



UNIVERSITY
OF WOLLONGONG
AUSTRALIA

University of Wollongong
Research Online

Faculty of Arts - Papers (Archive)

Faculty of Law, Humanities and the Arts

2010

Hierarchies and levels of reality

Patrick McGivern

University of Wollongong, patrickm@uow.edu.au

Alexander Rueger

University of Alberta

Publication Details

Rueger, A. & McGivern, P. 2010, 'Hierarchies and levels of reality', *Synthese*, vol. 176, no. 3, pp. 379-397.

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library:
research-pubs@uow.edu.au

Hierarchies and Levels of Reality

Patrick McGivern and Alexander Rueger

University of Alberta

Talk of levels or layers of reality is ubiquitous in science and in philosophy. It is widely assumed, for instance, that physical, chemical, biological, and mental phenomena can be ordered in a hierarchy of levels, and that what happens at the so-called micro level determines the goings on at the macro level. Questions about both reduction and emergence presuppose some structure of levels; without this, the very question of whether higher levels reduce to more basic ones could not be asked.

Often intuitions seem sufficient to establish the ordering and the nature of the levels. But, given the central importance of this picture, more should be said about how we identify levels, what populates them (objects, properties, facts, ...), what principles govern the hierarchy, and how widely any particular account of levels is applicable.¹

In this paper we address these questions, taking our starting point within (classical) physics. We examine the features of the hierarchies, ordering principles, and population of levels in physics and argue that the structure of levels suggested by physics is quite different from that assumed in widely assumed philosophical schemes, in particular Oppenheim and Putnam's (1958) account of levels of reality.

1. The Oppenheim-Putnam Hierarchy

¹ For an inspiring discussion of some of these questions from a (mainly) biological point of view see Wimsatt 1994.

According to the ‘working hypothesis’ defended by Oppenheim and Putnam (1958, 9), the universes of discourse for different branches of sciences can be ordered into a hierarchy of levels. They suggest the following hierarchy:

6. Social groups
5. (Multicellular) living things
4. Cells
3. Molecules
2. Atoms
1. Elementary particles

Each level is populated by the entities described by a particular branch of science, and entities at adjacent (‘proper’) levels are related by the part-whole relation, so that every member of level $i + 1$ can be decomposed into members of level i (for $i > 1$). However, Oppenheim and Putnam also stipulate that any whole that’s exhaustively decomposable into parts belonging to level i also belongs to i ; thus, the population at each level includes all entities at higher levels. The hierarchy is therefore ontologically conservative in the sense that moving up the ordering does not add any ingredients to the world. All the entities are already contained in the bottom layer, the level described by the most basic theories in physics.

The relations between successive levels were suspected to be reductive: the branch of science dealing with level i was supposed to be a ‘potential micro-reducer’ of the branch of science dealing with the next highest level $i + 1$. While it remains an empirical question whether reductions can actually be found between the branches of science dealing with each pair of adjacent levels, Oppenheim and Putnam argued that there was reason for optimism.

This account of levels has its attractions. Generating the hierarchy of layers by adding parts to form wholes which serve as parts for higher-level wholes, etc., leads to

some seemingly plausible consequences: entities increase in spatial extension as we move up the hierarchy, and the ordering seems to correspond nicely with widespread views about different sciences dealing with entities of different ‘complexity’. However, this scheme has in one sense not enough contact with scientific practice. The view of reduction with which Oppenheim and Putnam work is characterized by two main requirements (ibid., 5): for T^* to be reducible to T , (a) the ‘observational consequences’ of T^* must be derivable from T , and (b) T must be of greater ‘systematic power’ than T^* .² Apart from the often repeated objection that reductions in science only rarely (if ever) fit this view,³ this notion of reduction has a further shortcoming. One thing we should expect from a theory of reductive levels is an account of the ontological relations between the things populating different levels. However, on Oppenheim and Putnam’s account, the ‘theoretical’ parts of two theories T^* and T remain unconnected even when one theory is reducible to the other. Hence, successful reductions either have no ontological consequences, or they have the extreme ontological consequence of ‘eliminating’ the objects in the domain of the theory that has been reduced. We might think that the relevant ontological relation is the part-whole relation that is supposed to relate entities at different levels. However, on Oppenheim and Putnam’s account this relation is postulated to hold independently of successful reductions: the success or failure of reduction between levels tells us nothing about the parthood relation between the entities involved.

This is an unsatisfactory situation. Cases of successful or failed reductions should provide insight into the ontological relations between the objects of the theories involved.

