287–319. <https://doi.org/10.1093/bjps/axy015>.
- Damerau, F. J. (1964). A technique for computer detection and correction of spelling errors. *Communications of the ACM*, 7, 171–176. ACM. <https://doi.org/10.1145/363958.363994>.
- Dretske, F. I. (1972). Contrastive statements. *The Philosophical Review*, 81, 411. <https://doi.org/10.2307/2183886>.
- Glennan, S. (2017). *The new mechanical philosophy*. Oxford: Oxford University Press.

- Illari, P. M. K., & Williamson, J. (2011). Mechanisms are real and local. In *In Causality in the Sciences* (pp. 818–844). Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199574131.003.0038>.
- Illari, P. M. K., & Williamson, J. (2012). What is a mechanism? Thinking about mechanisms across the sciences. *European Journal for Philosophy of Science*, 2, 119–135. <https://doi.org/10.1007/s13194-011-0038-2>.
- Kaiser, M. I., & Crickel, B. (2017). The metaphysics of constitutive mechanistic phenomena. *The British Journal for the Philosophy of Science*, 68, 745–747. <https://doi.org/10.1093/bjps/axv058>.
- Kästner, L. (2017). *Philosophy of cognitive neuroscience: Causal explanations, mechanisms and empirical manipulations*. Berlin/Boston: De Gruyter.
- Krickel, B. (2018a). *The mechanical world* (Studies in brain and mind) (Vol. 13). Cham: Springer. <https://doi.org/10.1007/978-3-030-03629-4>.
- Krickel, B. (2018b). Saving the mutual manipulability account of constitutive relevance. *Studies in History and Philosophy of Science Part A*, 68, 58–67. <https://doi.org/10.1016/j.shpsa.2018.01.003>.
- Levenshtein, V. I. (1966). Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10, 707–710. <https://doi.org/10.1023/A:1022689900470>.
- Levy, A. (2014). What was Hodgkin and Huxley's achievement? *The British Journal for the Philosophy of Science*, 65, 469–492. <https://doi.org/10.1093/bjps/axs043>.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1–25.
- Milkowski, M. (2016). Explanatory completeness and idealization in large brain simulations: A mechanistic perspective. *Synthese*, 193, 1457–1478. <https://doi.org/10.1007/s11229-015-0731-3>.
- Monge, A., & Elkan, C. (1997). An efficient domain-independent algorithm for detecting approximately duplicate database records. In *The proceedings of the SIGMOD 1997 workshop on data mining and knowledge discovery*.
- Railton, P. A. (1980). *Explaining explanation: A realist account of scientific explanation and understanding*. Princeton University.
- Rice, C. (2015). Moving beyond causes: Optimality models and scientific explanation. *Nous*, 49, 589–615. John Wiley & Sons, Ltd (10.1111). <https://doi.org/10.1111/nous.12042>.
- Romero, F. (2015). Why there isn't inter-level causation in mechanisms. *Synthese*, 192, 3731–3755. <https://doi.org/10.1007/s11229-015-0718-0>.
- Wagner, R. A., & Fischer, M. J. (1974). The string-to-string correction problem. *Journal of the ACM*, 21, 168–173. <https://doi.org/10.1145/321796.321811>.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. New York: Oxford University Press.
- Woodward, J. (2013). Mechanistic explanation: Its scope and limits. *Aristotelian Society Supplementary Volume*, 87, 39–65. <https://doi.org/10.1111/j.1467-8349.2013.00219.x>.
- Woodward, J. (2018). Explanatory autonomy: The role of proportionality, stability, and conditional irrelevance. *Synthese*. Springer Netherlands: 1–29. <https://doi.org/10.1007/s11229-018-01998-6>.