Reductionism today

Michael Esfeld
University of Lausanne, Department of Philosophy
Michael-Andreas.Esfeld@unil.ch

(published in Maria Carla Galavotti, Elisabeth Nemeth and Friedrich Stadler (ed.):
European philosophy of science – Philosophy of science in Europe and the Viennese heritage.
Cham: Springer 2013, pp. 89-101)

Abstract

The paper outlines the main argument for ontological reductionism in today's discussion, claims that ontological and epistemological reductionism (theory reduction) stand or fall together and finally sketches out how today's most widespread form of reduction, namely functional reduction, can be developed into a fully-fledged theory reduction, thus taking up the programme of the Vienna circle in today's philosophy.

1. Introduction

The Vienna circle is associated with a strong reductionist programme in the philosophy of science, seeking to reduce all scientific theories to theories formulated in a physical vocabulary and to reconstruct the statements formulated in that latter vocabulary on the basis of protocol sentences and logics and mathematics. In particular, Carnap's *Logischer Aufbau der Welt* (1928) is evidence of this strong reductionist programme. This reductionist programme was considerably weakened when the heritage of the Vienna circle survived in the logical empiricism in the United States in the 1950s and the 1960s. Seminal works such as the Oppenheim and Putnam (1958) and Nagel (1961, in particular ch. 11) still pursue reductionism in the sense of the reduction of the theories of the special sciences to fundamental physical theories, but there no longer is the intention to reconstruct the vocabulary of physics on the basis of protocol sentences and logics and mathematics.

With the demise of logical empiricism, reductionism became unfashionable. However, during the last two decades or so, reductionism has come back on the agenda, but the focus is on ontology rather than on epistemology (theory reduction), since there always have been and still are strong arguments in favour of ontological reductionism. The thesis of this paper is that one cannot have a coherent ontological reductionism without epistemological reductionism, that is, establishing the principled possibility of reducing the theories of the special sciences to basic physical theories. In that sense, the spirit of the Vienna circle is still with us.

In the next section, I outline the main contemporary argument for ontological reductionism. I then explain why this argument is incomplete unless ontological reductionism is joint with theory reduction and sketch out how theory reduction can be defended in view of the objections that have been raised against it since the demise of logical empiricism (section 3). The material for this paper is based on the book Esfeld and Sachse (2011), which elaborates on the case for a comprehensive, but *conservative reductionism*. In this paper, I take for granted that reductionism is conservative instead of eliminativist: its aim is to show how the entities to which the special sciences are committed can exist in a physical world and how the theories of the special sciences can be true and can incorporate laws and explanations in such a world instead of having merely a heuristic or pragmatic value.

2. Ontological reductionism

Suppose that it is possible to define a basic physical domain of the world unambiguously: let it consist in all and only those physical properties that occur at space-time points. In order to obtain a complete microphysical description of the world, one would thus have to quantify over all space-time points and specify which physical properties occur at these points. Suppose now that the whole microphysical domain of the world is duplicated (see Figure 1). In other words, an operation takes place that duplicates the whole space-time including all and only those physical properties that are instantiated at space-time points. The world *w** thus created hence is microphysically identical with the real world *w*. Does *w** contain all that there is in *w*, that is to say, all the organisms, all the biological, psychological, social, etc. properties that there are in *w*, including the Vienna circle, the conference in Vienna in December 2011 and the book in which this paper is published? In other words, is *w** a duplicate *simpliciter* of *w*?

