# Chapter 9
# Reduction Beyond the Reduction Debate

Reduction reconciles diversity and directionality with strong unity. Diversity is descriptive, or conceptual in nature. Directionality is the directionality of explanatory dependence. Unity is cashed out in terms of identity. The direction of reduction depends upon features of the descriptions under which an object is presented by the expressions flanking the reduction predicate. These features of the descriptions under which an object is presented by the expressions flanking the reduction predicate can be captured by the notion of a property structure. This final chapter contains an outlook on the connection between reduction as conceived of here and the notion of a scientific level, physicalism, various forms of unification, grounding, and intervention. All these topics have been studied extensively; the present chapter thus merely suggests that the so explicated notion can be used to illuminate some aspects of these topics that clearly transcend the boundaries of the various reduction-debates in the philosophy of mind and the philosophy of science.

Section 9.1 introduces two notions of a scientific level and defines two notions of physicalism. Section 9.2 connects the explication of reduction to issues of unification, thereby completing the argument started in the previous chapter, Section 8.7, showing that the model of reduction proposed here is more fundamental than holistic notions of reduction. Section 9.3 suggests that there is a tight connection between reduction and grounding, and Sect. 9.4 characterizes one pragmatic dimension of reductive explanation: it gives access to specific intended interventions.

## 9.1 Reduction, Scientific Levels, and Physicalism

The explication of reduction proposed above sheds light onto the notion of a scientific level in a reductive hierarchy. Levels turn out to be conceptual levels, or levels of kinds of property structures. The physical level is distinct from the mental level in a reductive hierarchy because the former employs physical property structures, i.e. structures that present us with an object as physical, whereas the latter employs mental property structures that present us with an object as mental. This is

R. van Riel, *The Concept of Reduction*, Philosophical Studies Series 121,
DOI 10.1007/978-3-319-04162-9_9, 

the sort of levels a reductionist is committed to, not more and not less. We thus deliver some of the details needed to understand generic reduction-statements, such as 'mental properties reduce to physical properties' (see Sect. 5.6.2). To see how the notion of reduction proposed above connects to level-talk, consider the question of what physicalism consists in.

### *9.1.1 Physicalism*

What would the reductive physicalist want her notion of physicalism to achieve? It should (i) give an idea of how objects at different scientific levels differ, and (ii) give an idea of what the objects that constitute our world basically are. The framework proposed here suggests a disambiguation: In one sense, an object belongs to a scientific level insofar as it is presented under property-structures distinctive of that level. An object belongs to a scientific level in another sense insofar as it *can* be presented under property structures distinctive of that level. Physicalism is, then, the doctrine that any object is physical in the second sense, and that, in principle, everything reduces to physics. The property structures distinctive of physics are, according to physicalism, the most transparent with respect to the nature of objects that make up the actual world; based on these, we can explain all the rest.

To motivate this interpretation of two senses of physicalism, let me introduce a puzzle that is similar in spirit to the one that guided the discussion in Chap. 3. Within reductive hierarchies, we can distinguish different levels. The terms 'biological level', 'chemical level', 'psychological level' and so forth can be used to refer to levels within a reductive hierarchy (neglect sciences that clearly span levels for the moment). Moreover, within a reductive hierarchy, the biological level is clearly not the physical level. Similarly, *water, pain* and even *C-fiber stimulation* do not belong to the physical level in any reductive hierarchy they occur in. At the same time, one might want to hold that if everything is physical, then everything belongs to the physical level. This seems to be knowable *a priori*; or it seems to be an obvious condition on any appropriate explication of the thesis that everything is physical as well as of the notion of a physical level. Assume thirdly that physicalism might be true even in worlds with reductive hierarchies comprising more than the physical level. This seems equally plausible: Reductionism (based on the notion of reduction explicated above) should be compatible with physicalism. This requires there to be at least two levels – a reducing one and a reduced one. From these three assumptions, we get the following picture:

There is a possible reductive hierarchy in a physicalist world which is such that for some $x$, $x$ belongs to the, say, biological level of that hierarchy and, thus, not to the physical level of that hierarchy and is, thus, biological and not physical (since in a reductive hierarchy, it cannot be both at the same time). At the same time, since physicalism is true, $x$ is physical. So, doesn't it belong to the physical level? This pretty much looks like a contradiction.

This puzzle is of heuristic value, and the solution is easy to sketch: There are at least two different notions of *being physical* (or, more generally, of belonging to level *F*). The sense in which an object is not physical although it is identical to some object situated at the physical level, is this: It is not physical *under some property structure,* or *under some meaning,* but *physical under another.* So, being physical depends, in one sense, on being presented under a specific meaning (and being presented under a specific property structure the meaning gives access to). Thus, *being physical* would be a two place relation – an entity would be physical with respect to some property structures under which this entity is presented, and not with respect to other property structures determining this entity.[1] On this interpretation, predicates such as '_is physical', '_is biological' and the like are hyper-intensional: substitution of co-referential expressions may alter the truth-value.

The second notion is a notion of *being physical* (or *being an F*-entity, or *belonging to scientific level F*) that is independent of how an object is presented on a given occasion. I will come back to the distinction between two notions of being physical in a moment. Let us first consider what features of a property structure might determine which level it is associated with. I shall just outline a general framework, without trying to settle this issue finally, and I will dwell upon several examples.

### *9.1.2 Scientific Levels*

The idea is simple: different reductive levels are associated with a coherent pool of properties that form complex property structures at these levels. Any such coherent pool of properties forms the set of basic constituents of the property structures at the corresponding level. Thus, to which level a property structure belongs, or whether or not it spans levels, depends upon which properties are its basic constituents. Consider the following quote from Cussins:

> We perceive the world as beautiful or ugly, sweet or salty, happy or sad, brave or cowardly, intelligent or stupid; yet none of these properties figure in the world of the physical science. (Cussins 1992, 179)

If we replace 'in the world of' by 'as basic in the property structures of', we get an idea of how different properties might play different roles at different levels. Perceptual properties, or qualia-properties are basic in our everyday representations of the world. They are not basic in our physical representations of the world. Nevertheless, they may be physical in nature and, thus, be representable by a

[1] Note that this is in line with Fodor's conception of a kind, that is (i) tied to the notion of a law, and (ii) hence, at least intensional: Fodor suggests that 'it is a law that_' is not truth-functional (Fodor 1974, 109). This also seems to correspond to an idea suggested by Hempel in his (1969), where he argues that a purely ontological characterization of what it is to be physical, or biological, or the like, might fail.

property structure the basic constituents of which are properties that are basic in physical representations. Let me point to the fact that this assumption perfectly matches several ways of conceiving of levels which are implicit in the debate, thereby giving examples of what one might expect to be a *coherent pool of properties*:

(a) Early behaviorists argued that psychological terminology is analyzable or in some other sense conceptually reducible to behavioral terminology. If this is true, behavioral properties figure as basic constituents of property structures of psychology. In this respect, psychology would have to be distinguished from other sciences, in which other properties play a similar role.
(b) Modern functionalists conceive of the meaning of psychological terminology as being given by relations to other psychological terms and to behavioral terms. Again, basic constituents of such property structures will be behavioral properties (and some concepts of causation and an internal state). The difference to early behaviorists is not that we change the set of basic properties, but rather that we do not assume that there is a straightforward mapping function from psychological terms onto input–output relations without invoking sensitivity to different internal states which might influence each other. Thus, again, there is a coherent set of properties, which figure as basic in property structures employed by psychology according to the functionalist.
(c) Assume that qualia concepts cannot be conceptually reduced to behavioral concepts (see, for example, Papineau 2007). Then, there are genuine qualia properties that figure as basic in property structures under which qualia-terminology presents us with internal states. Let this level be the *phenomenal level*. Again, there is a coherent set of properties (namely: qualitative or phenomenal properties), which figure as basic in property structures characteristic of the phenomenal level. This is obviously not to say that these properties are not reducible. Rather, a description's feature of *being phenomenal* is described via the set of properties that are presented as basic in property structures of that level. Assuming that phenomenal properties are physical, a physical description of phenomenal states would present these properties under property structures that do not involve phenomenal properties *as basic constituents*; rather, they would determine phenomenal properties by giving their physical architecture.

I do not claim that these specific conceptions of scientific levels are correct. I merely claim that they perfectly match the assumption that there is one coherent set of properties which function as the basic constituents of the relevant property structures. But what does account for the differences between these pools of properties? This is a topic for another book. One might speculate that the relevant pool of properties depends upon the set of experimental or observational techniques that are applied investigating a specific set of phenomena – that would explain the occurrence of sciences that "branch" classical scientific levels as well as differences between the levels of the special sciences. Biochemistry consists in the study of chemical aspects of phenomena classically studied in biology, obviously with the application of techniques common in chemistry. Thus, intuitively, every science

comes with its basic kinds. Clearly, these basic kinds may gain special attention within a science and are, maybe, the science's primary target of investigation (this is, basically, Causey's point claiming that there is a sub-set of the relevant ontology a theory or science is *primarily* concerned with, although it may account for more phenomena (Causey 1972, 1977)). However, they do not form the set of kinds a science is concerned with *only*; rather, they figure as basic in property structures associated with this science. The kind *reproduction* or *being an organism*, or related properties, figure among the basic properties of some parts of biology, and it is this fact that makes these property structures the structures that are associated with biology. Behavioral properties and properties concerning "internal" relations form the basic properties of folk psychology, if functionalism is correct. Briefly: to characterize what makes a science belong to the reductive level it belongs to is to characterize it in terms of the properties that figure as basic in the property structures this science gives access to. In this sense, we can distinguish kinds of property structures. We can thus make sense of generic reduction statements (see Sect. 5.6.2). Let me repeat: *Mental* properties reduce to *physical* properties; For every property *x*, if x is determined by some mental property structure *z*, then there is a property *y*, a physical property structure *z**, such that *x*, when presented under *z* reduces to *y*, when presented under *z**.

*Prima facie*, we have reason to assume that the difference between levels stems from a difference in properties that function as basic constituents in the property structures characteristic of this level. In this sense, being an *F*-entity (with '*F*' standing for some term like 'biological', 'chemical' or the like) depends upon which basic properties form the property structure under which the object is presented. In the light of this, we can now turn back to the two notions of a scientific level and easily solve the "puzzle":

The Puzzle's Solution: Two Ways of Being Physical (at least)

Being an *F*-entity, i.e. being a physical, biological, neural ... entity, in one sense depends upon being presented by elements of the set of properties characteristic of the *F*–level:

*(Def. Being an F-entity)* *An entity a is an F-entity iff the meaning of 'a' gives access to a property structure of level F.*

This sense of being situated at a level is tied to characterizing something as being biological, physical and the like within *reductive hierarchies*. Thus, '_is physical' is, in this sense, hyper-intensional: It yields a truth only if the referent of the argument is conceptualized appropriately *by* the argument. Now, consider the second sense in which something might be *physical*. According to this sense of being physical, biological and so forth, we get the following criterion:

*(Condition – Being an F-entity*)* *An entity is an F-entity independent of how it is presented.*

This captures notions of being physical that rely on purely ontological descriptions of being physical. It is pertinent in explications of the notion of being physical

which refer to physics (Smart 1978; Braddon-Mitchell and Jackson 1996; Chalmers 1996), and those which describe being physical via paradigmatic physical objects and their properties (Meehl and Sellars 1956) (thus, the "second sense" of 'being physical' is, maybe, a family of different notions of being physical, which are tied together by taking things to be physical in some sense which is independent of how these things are presented). However, we can, relying on the hyper-intensional sense of belonging to a level, as described by *(Being an F-entity)*, give a criterion for belonging to a level in the second sense:

*(Def. Being an F-entity*)* *An entity is an F-entity iff it is determined by a property structure the basic constituents are genuine to the level F.*

Accordingly, an entity would be physical iff it is determined by a property structure the basic constituents of which are typical of the physical level (neglecting logical constituents). That is: It is physical if, in principle, it could be presented under a property structure characteristic of the physical level. And it is mental if it, say, makes for a phenomenal or behavioral difference.

This picture enables us to *explain what identifying some high-level property with a physical property consists in.* Some property P (= property P*) is at a higher level in a reductive hierarchy than P* iff the property structure associated with 'P' belongs to some higher level within that reductive hierarchy than the property structure associated with 'P*'. However, P (and, thus P*) might be a physical property. There is, then, no puzzle of how some property can be mental and, thus, in one sense non-physical, and at the same time physical. This also bears upon the relation between ontological and reductive hierarchies: Ontological hierarchies, whatever these are, do not perfectly match reductive ones. The former are insensitive to how its objects are presented, the latter are not.

This distinction is of utmost importance when we try to distinguish between purely ontological issues on the one hand (concerning, say, physicalism, monism and so forth) and issues that are sensitive to ways of presenting or describing an object (like issues of reduction). The concept of reduction is, thus, not the only notion in the field that is sensitive to meanings; in one of their uses, predicates such as '_is a physical kind' and '_is a psychological entity' are used to form sentences the truth of which is sensitive to semantic features other than designation or reference of their arguments. Given that the model of property structures captures the relevant aspect of meaning, levels are, at least partly, individuated by property structures.[2] Levels present things in different ways because they present things

[2]Here is an application of this idea. Under the assumption that property structures fully capture the cognitive role of meaning, so that there is no property structure associated with two cognitive roles, this can be used to describe one condition on the truth of *a priori* physicalism. A priori physicalism states that *from a statement of all physical facts, all truths about the mental are knowable on a priori grounds.* (Beckermann 2009, 162) For a priori physicalism to be true, there should be some a priori path from the property structures of the physical level to property structures at higher levels (if property structures are the primary targets of a priori operations). Then, a priori physicalism is

under different properties. Theories or other representational items at distinct levels may be concerned with the same set of entities and kinds and yet be distinct.