² Cf. also Kemeny/Oppenheim 1952.

³ E.g., Sklar 1967.

In fact, this sort of evidence may well be the only support we have for making ontological claims, such as claims about part-whole relations: we cannot postulate part-whole relations between lower and higher level theories independently of looking at the way reductions of the latter to the former work or fail. Accordingly, we'll now examine a case from physics where the success or failure of reduction has consequences for the ontological relations between levels.

2. A View From Physics

The equation of motion for a damped harmonic oscillator

$$m d^2x/dt^2 + kx + c dx/dt = 0$$

describes a system with mass (m), restoring force (kx), and damping ($c dx/dt$). It tells us that the restoring force and the damping are proportional to the displacement (x) and the velocity (dx/dt), respectively. The solution of the equation describes the *behaviour* of this system, that is, the distribution of properties (here: positions) of the system over time:

$$x(t) = A(t) \cos(\omega t - \delta),$$

a series of oscillations with frequency ω and phase δ and with a gradually decaying amplitude $A(t)$.

To define a hierarchy of levels, we can associate a particular *scale* with a given equation, such that the solutions to that equation will describe the system's behavior at that scale. There are various ways of doing this: depending on the system (or its equations), it is possible to have spatial, temporal, energy, momentum, or particle number scales, among others. Since scales are quantified, they have a natural ordering. In the

example, it could turn out that the amplitude $A(t)$ changes at a much slower time scale than the individual oscillations.

Once we use scales to characterize different theories, the most natural way to establish a reductive connection between them is in terms of a limit relation: when we let the scale which is characteristic for processes described by a theory T become negligible compared with the scale at which processes operate that are described by another theory T_0 , the behaviour described by T should approximate (and ultimately go over into) the behaviour described by T_0 . Thus we would say that a theory of processes at a large spatial scale reduces to a theory of processes at a smaller scale if the latter description goes over into the former (presumably less basic) theory when the ratio of the scales (small/large) approaches zero.⁴

More formally, we define:

A higher-level theory T_0 *reduces to* a lower-level theory T *iff* the solutions of the equations of T , $T(x)$, uniformly converge towards the solutions of the equations of T_0 , $T_0(x)$, in the limit of an appropriate parameter.

If the limit of $T(x)$ is not uniform (but rather ‘singular’), T_0 is not reducible to T .

Although the relata of the reduction relation are *theories*, given the limit notion of reduction, it is the *solutions* of the theories’ equations that are crucial since the limit operation is not defined for the equations themselves. And in terms of the ontological

⁴ For earlier discussions of this concept of reduction see Nickles 1973 and Batterman 1995. Note that we differ here from the way Nickles (1973) and others have understood the direction of the reduction relation. Nickles defines the “philosophers’ sense” of reduction as: T_0 reduces to T if T_0 can be derived from T (together with further premises); the direction of reduction here goes from the less basic to the more basic theory. The “physicists’ sense” of reduction, by contrast, requires that T reduces to T_0 if the limit of T (in an appropriate parameter) is T_0 ; the direction of the reduction relation is obviously reversed: it now goes from the basic to the less basic theory. We adapt the latter sense to the philosophers’ understanding by stipulating that the less basic T_0 reduces to the basic T if the limit of T is T_0 . There should be no problem with this definition as long as it is not confused with the original physicists’ sense of reduction.

commitments of the theories, it's the *behaviours* of systems of which we can ask questions concerning reduction, and hence it's the behaviours – rather than objects – that should populate a hierarchy of levels or scales.

Under fairly simple circumstances, behaviour at higher scales is often not reducible to behaviour at more basic scales. Consider the treatment of steady state heat conduction in a one-dimensional rod of length L .⁵ This system is described in terms of its temperature $T(x)$ and its thermal conductivity $k(x)$ which both depend on the spatial variable x . At the ends of the rod ($x=0$ and $x=L$) the temperature is held constant. The equation for heat conservation in this system is

$$(1) \quad (d/dx) [k(x)dT(x)/dx] = 0,$$

which tells us that the heat flux at any point in the rod ($= k(x) dT(x)/dx$) is constant.