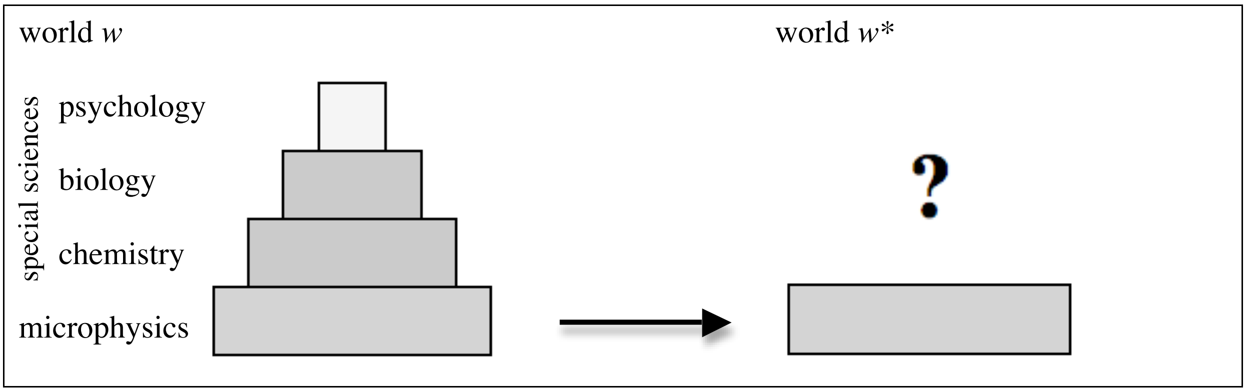

Figure 1: the world w* *on the right is an exact and complete microphysical duplicate of the world* w *on the left. Does* w* *contain everything that there is in* w*?*

There are good reasons to answer this question in the affirmative. We know that all objects that exist in the real world have developed from microphysical objects and are composed exclusively of microphysical objects. There can hence be no objects that exist in *w*, but that are absent in *w**. However, do complex macroscopic objects possess in *w** all the qualitative properties that they possess in *w*? In other words, is the operation that consists in projecting the whole domain of microphysical properties from *w* to *w** sufficient to guarantee that all the biological, psychological, social, etc. properties that there are in *w* exist also in *w**? Note that there is no question of a deterministic dynamics here: we stipulate that all the microphysical properties in the *whole* space-time of *w* be copied to *w**. The issue of what the development of the world *w* in time is like, whether it is deterministic or not, has therefore no bearing on this question.

If a certain biological, psychological, social, etc. property existed in *w*, but missed in *w**, we would search for a reason for this difference. That search would take us beyond the domains of biological, psychological, social, etc. properties: according to all what we know about the world, it is not possible that a duplicate of the world lacks only one phenotypic property – say, the yellow colour of the petals of a certain individual plant, these petals being red instead of yellow in *w** – without there being also a genetical or an environmental difference between the duplicate of the world (*w**) and the world (*w*). However, if there is such a difference, there also is some molecular and consequently some microphysical difference between *w** and *w* as well. Hence, in this case, *w** would not be an exact microphysical duplicate of *w*, but would differ from *w* in some microphysical detail.

By the same token, according to all what we know about the world, it is not possible that a duplicate of the world lacks only one psychological property – say, the thought of Barack Obama on 13 February 2012 that winning the presidential election in November 2012 will not be an easy matter. If this psychological property were absent in *w**, there would be further psychological differences between Barack Obama in *w** and Barack Obama in *w*, since any thought is linked up with further thoughts, as well as with emotions and finally with actions. Consequently, there would be some neurobiological difference or other between *w** and *w* in the brain state of Barack Obama at the indicated time as well as some behavioural difference and thus some molecular and finally some microphysical difference. Again, according to all what we know about the world, it is not possible that a duplicate of the world lacks only one economic property – say, that the Dow Jones Index slightly rises on 13 February 2012. If this economic property were absent in *w** – in other words, if the Dow Jones Index developed in a different manner in *w** on that day –, there would be some difference in the intentional attitudes and actions of persons between *w** and *w*, and thus some neurobiological and behavioural difference and hence finally some molecular including some microphysical difference. Consequently, in this case, *w** would again not be an exact microphysical duplicate of *w*.

We can sum up this reasoning in the following manner: there is no biological, psychological, social or economic difference without there also being a microphysical difference. In other words, everything that there is in the world globally supervenes on the microphysical domain. All biological, psychological, social or economic properties that are instantiated in the world have some effects (make a difference to the world), and they can have their effects only by also having physical effects down to microphysical effects.