This suggestion also points to how reduction relates to various aspects of *unification*. In general, this strategy of tying qualifications such as '_is physical', or '_is mental' to presentation under sorts of property structures enables us to account for descriptive *pluralism* in an ontologically *homogeneous* world, or: to maintain ontological *unity* while allowing for descriptive *pluralism*. The hints towards unification in the reduction literature are rather vague. The next section argues that given the explication offered here, a number of different sorts of unification are to be expected for cases of theory reductions (we thus not only provide further reason for the claim that the explication proposed here is fruitful; at the same time, we resume the issue of Chap. 8: the argument that the account proposed here is more fundamental than holistic accounts).

## 9.2 Reduction and Unification

The topic of unification in science forms a debate in its own right, especially since it has been described as being closely tied to explanation in general (Friedman 1974; Kitcher 1981, 1989). Scientific unity is discussed under several distinct foci, and one might easily get the impression that the debate is even more complex than the debate on reduction.[3] Here, we cannot give anything that comes even close to a full-blown account of how reduction relates to scientific unity in all these respects. Thus, I shall just point to how it relates to ontological, epistemological and explanatory unification. The three remaining conditions for directionality in holistic approaches (2–4) to reduction can thereby be shown to depend on reduction as explicated here.

---

true *only if a complete description of the world under physical property structures conceptually implies, or gives a priori access to, the set of truths about mental facts presented under mental property structures.*

It has been assumed that what makes a given property structure belong to the physical level within a reductive hierarchy is that it is constituted by elements of a specific set of properties which are presented as basic. Thus, a physical description of everything would be a description that gives access to everything *under these basic properties (and logical operations on these basic properties) only*. Now, higher-level descriptions give us entities under *their* specific property structures. In this specific sense, there has to be an a priori route from the physical level to higher-levels, if a priori physicalism is true.

[3]Critics of the unity-program often neglect the fact that this program did not consist in attempts to actually carry out unifications. Recall the Kemeny and Oppenheim-remark on a definition of reduction: such a definition would become "hopelessly complex" (1956, 13), if we were to capture all the relevant aspects of actual or possible reductions. In these and similar writings, it is quite common to appeal to an ideal demand rather than arguing that such general reductions could actually be carried out. At least, this point did not receive the attention it deserves. Although unification might not be a goal one should actually try to accomplish in any important sense, it might nevertheless be more than a regulative idea– it captures one aspect of the idea that some representation is more appropriate, or more basic than another.

### 9.2.1 Reduction and Ontological Unification

The notion of reduction, and the two sorts of notions of belonging to a scientific level *F* (a physical level, a biological level and the like) described in Sect. 9.1 give a precise idea of how to maintain *ontological unity* while enabling us to account for *descriptive pluralism* in a very specific sense of these terms. Ontological unity for a set of entities (kinds, states, objects and so forth) holds iff every element of this set can be described under one coherent set of property structures. Descriptive pluralism for a set of entities holds iff every element of this set can be described under at least two distinct coherent sets of property structures. The answer to the question of how this is possible is straightforward: Descriptive pluralism is required by reduction – each reductive level comes with its own property structures under which entities are presented at that level. With regard to their descriptions on given occasions, entities may be physical, biological and so forth and at the same time belong to one specific ontological category, like being physical. Different ways of being an *F*-entity allow for there to be different descriptions which give us just one (ontologically construed, or description-independent) kind of entities.

Given that ontological unity and descriptive pluralism hold at one world, *reduction* is the relation that ties together descriptive pluralism and ontological unity. In virtue of reducing to the lowest descriptive level (in virtue of being presentable under the property structures characteristic of this level), entities are ontologically situated at this level. *F*-entities (say, biological entities) reduce to *F**-entities (say, chemical entities) iff they are identical to *F**-entities and the *F*-descriptions or -property structures are adequately tied to *F**-descriptions or -property structures. In virtue of being presented under different property structures at different levels, the reduction task is an interesting one: We have to find the relevant cross-level identities. Once this is done, we may hope for reductive explanation and, thus for *ontological unification*. Picking up the topic of Chap. 8: Again, one of the features invoked in holistic accounts of the directionality of reduction, namely, ontological unification, is to be expected based on, and explained by the framework proposed here. Ontological unification occurs when we acquire knowledge about identities of objects that seemed to be situated at different ontological levels, which turn out to be mere reductive, and, hence, descriptive levels. The idea that reduction is tied to unification as a procedure of showing that there is unity where there seemingly was not has important *epistemological* implications.

### 9.2.2 Reduction and Epistemological Unification

Let us now turn to the epistemological value of the procedure of unification in science. *Carrying out a reduction goes together with epistemological unification.* The explication given above enables us to see *how it is possible that there are true sciences (or, more liberally: appropriate descriptions and terms) that are seemingly*

*unconnected, or are situated at different levels, although they are compatible with strict monism.* This is in need of further clarification. Epistemological unification, I suggest, consists in showing that what was presented by two different sorts of property structures (or meanings) can be *presented* by one. This captures the interpretation given for Nagel's characterization of reduction as an *assimilation of traits*.

Assume that two theories $T_R$ and $T_R$* both reduce to some theory $T_B$. (The same point could be made using just one reducing theory. However, using two makes the conceptual and the ontological aspect more transparent.) Assume, moreover, that $T_R$ and $T_R$* seemed 'unconnected' (Nagel 1961, 359). One criterion for seeming unconnected is being such that there is no (relevant) a priori or known metaphysical connection between statements of $T_R$ and $T_R$*. If there were no a priori connections, then, according to the model outlined above, both theories were associated with relevantly different property structures. Now, introducing bridge-principles which state identities and subsuming what was described by $T_R$ and $T_R$* under *one common conceptual framework*, (presenting their respective fields under property structures which use the same resources of basic properties), namely that of $T_B$, is clearly some sort of *unification of the conceptual resources*. What seemed unconnected before is now shown to be connected – to be representable under one common conceptual framework. Thereby, the conceptual aspect of unification is captured; since knowledge is mediated by conceptual presentation, we could describe this sort of unification as epistemological: One unified body of knowledge (in terms of $T_B$) captures what involved a body of knowledge in terms of $T_R$ and a body of knowledge in terms of $T_R$* before.

This conceptual or epistemological aspect goes together with the ontological aspect of unification described in the previous section. By the reduction of $T_R$ and $T_R$* to $T_B$ it is shown that the difference in *meaning* of the languages of $T_R$ and $T_R$* (and $T_B$) bears no relevant ontological implications. There just is one ontological sort of entities, although there are different sorts of entities with respect to how they are presented. Thus, the intuition that reduction is tied to unification, or that reductions go together with unification, can be given a precise explication. Two aspects of unification, namely, a conceptual aspect, which bases an epistemological, and an ontological aspect can thereby easily be explained using the model outlined here. Again, a holistic criterion turns out to be derivative on the explication proposed above. In addition, reduction amounts to what one may want to label *explanatory* unification.