At the lower, 'microscopic' level, we assume that the system consists of individual 'atoms', separated by empty space forming a periodic lattice with a period of length $P = \epsilon L$, with $\epsilon \ll 1$. The microscopic conductivity $k(x)$ will then be a rapidly oscillating function of position: high around the location of each atom, low in the interatomic spaces. This behaviour of the conductivity is indicated by writing k as a function of x and x/ϵ . The dependence on x/ϵ manifests itself as rapid variations because $d/dx (k(x/\epsilon)) = (1/\epsilon)d/dx(k(x/\epsilon))$, that is, the derivative of k is large for small ϵ . At the *macroscopic* level, the level at which we make measurements about the system's temperature, etc., no such fast variations are observed: temperature and conductivity at this level are typically slowly changing functions of position. The temperature distribution at the micro level is thus described by:

$$(2) \quad (d/dx) [k(x, x/\epsilon)dT(x)/dx] = 0.$$

⁵ See, for instance, Holmes 1995, 224ff., and Rueger 2006.

Note that any *solution* of (2), will represent the property of having such-and-such a distribution of micro-level temperature.

The aim now is to find, starting from eq. (2), a *macroscopic* heat equation which describes the behaviour (temperature distribution) at the macroscopic level. Such an equation would no longer contain the rapidly varying $k(x, x/\varepsilon)$ but instead a slowly varying macro conductivity $K(x)$. The obvious way of going from a discrete description at a small scale, eq. (2), to a continuous description at a larger scale, is to let the ratio $\varepsilon = P/L \rightarrow 0$. This ‘continuum limit’ reflects the intuitive requirement that the macroscopic representation smoothes out the details at the micro level. We therefore seek an expansion of the solutions of (2) in terms of the small parameter $\varepsilon = P/L$ and expect to retain, in the limit $\varepsilon \rightarrow 0$, the solution of the macroscopic equation, $T_0(x)$:

$$(3) \quad T(x) = T_0(x) + \varepsilon T_1(x) + \varepsilon^2 T_2(x) + \dots$$

A successful expansion of this sort, where $T(x)$ uniformly converges towards $T_0(x)$, would suggest that the micro behaviour can be represented as the macro behaviour plus some small corrections.

Letting the parameter ε go to zero, however, will result in more and more rapid oscillations of the coefficient $k(x, x/\varepsilon)$, hence of $T(x)$, and in general, $T(x)$ will show a ‘singular’ dependence on ε . In general – that is, unless $k(x, x/\varepsilon)$ is chosen in special ways – the solution $T(x)$ will therefore not converge uniformly to $T_0(x)$ in the limit $\varepsilon \rightarrow 0$.

That is, we have

$$\lim T(x) \neq T_0(x), \text{ for some } x.$$

This indicates that the micro behaviour of the system cannot be represented as essentially the macro behaviour, described by $T_0(x)$, plus small corrections. $T(x)$ and $T_0(x)$ differ

considerably, even for the smallest nonzero values of the parameter. $T_0(x)$ is not reducible to $T(x)$.

Nevertheless, one can still construct a *uniformly valid approximation* of $T(x)$ in the form of a power series, starting with the macro solution $T_0(x)$. To do this we introduce *two* length scales in the micro description, the macroscopic scale x and a microscopic scale $y = x/\varepsilon$. Thus we consider an asymptotic expansion of T as a function of *two independent* variables:

$$(4) \quad T(x, y) = T_0(x, y) + \varepsilon T_1(x, y) + \varepsilon^2 T_2(x, y) + \dots$$

This procedure gives us

$$(5) \quad T(x, y) = T_0(x) + \varepsilon T_1(x, y) + \varepsilon^2 T_2(x, y) + \dots$$

which shows that $T_0(x, y) = T_0(x)$ is indeed independent of the microscopic variations measured by y ; it depends only on the macro level position x , as a macro level quantity should. But it also shows that the solution of the micro equation (2) cannot be written in terms of the micro scale y alone; the higher order terms like $T_1(x, y)$ will contain both scales. This, again, shows that reduction of the macro behaviour to the micro behaviour is not possible here.

3. Interpreting the Failure of Reduction

How should we interpret this result? In analogy with the disagreement between ‘reductive’ and ‘non-reductive’ physicalists in the philosophy of mind, we can imagine two general attitudes toward the failure of reduction here. A *reductionist* about levels will take the failure of reduction as a basically epistemic issue: the macro description is at least slightly wrong, and the only correct description of the system is the microscopic

one – although we don't have solutions purely at the micro level; the best we can do is produce approximations of the 'mixed' sort $T(x, y)$. The hierarchy of levels would then become a hierarchy of (slightly) erroneous views of the basic reality, with no serious ontological consequences. While this is a consistent view, it pays the price of creating a mismatch between metaphysics and scientific practice, since the mixed solutions are the best descriptions we have of the system's behaviour.