However, in making this statement, we face a serious problem: for any physical change, there is a complete physical cause (insofar as there is a cause at all). Any physical change comes under physical laws, and these laws contain only physical variables. Even if these laws are not deterministic, but only probabilistic, they indicate the complete probabilities for the physical change in question. No biological, psychological, social or economic variables can influence the probabilities for the occurrence of certain physical changes without these probabilities being at the same time completely fixed by physical variables. This causal completeness of physics is known since the advent of Newtonian mechanics and employed since Leibniz' *Monadology* (1714, § 80) to refute dualist interactionism. We can illustrate this problem in the following figure (see Kim 1998, ch. 2, and 2005, ch. 2, for an elaborate argument):

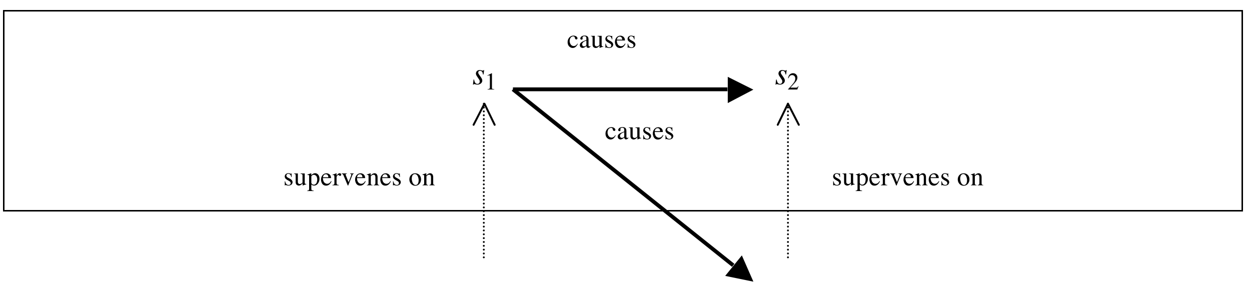

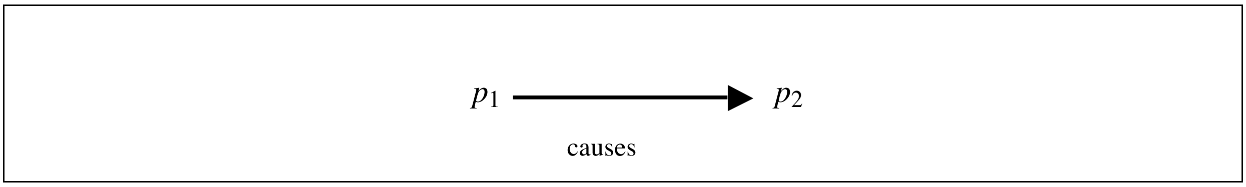

Figure 2: The domain of the properties with which the special sciences deal – the domain of biological, psychological, social, economic properties, etc. – globally supervenes on the domain of physical properties. A property token s_1 *of a special science causes another property token* s_2 *of a special science and thereby causes also a physical property token* p_2 *(the supervenience base for* s_2*). However,* p_2 *also has a complete physical cause* p_1.

There are exactly two possibilities to solve this problem if one admits global supervenience. The first possibility consists in maintaining that the properties (in the sense of property tokens) with which the special sciences deal are not identical with physical properties (that is, $s_1 \neq p_1$ and $s_2 \neq p_2$). Consequently, the properties of the special sciences *overdetermine* all their effects. For all physical effects of properties of the special sciences there is a complete physical cause. Furthermore, the supervenience of the domain of the properties of the special sciences on the domain of physical properties implies that there are sufficient physical conditions for the occurrence of any property token of a special science in the world. In other words, if s_2 supervenes on p_2, then p_2 is a sufficient condition for the existence of s_2: given p_2, s_2 cannot fail to exist.

However, it is not the case that physical properties also overdetermine the effects in the domain of the special sciences. The situation is not symmetrical. There are purely physical effects for which there are no causes in the domain of the special sciences. But there is nothing in the domain of the special sciences for which there are not completely physical sufficient conditions. Consequently, it is only the supposed causal efficacy of the properties of the special sciences that results in the mentioned overdetermination. The conclusion therefore is the following one: a situation (or a possible world) in which the properties of the special sciences do not cause anything and thus are epiphenomenal would be indiscernible from the situation depicted in figure 2.