### *9.2.3 Reduction and Explanatory Unification*

The explanatory priority of the property structures given in the *explanans* of a reductive explanation corresponds to the degree to which these properties are responsible for how the reduced object behaves. One interesting form of explanatory

unification can easily be characterized as follows, building on an idea already hinted at in Sect. 8.7.2. We already know that in a theory reduction, the range of explanation is expanded. Let me repeat. Water reduces to $H_2O$. Assume, again, that freezing reduces to the formation of lattice structures of molecules. Why does water freeze below 0 °C under normal conditions? It does so *because* $H_2O$ molecules form (stable) lattice structures below 0 °C under normal conditions. Given our every day theory of freezing events, we can explain why, say, in Northern Europe, lakes sometimes freeze during winter. They do so *because* water freezes below 0 °C, and during winter, temperature sometimes falls below 0 °C in Northern Europe. If so, by the transitivity of explanation, in Northern Europe, lakes sometimes freeze during winter *because* sums of $H_2O$ molecules form lattice structures when the temperature falls below 0 °C (under normal circumstances). This follows from the explication of reduction and the transitivity of explanation:

If *F*-ness reduces to *G*-ness, then for every *x* that is *F*, *x* is *F because x* is *G*. Assume that '*F**' is a term of our reduced theory. Then, by the transitivity of explanation, for every truth of the form '*x* is *F** because *x* is *F*', there is a corresponding truth '*x* is *F** because *x* is *G*'. Thus, transitivity plus reduction ensures that the range of explanation is preserved. Reduction + transitivity + the idea that a constitutive property structure is always more informative as to the nature of the object (for example, by tracking mereological dependence) than the higher-level property structure ensures that the range of explanation is expanded. Now, applying the notion of a scientific level, we arrive at explanatory unification: If property structures at a level come in a uniform format, then the idea that reduction guarantees *explanatory uniformity* comes for free. By a reduction, we get a greater range of our conceptually coherent *explanatory resources*. One coherent set of property structures enables us to explain what we used to explain using a different set of property structures. In this sense, we get explanatory unification, which is, somewhat like epistemological unification, a side effect of conceptual unification. What we gain by reductions (in this sense) is *transparency* with respect to the contribution of properties to what we want to explain. Reduction goes together with conceptual (and, hence, explanatory and epistemological) unification as well as with ontological unification. Hence, another holistic criterion turns out to be derivative on the explication proposed in the first part of this book.

As a concluding remark on unification, it should be noted that a great deal of the debate on unification is concerned with the scientific *value* of unification (an idea which has been attacked by Darden and Maull (1977), Cartwright (1983, 1989), and Dupré (1993)). Let me just point to the fact that the interpretation of the connection between reduction and unification I proposed by no means goes together with the idea that unification is the *only* goal of science. However, it *shows* that it is at least one goal (which might even be incompatible with others).

Recall the distinctions suggested in the introduction concerning the connection between realist/anti-realist and conservative/eliminative versions of reductionism, that correspond to different interpretations of the alleged truth or relevance of high-level theories or sciences (see Table 9.1).

**Table 9.1** Monisms and reductionisms

| Monisms and reductionisms | Anti-realism | Realism |
|---|---|---|
| Eliminativism | Monism-1: Eliminativism/replacementism | Monism-2: Unificationism |
| Conservatism | Monism-3: Epistemic/pragmatic non-reductivism | Monism-4: Conservative reductionism |

In the introduction, I mentioned that there is no conservatism or eliminativism (building on normative constraints on reductions) *tout court*. The book deals with realist versions of reduction, i.e. versions that assume that in principle, high-level sciences or theories or descriptions may be true, or that the entities purportedly referred to by these theories and descriptions possibly exist. For these cases, we should explicitly distinguish different respects in which eliminations might be appropriate or misguided. Elimination (in the sense of getting rid of a theory or a way of talking or conceptualizing) is, in one specific respect, a goal worth aiming at, because it will force us to reject less transparent theories and replace them by more transparent ones (or move to reductive levels the property structures of which are explanatorily prior to the property structures associated with higher levels). Once we give primacy to what is explanatorily prior, we clearly gain knowledge about the structure of our world, if it is monistic. As pointed out above, once we get a reduction, we get unification, and as long as we regard unification as a goal in itself, this might give further reason to believe that reductions are relevant and worth aiming at.

On the other hand, there are many respects in which elimination is *not* a goal worth aiming at. Or, as Elliot Sober put it: "The reductionist claim that lower-level explanations are *always* better and the antireductionist claim that they are *always* worse are both mistaken." (Sober 1999, 560, italics in the original). One problem is, clearly, that once descriptions become too complex, we would be utterly lost. In this respect, some reductions might not be worth aiming at. Accordingly, (and again) we should distinguish different respects in which levels of description are appropriate. Some (possible) reductions would supposedly contradict criteria of pragmatic appropriateness (see also van Gulick 1992; Friedman 1982, 17; Wimsatt 1976).

One cognitive role of high-level descriptions (which shows that they are indispensable in one respect) can nicely be illustrated using a quote from Lennon (1992):

> The distinct physical groundings of a given intentional kind will be *shapeless* from the viewpoint of physical theory. [...] Our classifications into intentional kinds provides a pattern of conceptualization which yields empirical generalizations and conditional dependencies. These patterns of dependency are at the heart of causal explanation, and could not be captured without the employment of intentional notions. (Lennon 1992, 226)

What sort of modality is involved in 'could not be captured without the employment of intentional notions' (which are not to be confused with kinds, as this passage suggests)? A modest interpretation would suggest that it is *us* who could not capture the relevant dependencies without using these (or similar) notions, due to

our cognitive shape. This interpretation is compatible with reduction as the notion is spelled out here: For us the complex physical descriptions of high-level kinds might look shapeless. But this is not necessarily so. If we were able to grasp the property structures revealed by the relevant low-level descriptions, the apparent reason to assume that these descriptions look shapeless might disappear. Therefore, we should not assume that if one representational item is more appropriate with respect to, say, practical or didactical purposes, it is also a more appropriate device in the sense that it represents the relevant domain in a more transparent way. What reduces in virtue of its representational structure does not necessarily "reduce" when questions of accessibility are addressed. More precisely, possibly implicit aspects of *normativity* in the notion of reduction are limited and context-sensitive. There isn't just one goal. Gaining transparency and unification form one goal, systematization (Kemeny and Oppenheim 1956; Kitcher 1981, 1989) is another. In principle, it would be great to achieve these goals all at once. But due to our limited cognitive capacities, this hope seems to be futile. So, we should be flexible in our use of notions of *eliminative* or *conservative reduction* (in the pragmatic sense): Reduction is clearly not the *only* goal of science. Let us now turn to another issue, that of connecting reduction to what has recently been discussed under the title 'grounding' in metaphysics.

## 9.3 Reduction and Grounding

This section briefly connects the book's topic to a discussion that emerged in a completely different area of philosophy – a topic closely related to ontological dependence. The goal of this section is modest; it argues that, under a very general conception of what theories of grounding are after, reduction is a promising candidate for counting as a *variant* of grounding.