An *anti-reductionist*, by contrast, will take the failure of reduction seriously and accept that macro behaviour is not just a product of special epistemic circumstances. In this case, there remains a choice to be made concerning the attitude one should take towards the 'mixed' solution $T(x, y)$. There are two alternatives:

(AR1) Although we accept nonreducibility of macro to micro behaviour, we do not take the mixed solution seriously but regard it as an approximation to the solution of the micro equations, a substitute for a purely microscopic solution which we are unable to obtain. So, in effect, we treat $T(x, y)$, despite the fact that it depends on both scales, *as if* it represented a solution at the micro scale only. We are then in a position to discuss the relation of the macro behaviour $T_0(x)$ to the micro behaviour, even though we do not have a solution of the equations purely at the micro level.

(AR2) We do take the mixed solution seriously, that is, we do not regard it as a substitute of an unavailable purely micro solution but as a description of the behaviour of the system that is indeed composed of ingredients from both scales. The behaviour of the heated rod, the distribution of temperature over the length of the system, consists of several components, one of which is the purely macro behaviour $T_0(x)$. In this case, the question of how the macro behaviour and the behaviour of the system relate to each other

has an obvious answer: they are not entirely distinct; their relation is one of *part to whole*. That's what a literal reading of the expansion eq. (4) indicates: $T_0(x)$ is 'contained' in $T(x, y)$.⁶ And although the part-whole interpretation of the relation of macro behaviour to the behaviour of the system as represented by $T(x, y)$ is most natural if we take the latter seriously according to option (AR2), the same interpretation can be used if we choose option (AR1), i.e., when we take $T(x, y)$ to be a substitute for a purely microscopic description of the system's behaviour.⁷

The claim that the macro behaviour is a part of the micro behaviour might sound paradoxical or at least peculiar because we tend to think of parts and wholes in spatial terms. On a spatial view, the parts have to be smaller than the whole; the parts are spatially included in the whole. This, of course, cannot be the right sense of parts and wholes in our case of the relation of behaviours. But this isn't the only way to think about parthood. There are a variety of senses of 'parthood' that do not imply a particular spatial relation between parts and wholes. Thus, in the case of vector and Fourier components, the parts are clearly not to be understood as spatially smaller than the whole. One particularly suggestive way to understand parthood in this non-spatial sense is in

⁶ It might seem arbitrary to interpret an equation like

$$T(x) = T_0(x) + \varepsilon T_1(x) + \dots$$

as representing a part-whole relation in which $T(x)$ is the whole. The equation itself is symmetric in the sense that we can just as well write

$$T_0(x) = T(x) - \varepsilon T_1(x) + \dots,$$

so that it looks as if $T_0(x)$ might be the whole and $T(x)$ a component. But there is no arbitrariness here. Although the equations are symmetric, the asymmetry required for our interpretation is introduced by the perturbation approach itself. We are looking for a representation of the system's (total) behaviour, an appropriate solution of the equations of motion, which is $T(x)$. $T_0(x)$, by contrast, solves the equations of motion only approximately, at the lowest order (e.g., ε^0) of the perturbation theory; the complete solution is $T(x)$ and therefore we are justified in interpreting $T(x)$ as the whole and $T_0(x)$ as a component.

⁷ That a certain sort of behaviour is a part of another sort is, of course, familiar in physics from cases where the behaviours are represented as vectors. In these cases, the vector components relate to the resultant vector as parts to the whole. Similarly, when we decompose some behaviour into its Fourier components we have examples of part-whole relations between behaviours.

terms of the sets of causal powers associated with the properties or behaviors involved. For instance, the claim that the macro behaviour is part of the micro behaviour could be understood as the claim that the set of causal powers characteristic of the macro behaviour is included, as a subset, in the set of powers that constitute the micro behaviour.⁸

Interpreting the relation of $T_0(x)$ and $T(x, y)$ as a part-whole relation introduces a *second ordering principle* for hierarchies of levels. So far we considered only the spatial scales at which the behaviour takes place; these define a hierarchy of micro and macro behaviours. Now we have, in addition, a part-whole hierarchy in which the macro behaviour, $T_0(x)$, is a part or component of the behaviour, $T(x, y)$. Call $T(x, y)$, which we can take either seriously as behaviour at two scales (attitude II), or as a mere substitute for behaviour at the micro scale alone (attitude I), the *total* behavior of the system. From the point of view of the part-whole ordering, the macro behavior would be located at the lower level, the total behaviour at the higher level.