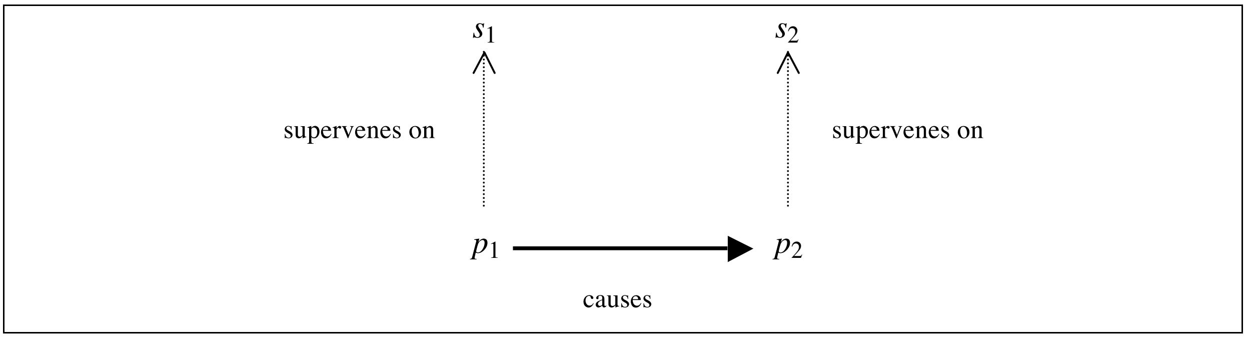

Figure 3: The properties of the special sciences being epiphenomenal.

Overdetermination is hence not a solution of the above mentioned problem that is in the position to establish the causal efficacy of the properties with which the special sciences deal.

This result is independent of the theory of causation that one favours (for a contrary view see Ladyman 2008). For all physical property tokens, there are complete physical causes (insofar as there are causes at all). The laws in which types of physical properties figure are strict laws. The laws of the special sciences, by contrast, are never strict. Even if one endorses a theory of causation according to which causal relations are completely captured by counterfactual propositions, there is the mentioned failure of symmetry, since the laws of nature figure prominently among the truth conditions of the counterfactual propositions in question. That is to say, the counterfactual propositions that are about causal relations among physical property tokens have a privileged status, since they are backed up by strict laws. There are of course also true counterfactual propositions linking supervenient property tokens with subsequent physical property tokens. However, there is no argument visible why these propositions should express a causal relationship (see Esfeld 2010 for an elaboration and Harbecke 2011 for a counter-argument).

The other possibility to solve the above mentioned problem consists in maintaining that the properties in the sense of the property tokens with which the special sciences deal are *identical with* physical properties. More precisely, each property token in the domain of a special science is identical with a configuration of physical property tokens. Supervenience is compatible with identity: the relationship of supervenience does not exclude that all the properties in the domain of supervening properties are identical with properties in the domain of the supervenience base. Instead of the three arrows of causation drawn in figure 2 above, there hence is only one causal relation; but there are different descriptions of the property tokens standing in the causal relation in question:

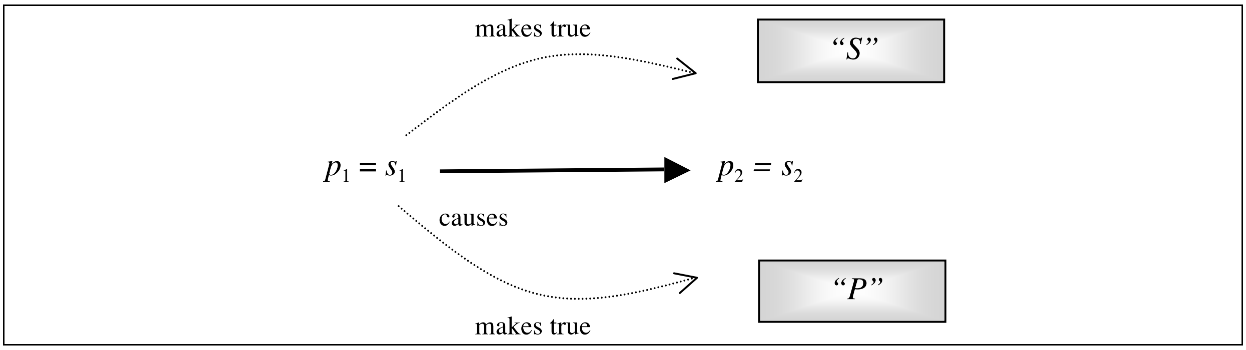

Figure 4: The property tokens of the special sciences are identical with physical property tokens. One and the same property tokens make true descriptions of different types.