Classical accounts of non-causal and, usually, non-conceptual dependence relations have tried to come up with explications in terms of metaphysical modality (see, for example, Simons 1987). Recently, this view has been attacked by, for example, Correia (2005), and Fine (2005).[4] There are two sorts of criticism of this approach that are relevant for the present discussion. One criticism concerns what Fine dubbed "modal-mania" in metaphysics (2005, 9). Roughly, it is argued that modal notions cannot be used to define a relevantly *directional* relation, or, in the syntactic sense, a directional expression. Thus, focus shifted towards explanatory concepts. And this is where the second criticism enters: Fine (2001) argued that 'because' and similar expressions do not express a *relation*. This latter point will be ignored for the present discussion, although it should be noted that if it is correct, then the *reduction*-relation as described here is not that of explanatory dependence. The reduction relation is

[4]For an overview, see Correia (2008).

then conceptually tied to non-relational explanatory concepts.[5] Moreover, I will not question the first criticism. Rather, I want to briefly address the question of how reduction relates to grounding, if grounding is not to be cashed out in counterfactual terminology.

In a nutshell, it will be suggested that one promising job-description of grounding or grounding-relations is this: grounding reconciles *diversity* and *directionality* with *committal unity*. Depending on what we take to be the primary relata of grounding, grounding is *committal* with respect to the grounded and the grounding entity; that is: either with respect to the truth of what is grounded and what grounds or to the existence of what is grounded and what grounds. Put differently, grounding, as understood here, is distinct from elimination. If *that p* is grounded in *that q*, (and it is true *that q*) then *that p* is true. Insofar as strong unity is a form of *committal* unity, reduction, as explicated here, is a special case of grounding.[6]

Jenkins (2011) has proposed a very general argument for the claim that reduction is a special case of grounding. She suggests that we might *want*, in principle, our notion of grounding to be neutral with respect to the question of whether or not if *a* is grounded in *b*, *b* and *a* are identical.[7] This section is supposed to further motivate this idea, reflecting on the relation between the two job-descriptions for 'ground' and 'reduction'. To begin with, recall the *prima facie* plausible worry that immediately arises, which is based on an idea already hinted at (Chap. 3):

If *a* reduces to *b*, then $a = b$. Therefore, reduction has the formal properties of identity: reflexivity, symmetry, and transitivity. Following this line of thought (*cf.* Trogdon 2013, Section 6), one might go on stating that grounding is *irreflexive*. Therefore, reduction is not a variant of grounding, even if it does not exhibit the structural features of identity (as has been argued above). Reduction is not irreflexive either; so, if grounding is, then reduction is not a variant of the grounding relation. But is grounding irreflexive? Obviously, we should assume that if an instance of '*a* is grounded in *b*' expresses a truth, then the corresponding instance of '*b* is not grounded in *a*' expresses a truth as well. However, this does not, without further ado, imply that the grounding relation is irreflexive; Jenkins (2011), for example, argues that it is not irreflexive, reflecting on cases of reduction. How can we decide? Isn't this a matter of convention, or, again, stipulation? To assess the question of whether or not reduction is a form of grounding, we should first get clear about what this question actually consists in.

---

[5]This raises further subtle issues: If expressions such as 'because' do not express relations, one may ask whether or not expressions such as 'grounds' or 'ontologically depends upon' are themselves appropriately explicated in terms of relational concepts (*cf.* Fine 2001). Reduction itself may be inadequately described in relational terminology; however, a decision on that issue would take us too far from our present target. If explanatory dependence is not relational, then the explication proposed above should be rephrased accordingly.

[6]Note that the use of 'reduction' deviates from the way it is introduced by Fine in his 2001.

[7]Her discussion of questions regarding asymmetry and reflexivity arrive at conclusions similar to the ones suggested above.

### 9.3.1 Reduction as Grounding – The First Argument

'_grounds_' and '_reduces to_' are more or less technical expressions a characterization of which should take the form of an explication (Chap. 2). Correspondingly, the question of whether or not a technical concept *F* is a variant of a technical concept *G*, could be rephrased as follows: According to the best explications of *F* and *G*, is *F* a variant of *G*? Obviously, an explication of the concept of *grounding* cannot be given here. We thus cannot give a straightforward answer to the question of how reduction relates to grounding. Two alternatives suggest themselves: We could go through the different candidate explications for *grounding* and check for each of them whether or not it covers reduction as well, and then give a conditional answer: given the appropriateness of the account, reduction is/is not a case of grounding. One problem with this procedure is that it does not provide the resources for an answer to the following question: *Should* grounding be regarded as covering cases of reduction?[8]

Any fruitful characterization of grounding should, like any other explication, comply with a job-description. The question of whether or not reduction is a form of grounding can then be construed as a question concerning the respective concepts' job-descriptions: Is fitting the job-description for reduction a way of fitting the job-description for grounding? If the above job-description is correct, reduction is a form of grounding: Whenever we have reconciliation of directionality and diversity with *strong* unity, we have reconciliation of directionality and diversity with *committal* unity. Thus, we now have to motivate the idea that the proposed job-description for grounding forms at least one appropriate way to characterize the idea underlying grounding talk.

It will prove useful to think of the predicate '_grounds_' and its variants as expressions whose definitions have to correspond to a number of uses of *other* expressions that are taken from samples of ordinary, scientific, and philosophical discourse, which, under an appropriate interpretation, inform the job-description. Schaffer suggests the following list: '(i) the entity and its singleton, (ii) the Swiss cheese and its holes, (iii) natural features and moral features, (iv) sparse properties and abundant properties, and (v) truthmakers and truths.' (Schaffer 2009, 375)

These samples comply with the job description already mentioned, in slogan form: *Grounding reconciles diversity and directionality with committal unity* (that it is an atemporal or synchronic relation should go without saying). This rough

[8]To answer this question, we could try to check explications of grounding on the market; however, grounding is often regarded as primitive (Schaffer 2009; Fine 2001). But even if there were "reductive" characterization, like in the case of reduction, one should not treat the characterizations as stipulations; rather, they are *proposals*. One might even feel tempted to revise a given proposal in the light of the fact that they cover, or do not cover a notion of reduction (this seems to be Jenkin's move). Thus, showing that reduction is covered by alleged characterizations of grounding is not enough to account for the relation between reduction, as construed here, and the notion of grounding the candidate characterizations are after. The alternative strategy is maybe less straightforward, but it seems to perfectly suit the present purpose.

description can be further motivated reflecting on further examples; the expressions 'grounds' and 'is grounded in' can be employed to capture metaphysical relations that are expressed by or correspond to metaphysical explanations, and their use is informed by the use of other bits of philosophical jargon, such as 'constitutes', or 'brings about', and more technical expressions such '_is an element of_', and, maybe, '_is an epiphenomenon of_'. Here are a few examples:

(i) The field exists *because* the individual plants exist.
(ii) My headache occurs *because* a number of physiological events occurred.
(iii) This statue is beautiful *because* it has specific physical properties.
(iv) The heart pumps blood *by* regular muscle-contraction.
(v) This statue is *constituted by* a chunk of matter.
(vi) My mental processes *do not exist over and above* neural processes.
(vii) The mayor's announcing you husband and wife *brings it about* that you become husband and wife.
(viii) My current mental state *is an epiphenomenon* (given a non-causal interpretation of 'epiphenomenon') *of* my current brain-state.