However, it is often objected that if one takes the properties with which the special sciences deal to be identical with physical properties, then instead of vindicating the causal efficacy of the former, one *de facto* eliminates them, retaining *only* physical properties in one's ontology. But eliminativism is distinct from the identity theory: identity is a logical relation that is symmetric. If all the properties with which the special sciences deal are identical with physical properties, then some physical properties *are* properties of the special sciences. In general, if all *A*s are identical with *B*s, then both *A*s and *B*s exist, and some *B*s are identical with *A*s. Furthermore, it does not make sense to ask whether a given object in the domain of a special science brings about certain effects in virtue of its physical properties or in virtue of its special sciences' properties, since both are the same on the identity theory. In general, if the property of being A is the same as the property of being B, then all the effects that an object brings about qua being A are the effects that it brings about qua being B, and *vice versa*.

The identity claim as such does hence not provoke any eliminativism worry. The point at issue behind the eliminativism objection raised against the identity theory is the question how the properties with which the special sciences deal can be identical with physical properties. That question is well-taken. Even if the identity of biological, psychological, social or economic properties with physical properties in the sense of token identity is a fact, it is not a brute fact, but we need a theory that explains how it is possible for that identity to obtain. Functionalism is such a theory. If the properties with which the special sciences deal are functional properties, their essence is a causal role, that is, consists in bringing about certain effects given normal conditions. If the physical properties are also causal properties (that is, if their essence is or includes the disposition to bring about certain effects), then the causal role of certain complex configurations of physical properties can, given normal conditions, be identical with the causal role that is the essence of certain properties of the special sciences (see Esfeld and Sachse 2011, ch. 2, for an elaborate argument for this view).

In sum, if one accepts the premises of global supervenience and causal completeness of the physical domain, then one can build up a strong argument for *ontological reductionism*, that is, for the view that all the properties that are instantiated in the world are either physical property tokens or identical with (complex) configurations of physical property tokens.

3. Theory reduction

Ontological reductionism as set out in the previous section can also be accepted by some people who see themselves as physicalists, but who reject reductionism, thus endorsing the position that is known as non-reductive physicalism. Even Fodor (1974), in his famous argument against the unity of science as proposed by Oppenheim and Putnam (1958), can be read as accepting token identity, that is, the claim that every token of a special sciences' property is identical with some (complex) physical token. What Fodor (1974) – and non-reductive physicalists in general – reject, is type identity: there is no identity between types of the special sciences – such as biological or psychological types – and physical types. The reason is multiple realization: tokens coming under one and the same type of a special science can come under widely different physical types. However, this position provokes the following objection: insofar as there is only a reduction of tokens, but not of types, the problem remains whether the properties on which the special sciences focus cause anything qua biological, or psychological properties, etc.

Let us briefly consider the background of that objection: Donald Davidson (1970) claims in his famous paper "Mental events" that mental events are identical with physical events. More precisely, all events admit a physical description, and some events admit also a mental description. It is not possible to reduce the mental to a physical description. This position is known as anomalous monism – monism, because all events are physical, anomalous, because there is no nomological connection between the mental and the physical description of events, which would enable a reduction of the former to the latter. This position is widely recognized to fail due to the following objection: it cannot show that events cause anything insofar as they are mental events. Fred Dretske highlights this problem by conceiving the following example: the voice of a soprano singer causes a thin glass to shatter. This effect occurs in virtue of the amplitude and frequency of the sounds. The meaning of the sounds is irrelevant to this effect (Dretske 1989, pp. 1-2). The same applies to events insofar as they are mental in Davidson's anomalous monism according to a widely recognized objection (see the papers in Heil and Mele 1993).

The ontological reductionism sketched out in the preceding section differs from Davidson's anomalous monism in that what is identical are not events, but property tokens. Nonetheless, Paul Noordhof (1998) objects to David Robb (1997) that in the same way as it is reasonable to ask whether a Davidsonian mental event causes anything qua mental, it is reasonable to ask whether a mental property token causes anything qua being a token of a mental type. Robb (2001) retorts that if identity is applied to those entities in virtue of which an object or event causes something, namely property tokens, it makes no sense to raise the qua-question for these entities, since they are already the most fine-grained ones.

Even though that reply is correct, there remains a problem. Let us assume for the sake of the argument that all that exists in the world are particulars (objects having property tokens); types then are concepts that seize salient similarities among the property tokens that objects have, whereby such salient similarities can amount to natural kinds. Multiple realization then is the epistemological fact that tokens coming under one single concept of the special sciences often come under different physical concepts. The concepts of the special sciences and the corresponding physical concepts differ not only in meaning, but they are also not coextensive. Consequently, it is not possible to reduce the concepts of the special sciences to physical concepts.

However, in that case, the problem that haunts Davidson and that Noordhof raises against Robb reappears: it has to be possible to relate the different descriptions in a systematic, reductive manner, if they are descriptions that are made true by one and the same token in the world and if each of them is to have a scientific quality (that is, to provide for law-like generalizations that are projectible, support counterfactuals, yield causal explanations, etc.). Otherwise, it could not be vindicated that these descriptions are about the *same* entities in the fine-grained sense of the same tokens instead of being about *different* properties that objects have. Consequently, the ontological reductionism set out in the preceding section cannot stand on its own. *Reductionism cannot be had in a piecemeal way*. Ontological and epistemological reductionism (theory reduction) stand together, or fall together.

But how can a theory reduction be possible given multiple realization? Nagelian reduction (Nagel 1961, ch. 11) has been superseded by functional reduction (see notably Lewis 1994 and Kim 1998, ch. 4, 2005, ch. 4 & 5). Let T_1 be a theory of a special science and T_2 be a physical theory that covers the domain of objects with which the special science in question deals. Functional reduction then proceeds in three steps:

1) One defines the property types in the domain of T_1 in a functional manner by indicating in terms of T_1 notably the characteristic effects of the tokens that come under these types – to put it differently, the causal roles that tokens of these property types exercise.

2) One looks for realizers of these causal roles in the domain of the properties of T_2. The realizers of the functionally defined property types may differ in physical composition.

3) One explains in each case – that is, for each token – why there is a functional property falling in the domain of T_1 instantiated by describing how a configuration of properties in the domain of T_2 present in the situation under consideration brings about the effects that are characteristic of the functional property type in the domain of T_1 in question.

Functional reduction hence offers in each case, multiple realization notwithstanding, a causal explanation of why there is a property token present falling in the domain of T_1 by telling us how the effects that are characteristic of the property type in question are brought about (cf. what Chalmers 1996, pp. 42-51, calls a reductive explanation). Functional reduction thereby explains why there are properties falling in the domain of T_1 instantiated in the world and thus shows how T_1 is about salient properties. Functional reduction supports one-way conditionals of the sort that everything that comes under a physical type P_1 also comes under a special science type S ($\forall x\ (P_1x \rightarrow Sx)$); but it does not support biconditionals linking the types P_1 and S, since some tokens that come under S come under the physical type P_2 instead of the physical type P_1.

Though one-way conditionals are sufficient for the discovery of realizer types and reductive explanations, they are not sufficient for reducing T_1 to T_2, even if the domain of objects of T_1 is a proper part of the domain of objects of T_2 (T_2 may be a fundamental and universal physical theory). The reason is that one cannot deduce the laws of T_1 from the laws of T_2: there are no concepts available in T_2 that are coextensive with the concepts proper to T_1 and in which the laws of T_2 can be formulated, insofar as they are pertinent for that part of the domain of objects of T_2 that is identical with the domain of objects of T_1. Consequently, one cannot deduce the laws of T_1 from laws of T_2: the concepts figuring in fundamental and universal laws of nature (such as e.g. the laws of gravity or electromagnetism) are too general in order to deduce the laws of T_1 from these laws, and the concepts proper to T_2 that seize particular realizer types of property types of T_1 and laws or law-like generalizations formulated in terms of these concepts are too specific to capture the property types on which T_1 focuses: if there is multiple realization, several concepts proper to T_2 are needed to cover the extension of a single concept proper to T_1.