I submit that none of the statements in the group (i)–(iv) expresses a truth unless the explanantia and the explananda are true and/or the entities the explananda are about exist (one may describe this as follows: theories of grounding buy into the factivity of explanations); moreover, none of the members of the group (v)–(viii) is true unless the relevant entities exist. Thus, if grounding is supposed to capture these fragments of discourse, it is *committal*. Moreover, we get unity by directionality: The cases suggest that there is a fundamental level that fixes the rest. In this respect, (i)–(viii) are similar to the examples on Schaffer's list. To the extent that this already gives a very general notion of grounding, reduction, as defined above, is a particular form of grounding. It is a form of non-causal, non-conceptual dependence, that is possibly cashed out in terms of explanation and it is committal in that it does not rely on a purely eliminative stance towards the grounded/reduced objects or propositions. Grounding reconciles diversity and directionality with committal unity. Reduction does so, too, in a specific way: Diversity is conceptual, and unity comes in the form of identity. If these job-descriptions are appropriate, we have a very general reason to assume that reduction is a case of grounding. Jenkins proposes an additional argument for the thesis that metaphysical dependence or grounding is (best construed as being) neutral with respect to identity.

### *9.3.2 Reduction as Grounding – The Second Argument*

Jenkins' argument seems to rely on strategic considerations (Jenkins 2011, 269f. – here adapted to the present list of examples): A reductive interpretation of examples similar to those on the list discussed in the previous section seems at least conceivable. If so, and if we want our notion of dependence or grounding to cover

these cases, we would be well advised to describe dependence or grounding as being neutral with respect to reduction. Thus, one may *want* to describe grounding in a neutral way, so that by embracing a grounding claim, one is not committed to the non-identity of the *relata*. After all, the cases decide whether or not we are faced with identity, so the conceptual apparatus we employ to approach these cases should remain neutral in this respect (Jenkins 2011, 270). But why not adopt a different (terminological) strategy? Why not employ 'reduction' and 'dependence' or 'grounding' in competing ways, so that reduction and grounding turn out to be mutually exclusive? Then, instead of phrasing the issue in terms of whether or not a particular case of grounding is a case of reduction, we would have to assess the question of whether or not a given case is a case of grounding *or* reduction. Apparently, we do not increase or decrease the expressive power of our theoretical apparatus by opting for this alternative strategy. As a thesis about how we should construct our conceptual apparatus, the point is primarily terminological in nature. But terminological decisions may be strategically relevant. In the present case, having a notion that captures the similarities may form a strategic advantage – we do not lose sight of the possibility that we may be convinced that some *x* is grounded in *y* (because the case bears some signature features of grounding), and still, *x* may turn out to reduce to *y*, without this having any relevant effect on our conviction that *x* is grounded in *y. Prima facie*, reductive interpretations are conceivable for various candidates of grounding claims: On some accounts, the field will turn out to be identical to (the sum of) the individual plants (i), just like my headache will turn out to be identical to the relevant physiological events (ii). The latter has been offered as a paradigm case for identity based reduction (Nagel 1961, 366). Cases similar to (i) seem to fit into a reductionist picture as well (though some may hesitate to assume that the field is nothing but the individual plants): A given amount of water reduces to the (sum of) $H_2O$ molecules, which constitute it, and it exists *because* the latter exist. (iii) seems to form a paradigm case for reduction (from the reductionist's perspective). Similarly, as has been extensively argued in Chap. 4, 'by' explanations can be given a reductive interpretation: A heart's pumping blood *might be identical to* rhythmic muscle-contraction in an appropriate environment. And at least *prima facie*, the way the term 'constitution' is employed may have instances where *what is constituted is identical to what constitutes it*: Early type-identity theorists took (complete) constitution-relations to *be* identity-relations (Place 1956, 1960; Smart 1959). According to this interpretation, our mental life is constituted by part of our physical life; at the same time, it just *is* this part of our physical life. Similarly, it goes without saying that the idea that mental processes *do not exist over and above* (and are, therefore, grounded in) neural processes has a straightforward reductionist interpretation. There may be uses in which 'brings about' cannot be coherently interpreted reductively, like in the truth-maker case alluded to above. But why shouldn't a reductionist about phenomenal properties happily admit that pains are *brought about by* (and are, hence, grounded in) C-fiber firings, although pains *are nothing but* C-fiber firings? What about more technical philosophical conceptions, such as *being an epiphenomenon of*, various forms of *dependence* or *being an element of*? Reductive interpretations of epiphenomena, sets and their elements

as well as particularized properties and their bearers seem rather counter-intuitive. For epiphenomenalism, which is introduced as an anti-reductionist position, such an interpretation cannot be given without changing the subject. Some nominalists, however, famously argued that sets can be reduced to non-abstract objects (as, for example, sketched in Lewis 1970); should we, then, give a definition of grounding which is such that no nominalist can consistently hold that Sokrates *grounds* {Sokrates}? Again, it might be a good idea to remain neutral on whether or not our use of 'grounding', when used to capture a relation between a set and its elements, commits us to the assumption that a set is distinct from its elements. In general, it seems to be a promising strategy to describe grounding in a way that leaves open the question of whether or not the "*relata*" are possibly identical. Issues of identity should be settled in the light of a study of the *relata*, rather than by a conceptual decision. Thus, Jenkins seems to be right; her argument that paradigmatic cases of grounding relations can be given a reductive interpretation is persuasive. In the light of a comparison between the respective job-descriptions for grounding and reduction, this seems to be quite natural.

So, *prima facie*, a number of, but not all cases that seem to be adequately describable in terms of grounding are conceptually compatible with or even perfectly match a reductionist interpretation. If so, then grounding comes in two forms, in a reductionist and in a non-reductionist form. The (more general) grounding relation is then *underdetermined* with respect to the issue of identity.

Let me point to one alternative job description for grounding, a description which would clearly rule out a reductive interpretation. If the study of grounding is the study of *dependencies among different metaphysical layers of reality*, where talk of 'layers' does not allow for cross-layer identity, then reduction is not a case of grounding (and what merely appears to be *grounded* may turn out to *reduce* instead). Moreover, some theorists might want to argue that whenever we have diversity, directionality, and committal unity, we do not have identity; this is a substantive issue, and not a terminological one. I hope that the present book's first part has shed light on how we can conceive of a coherent explication of reduction that turns out to be compatible with diversity, directionality, and committal unity based on identity. On a liberal understanding, the grounding relation (or the way philosophers employ the term 'ground') is more encompassing: It reconciles diversity and directionality with metaphysically committal unity, either by identification *or* by dependence.[9]

The grounding debate is one of the last in which true armchair philosophy seems to flourish. Let us now turn to a final application of the explication of reduction, an application to interventionist models of causation and dependence. This debate comprises two aspects, at least. First, it aims at a counterfactual

[9]However, whereas there may be an interesting *a priori* route to metaphysical theories about the relation between a cheese and its wholes, particularized properties and their bearers, natural properties and moral or aesthetic properties, and sets and their elements, there is no such route to a theory about the relation between the mental and the physical; this seems to be a purely *empirical* matter.

definition of causation in terms of intervention. Secondly, it has been referred to in order to illuminate a pragmatic aspect of the scientific relevance of mechanistic explanations (Craver 2006). In the next and final section, I will briefly argue that reductive explanation does, once achieved, form a perfect epistemic basis for interventions.