Nonetheless, there is a possibility to turn the mentioned one-way conditionals into biconditionals. Consider configurations of physical property tokens that all come under a single special science type S, but that make true different physical descriptions (P_1, P_2, P_3, etc.). Each of the physical descriptions P_1, P_2, P_3, etc. expresses a minimal sufficient condition to bring about the effects that characterize S under normal conditions. However, configurations coming under P_1 differ from configurations coming under P_2 and from configurations coming under P_3 in their physical composition. Differences in physical composition imply differences in the way in which the effects characterizing S are brought about. For each such difference, environmental conditions are conceivable in which that difference is relevant even on the level of abstraction at which S is situated.

Consider classical genetics. Natural selection explains why there is multiple realization in the domain of classical genetics (see Papineau 1993, p. 47): depending on the environmental conditions, only some of the causal powers of a given molecular configuration, realizing a property of the type S of classical genetics, are pertinent for selection. Against this background, it is reasonable that the proper concepts of classical genetics abstract from molecular differences. There are for instance molecular differences among DNA sequences possible that, under certain cellular conditions, do not amount to phenotypic (functional) differences. Nonetheless, these molecular differences are different ways to bring about the effects that define S. But this implies that there is at least one difference in the production of side effects that are systematically linked with the main effects in question – such as different causal interactions with the molecular environment within the cell during the causal process from a gene to the production of its characteristic phenotypic effects. For any such difference in side effects, there is a molecular environment possible in which that difference leads to a detectable functional difference within the scope of classical genetics and the evolutionary context because any such difference may become pertinent to selection in certain environments (see Rosenberg 1994, p. 32). Consequently, that difference can in principle also be considered in terms of the concepts that are proper to classical genetics. Hence, for any type S of T_1 (that is multiply realized by P_1, P_2, P_3, etc.), it is possible to conceive functional sub-types S_1, S_2, S_3, etc., taking those side effects into account.

These sub-types are no longer multiply realizable, since any molecular difference that is relevant to distinguish between different types of realizers leads to specific functional differences. The functionally defined sub-types of a special science hence correspond to one physical type each. They are nomologically coextensive with physical types and thus make it possible to reduce theories of a special science to physical theories in a functional manner.

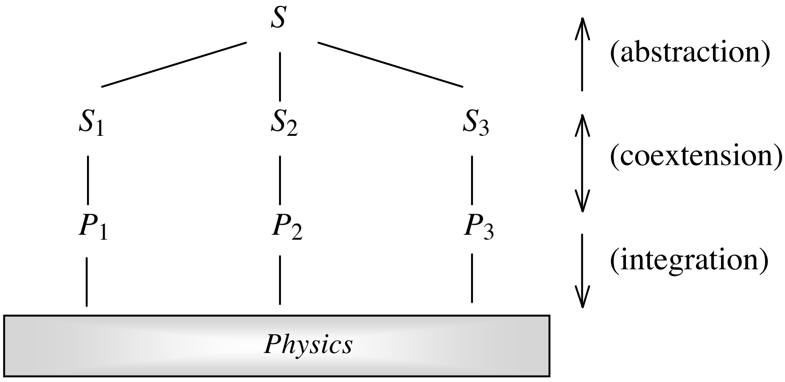

Figure 5: The reduction of a theory of a special science to a physical theory via functional sub-types.

More precisely and more generally speaking, (1) within a physical theory T_2, one builds the concepts P_1, P_2, P_3, etc. capturing the differences in composition among the physical configurations that are all described by the same concept S in T_1. (2) One makes S more precise by building functional sub-concepts (sub-types) S_1, S_2, S_3, etc. of S, seizing the systematic side effects linked to the different ways of producing the effects that define S. Provided that one such functionally defined sub-concept can be construed for each type of realizer of S in such a way that the former grasps the functional differences to which the latter give rise under certain circumstances, it follows that these sub-concepts S_1, S_2, S_3, etc. are nomologically coextensive with the concepts P_1, P_2, P_3, etc. (3) One can reduce any concept S of T_1 to T_2 via S_1, S_2, S_3, etc. and P_1, P_2, P_3, etc. Starting from T_2, one builds P_1, P_2, P_3, etc. and then deduces S_1, S_2, S_3, etc. from P_1, P_2, P_3, etc. given the nomological coextension. One gains then S by abstracting from the conceptualization of the functional side effects contained in S_1, S_2, S_3, etc., retaining only the main functional specification they have in common, which is nothing but the functional definition of S. This abstraction step depends on what the world is like – that is, what salient normal environmental conditions there are – rather than on our heuristic and practical aims. It enables thereby to highlight genuine causal similarities in the world that are brought out by the special sciences. As regards the laws, one can formulate the laws of T_1 in terms of S_1, S_2, S_3, etc. by adding more functional details. Given the nomological coextension, one can deduce these sub-type laws from the laws of T_2, couched in terms of P_1, P_2, P_3, etc. and then gain the laws of T_1 formulated in terms of S by a theory-immanent abstraction from functional details (that are not relevant in many environmental contexts) (see Esfeld and Sachse 2011, chapters 4 and 5, for a detailed account).

In conclusion, functional reduction, thus conceived, can achieve the same as Nagelian reduction, namely a fully-fledged theory reduction. As mentioned above, ontological and epistemological reduction stand or fall together. Nonetheless, one can still have both, as envisaged in the Vienna circle, although today's argumentative focus is different.

References

Carnap, Rudolf (1928): *Der logische Aufbau der Welt*. Leipzig: Meiner.

Chalmers, David J. (1996): *The conscious mind. In search of a fundamental theory*. New York: Oxford University Press.

Davidson, Donald (1970): "Mental events". In: L. Foster & J. W. Swanson (eds.): *Experience and theory*. Amherst: University of Massachusetts Press. Pp. 79-101. Reprinted in D. Davidson (1980): *Essays on actions and events*. Oxford: Oxford University Press. Pp. 207-225.

Dretske, Fred I. (1989): "Reasons and causes". In: J. E. Tomberlin (ed.): *Philosophical Perspectives 3: Philosophy of mind and action theory*. Oxford: Blackwell. Pp. 1-15.

Esfeld, Michael (2010): "Causal overdetermination for Humeans?" *Metaphysica* 11, pp. 99-104.

Esfeld, Michael and Sachse, Christian (2011): *Conservative reductionism*. New York: Routledge.

Fodor, Jerry A. (1974): "Special sciences (or: The disunity of science as a working hypothesis)". *Synthese* 28, pp. 97-115.

Harbecke, Jens (2011): "Mind in a Humean world". Forthcoming in *Metaphysica*, DOI 10.1007/s12133-011-0086-2.

Heil, John and Mele, Alfred (eds.) (1993): *Mental causation*. Oxford: Oxford University Press.

Kim, Jaegwon (1998): *Mind in a physical world. An essay on the mind-body problem and mental causation*. Cambridge (Massachusetts): MIT Press.

Kim, Jaegwon (2005): *Physicalism, or something near enough*. Princeton: Princeton University Press.

Ladyman, James (2008): "Structural realism and the relationship between the special sciences and physics". *Philosophy of Science* 75, pp. 744-755.

Lewis, David (1994): "Lewis, David: Reduction of mind". In: S. H. Guttenplan (ed.): *A companion to the philosophy of mind*. Oxford: Blackwell. Pp. 412-431. Reprinted in D. Lewis (1999): *Papers in metaphysics and epistemology*. Cambridge: Cambridge University Press. Pp. 291-324.

Nagel, Ernest (1961): *The structure of science. Problems in the logic of scientific explanation*. London: Routledge.

Noordhof, Paul (1998): "Do tropes resolve the problem of mental causation?" *Philosophical Quarterly* 48, pp. 221-226.

Oppenheim, Paul and Putnam, Hilary (1958): "Unity of science as a working hypothesis". In: H. Feigl, M. Scriven and G. Maxwell (eds.): *Concepts, theories, and the mind-body problem. Minnesota Studies in the philosophy of science. Volume 2*. Minneapolis: University of Minnesota Press. Pp. 3-36.

Papineau, David (1993): *Philosophical naturalism*. Oxford: Blackwell.

Robb, David (1997): "The properties of mental causation". *Philosophical Quarterly* 47, pp. 178-194.

Robb, David (2001): "Reply to Noordhof on mental causation". *Philosophical Quarterly* 51, pp. 90-94.

Rosenberg, Alexander (1994): *Instrumental biology or the disunity of science*. Chicago: University of Chicago Press.