## 9.4 Reduction and Intervention

Some interventions on a dependence-base result in a change in the dependent object. The following brief reflection on the notion of an intervention will enable us to adequately describe the connection between reduction and the pragmatically relevant aspect of interventions.

The literature on the pragmatics as well as the theoretical relevance of interventions in the philosophy of science focused on causal and mechanistic contexts, where it is taken for granted that the benefits of interventions stem from the (theoretically and pragmatically relevant) fact that when intervening on an object (or a variable in a representation of a causal system), we thereby change objects (or values of variables) that causally or mechanistically depend upon the object we intervened upon.[10]

In the debate on mechanisms, reference to interventions is frequent when it comes to distinguishing between mechanistic explanations and mere models of mechanisms that do not adequately capture the target entity's organization (*cf.* Craver 2006). Thus, it is introduced as an epistemological criterion to distinguish between mechanistic (or correct mechanistic) and non-mechanistic (or incorrect mechanistic) explanations. Moreover, it plays a role when it comes to judging issues regarding top-down "causation". The idea is that for there to be top-down "causes", we should find an intervention on the overall-mechanism that somehow alters its organization (cf. Craver 2007). These discussions do not directly bear upon reductive explanation, since they seem to presuppose that there are *two* objects involved in the intervention, one that is dependent and another one that forms the dependence base. However, necessarily, if $x$ reduces to $y$, then an intervention on $x$ is an intervention on $y$. This is trivially true, because (i) for $x$ to reduce to $y$, $x = y$, and (ii) 'intervenes on' is extensional. Note that this changes once we take interventions to be interventions on variables in a representation of a system: The intervention on one variable results in a change in another variable does not ensure that this

[10]The vast literature on this topic mainly refers to Pearl (2000) and Woodward (2003). I assume that the reader is familiar with the basics. The task of finding definitions for causation will not be our concern here. Similarly, the role interventions may play in the discovery of reductions will not be addressed. If they do, the causal graphs that represent a system should be neutral with respect to which dependence relations they represent (see also footnote 11).

dependence mimics a dependence between two objects in the target system – the might just be *one* represented by *both*.

Thus, we have to look elsewhere for the connection between reduction and intervention. One interesting thought, pertinent, for example, in Woodward (2003) and Craver, is this: A *deeper* explanation enables us, at least in principle, to answer questions of the following form: 'What if things had been different?' This is supposed to capture the following connection: If we are familiar with the causal architecture, or the nature, or a constitutive property structure of an object, we can manipulate things *in an informed way*. This suggests a tight connection between reductive explanation and the pragmatic relevance of scientific interventions. In order to make the connection between interventions and reduction transparent, we need an idea of how interventions may become description relative. Once this idea is made precise, the connection between reduction and interventions can easily be explicated.

Interventions are description-dependent insofar as they are *intended* actions that are based on specific explicit representations of the object we intervene on. Necessarily, any intervention on water is an intervention on $H_2O$, and *vice versa*. However, an *intended* intervention can, intuitively, be construed as an intervention that is relative to a certain representation of the object intervened on. Let an *intended intervention* be an intervention on $x$ that is planned with respect to a representation by a property structure that presents us with $x$. The hyper-intensionality is here inherited from the hyper-intensionality of intentionality. If I intend to intervene upon the temperature of a piece of ice, I do not thereby intend to intervene upon the kinetic energy and the lattice structures of the sum of corresponding $H_2O$ molecules. Intended interventions are always description-, and, hence, property-structure-relative.[11]

Now, the relevance of reductive explanations for scientific interventions becomes transparent: Once we *learn* that an object $x$, when presented under a property structure *PS1*, reduces to $x$, when presented under a property structure *PS2*, we get conceptual access to a number of intended interventions with respect to *PS2*. Assume that we learn that, say, the occurrence of fever is a complex event, involving the release of PGE2 which, acting on the hypothalamus, causes the body temperature to raise. Acquiring this knowledge, and acquiring the knowledge that fever occurs in virtue of the occurrence of this mechanism, we are in a position

[11]Note that interventionist models of causation talk about representation-dependence as well. However, it is noteworthy that here, representation dependence concerns the different causes that are presented in a causal graph, rather than the different ways one and the same cause may be presented in a causal graph. A possible problem for interventionist definitions of causation is that they cannot properly distinguish between causation, on the one hand, and other dependence relations on the other. Unless we stipulate that in a causal graph, any variable has to represent another entity, and that synchronic dependence relations are excluded by fiat, they seem to be highly problematic (similar worries have been risen in the discussion of how interventionism relates to the causal exclusion argument and, hence, to epiphenomenalism; *cf.* Baumgartner (2010) for a critical examination of this discussion).

to plan intended interventions on the mechanism. Note that this does not come for free with the truth of an identity statement: Assume that fever = the medical sign which is most often mistaken for a disease (in the actual world). Knowledge of this identity statement does not enable us to intervene on fever. Why is this so? One may suggest that this is so *because* being a fever does not *depend* on being the medical sign that is most often mistaken for a disease (in the actual world). Reduction is a form of dependence. Knowledge of reduction statements puts one in the position to plan intended interventions one might not have access to otherwise. Interestingly, such knowledge also enables us to get access to intended interventions with respect to the constitutive property structure, which may have an effect on the properties under which the same entity is presented at the reduced level. Assume that the folk-concept of fever corresponds to a property structure that represents fever in terms of its symptoms. Once we come to know what fever is on a physiological level – in terms of a physiological property-structure –, we learn how to intervene upon it by designing effective drugs. These drugs enable us to intervene upon the symptoms by intentionally intervening upon the physiology. Thus, reductive explanation connects to intentional intervention as follows: If a subject *s* knows that *x*, when presented under *PS1*, reduces to *x*, when presented under *PS2*, then *s* is (at least under favorable circumstances)[12] in a position to intentionally intervene upon *x* with respect to *PS1* by intentionally intervening on *x* with respect to *PS2*.

Since at the reducing level, property structures give access to *constitutive* property structures, knowledge of a truth about a reduction relation increases, at least in principle, the set of possible intended interventions on that object. Thus, reductive explanations play an important part in the pragmatically relevant outcome of scientific enterprises. Even a partially correct reductive explanation would suffice: As long as part of the low-level representation – say, fever under some of its physiological properties – is correct, an intervention along the lines just described may succeed. Correspondingly, a partial reduction (i.e. where a whole is not fully re-described in terms of its constituents) will be sufficient. This is not sufficient for full-blown reduction. But it may be sufficient for full-blown intended interventions.

## 9.5 Conclusion

In this chapter, it was shown that the concept explicated in the first part of this book can be fruitfully applied to topics outside the reduction debate, and the argument proposed in Chap. 8, that the account proposed here is more fundamental than its

[12]That is: The intervention must in principle be possible and the subject needs the relevant means and skills. Knowledge itself does not guarantee the possibility of an intentional intervention. However, it is required for this kind of intervention. Moreover, one should, maybe, not stress the notion of knowledge to much in this context. Even true believe might be sufficient.

holistic rivals, has been continued. The answer to the guiding question of this second part of the book has thus been completed:

| | |
|---|---|
| Q2: | How can this explication be further motivated? |
| ... | |
| Th. 7: | It is as committal as and more fundamental than rival explications. |
| Th. 8: | It sheds light on closely related issues, such as reduction and unification, pragmatic benefits of reduction, and notions of scientific levels. |

This chapter continued the debate on the relation between holistic approaches to reduction (Chap. 8), with an eye on the relation between reduction and forms of unification, and it applied the explication proposed here to issues such as reduction and grounding, reduction and intervention, reduction and physicalism, and reduction and the notion of scientific levels. One notion of a scientific level is tied to modes of presentations, in a way that can, again, be illuminated referring to property structures. Unification is to be expected in reductions: We subsume apparently different phenomena under one coherent conceptual scheme, achieve ontological unification and epistemic as well as explanatory unification. Reduction is a cognate, or a version of grounding, and knowledge of true reduction statements has an impact on the set of intended interventions we have access to.

## Bibliography


Baumgartner, Michael. 2010. Interventionism and epiphenomenalism. *Canadian Journal of Philosophy* 40: 359–384.

Beckermann, Ansgar. 2009. What is property physicalism? In *The Oxford handbook of philosophy of mind*, ed. Sven Walter, Brian P. McLaughlin, and Ansgar Beckermann, 152–172. Oxford: Oxford University Press.

Braddon-Mitchell, David, and Frank Jackson. 1996. *Philosophy of mind and cognition*. Oxford: Blackwell Publishing.

Cartwright, Nancy. 1983. *How the laws of physics lie*. Oxford: Oxford University Press.

Cartwright, Nancy. 1989. *Nature's capacities and their measurement*. Oxford: Oxford University Press.

Causey, Robert. 1972. Attribute identities in microreductions. *Journal of Philosophy* 64: 407–422.

Causey, Robert. 1977. *Unity of science*. Dordrecht: Reidel.

Chalmers, David. 1996. *The conscious mind*. Oxford: Oxford University Press.

Correia, Fabrice. 2005. *Existential dependence and cognate notions*. München: Philosophia.

Correia, Fabrice. 2008. Ontological dependence. *Philosophy Compass* 3(5): 1013–1032.

Craver, Carl. 2006. When mechanistic models explain. *Synthese* 153: 355–376.

Craver, Carl. 2007. *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford: Oxford University Press/Clarendon Press.

Cussins, Adrian. 1992. The limitations of pluralism. In *Reduction, explanation, and realism*, ed. Kathleen Lennon and David Charles, 179–224. Oxford: Oxford University Press.

Darden, Lindley, and Nancy Maull. 1977. Interfield theories. *Philosophy of Science* 44: 43–64.

Dupré, John. 1993. *The disorder of things. Metaphysical foundations of the disunity of science*. Cambridge, MA: Harvard University Press.

Fine, Kit. 2001. The question of realism. *Philosophers' Imprint* 1: 1–30.

Fine, Kit. 2005. *Modality and tense. Philosophical papers*. Oxford: Oxford University Press.

Fodor, Jerry. 1974. Special sciences: Or the disunity of science as a working hypothesis. *Synthese* 28: 97–115.

Friedman, Michael. 1974. Explanation and scientific understanding. *Journal of Philosophy* 71: 5–19.

Friedman, Kenneth. 1982. Is intertheoretic reduction feasible? *The British Journal for the Philosophy of Science* 33: 17–40.

Hempel, Carl G. 1969. Reduction: Ontological and linguistic facets. In *Philosophy, science, and method: Essays in honor of Ernest Nagel*, ed. Morton White, Sidney Morgenbesser, and Patrick Suppes, 179–199. New York: St. Martin's Press.

Jenkins, Carrie S. 2011. Is metaphysical dependence irreflexive? *The Monist* 94: 267–276.

Kemeny, John G., and Paul Oppenheim. 1956. On reduction. *Philosophical Studies* 7: 6–19.

Kitcher, Philip. 1981. Explanatory unification. *Philosophy of Science* 48: 507–531.

Kitcher, Philip. 1989. Explanatory unification and the causal structure of the world. In *Scientific explanation*, ed. Wesley Salmon and Philip Kitcher, 410–505. Minneapolis: University of Minnesota Press.

Lennon, Kathleen. 1992. Reduction, causality, and normativity. In *Reduction, explanation, and realism*, ed. Kathleen Lennon and David Charles, 225–238. Oxford: Clarendon.

Lewis, David. 1970. Nominalistic set theory. *Noûs* 4: 225–240.

Meehl, Paul E., and Wilfrid Sellars. 1956. The concept of emergence. In *The foundations of science and the concepts of psychology and psychoanalysis*, ed. Michael Scriven and Herbert Feigl, 239–252. Minneapolis: University of Minnesota Press.

Nagel, Ernest. 1961. *The structure of science. Problems in the logic of explanation*. New York: Harcourt, Brace & World, Inc.

Papineau, David. 2007. Phenomenal and perceptual concepts. In *Phenomenal concepts and phenomenal knowledge*, ed. Torin Alter and Sven Walter, 111–144. Oxford: Oxford University Press.

Pearl, Judea. 2000. *Causality: Models, reasoning, and inference*. Cambridge: Cambridge University Press.

Place, Ullin. 1956. Is consciousness a brain process. *British Journal of Psychology* 47: 44–50.

Place, Ullin. 1960. Materialism as a scientific hypothesis. *Philosophical Review* 69: 101–104.

Schaffer, Jonathan. 2009. On what grounds what. In *Metametaphysics: New essays on the foundations of ontology*, ed. David Chalmers, David Manley, and Ryan Wasserman, 347–383. Oxford: Oxford University Press.

Simons, Peter. 1987. *Parts*. Oxford: Clarendon.

Smart, John J.C. 1959. Sensations and brain processes. *Philosophical Review* 68: 141–156.

Smart, John J.C. 1978. The content of physicalism. *Philosophical Quarterly* 28: 339–341.

Sober, Elliot. 1999. The multiple realizability argument against reductionism. *Philosophy of Science* 66: 542–564.

Trogdon, Kelly. 2013. An introduction to grounding. In *Varieties of dependence*, ed. Miguel Hoeltje, Benjamin Schnieder, and Alexander Steinberg, 97–122. München: Philosophia Verlag.

Van Gulick, Robert N. 1992. Nonreductive materialism and the nature of intertheoretical constraint. In *Emergence or reduction? essays on the prospects of non-reductive physicalism*, ed. Ansgar Beckermann, Hans Flohr, and Jaegwon Kim, 157–179. Berlin: de Gruyter.

Wimsatt, William C. 1976. Reductive explanation: A functional account. In *PSA: Proceedings of the Biennial meeting of the Philosophy of Science Association, 1974*, ed. G. Pearce, A.C. Michalos, C.A. Hooker, and R.S. Cohen, 671–710. Dordrecht: Reidel.

Woodward, James. 2003. *Making things happen. A theory of causal explanation*. Oxford: Oxford University Press.