But isn't there a serious problem now? On interpretation (AR1), the two orderings seem to give inconsistent results with respect to the same behaviours: in terms of spatial scales, the macro behaviour should be 'higher' in the hierarchy than the micro behaviour ($T_0(x)$ above $T(x,y)$); in terms of parts and wholes, $T(x, y)$, the whole, should be 'higher' than $T_0(x)$, one of the parts. Doesn't the existence of two inconsistent level hierarchies constitute evidence against an ontological reading of what populates the levels? The case

⁸ This way of thinking about parthood has been frequently applied in metaphysics and the philosophy of mind. Discussion of the relation of higher and lower level properties in these areas is often phrased in terms of 'realization': a lower level property P realizes a higher level property Q just in case the causal powers associated with Q are a subset of the causal powers associated with P. And furthermore, an instantiation of P includes as a part the instantiation of Q. By analogy, we claim that the micro behaviour realizes the macro behaviour (in this sense) and that the former includes the latter as a part. See Shoemaker 2001, 78ff. (with further references). Compare also Yablo 1992.

might seem similar to the ‘incompatible models argument’ against scientific realism, which says that a realist attitude towards models in science is mistaken because scientists routinely and successfully employ models of one and the same natural system which are incompatible with respect to the properties they assign to the real system.⁹ Given such inconsistent attributions, models should obviously not be interpreted realistically. The analogous argument in the case of hierarchies of levels would be: since the pair of behaviours, $T(x, y)$ and $T_0(x)$, are ordered in incompatible ways in the spatial scales and in the part-whole hierarchy, the behaviours should not both be taken as real but rather as different descriptions (from different points of view) of the real behaviour.¹⁰

The analogy of the arguments, however, is only apparent and does not justify the common anti-realist conclusion. What is significantly different in the levels case is the existence of two distinct ordering principles that underlie the hierarchies. There is no equivalent to this feature in the incompatible models argument. Compare the ordering of a set of objects according to a colour scale (say, red-yellow-blue) and according to size (large-small). This can, under easily imaginable circumstances, lead to incompatible hierarchies. There may be red objects, that is, high-colour level objects which come to populate the lower-size level, and so on. Since the ordering principles pick out different aspects or features of the objects, there is no incompatibility in this case analogous to the incompatibility in the models argument. Similarly for spatial scales and part-whole hierarchies of behaviours. The large-scale behaviour exists at a higher level than the

⁹ For a presentation of the argument see, for instance, Morrison 2000, 47ff.

¹⁰ Wimsatt (1994, 265ff.) has discussed this sort of problem for a layered view of the world and concluded that in those cases where multiple hierarchies are encountered, the ontological concept of levels is no longer applicable and should be replaced by the notion of ‘perspectives’. His discussion, however, does not consider the possibility of nonspatial part-whole relations.

small-scale behaviour in an ordering of spatial scales. But that's compatible with the former behaviour being a part in the latter if we spell out the part-whole relation in non-spatial terms. So the presupposed hierarchy of behaviours in terms of an increasing spatial scale does not get upset by the non-spatial part-whole relations between behaviours. The problem of inconsistent orderings would arise if we thought of the hierarchy in terms of objects of increasing spatial extension. A macro object, obviously, cannot be a (spatial) part of a micro object; the macroscopic heated rod is not a part of its microscopic constituents. But a macro behaviour can be contained, as a component structure, in a micro behaviour.

What if we adopt (AR2), the view that the behaviour described by $T(x, y)$ cannot be associated with only one scale? *This behaviour does not really populate a definite level in a spatial scales hierarchy* – but it still is an element in a part-whole hierarchy. Thus it is not correct to interpret the result of the multi-scale calculation as showing that the macro behaviour is a part of the *micro* behaviour; it is a part of the 'mixed' behaviour $T(x, y)$ which has no definite place in the micro-macro hierarchy. Since $T(x, y)$ cannot be ordered in terms of micro and macro spatial scales as lower than $T_0(x)$, there is no inconsistency with the ordering of behaviours in terms of parts and wholes where $T(x, y)$ is higher than $T_0(x)$.

Finally, compare the situation in these singular cases with the case of a regular perturbation expansion. If this approach is successful, we have $T(x) = T_0(x) + \varepsilon T_1(x) + \dots$, and we regard $T_0(x)$ as reducible to $T(x)$. This can, of course, also be interpreted as a whole, $T(x)$, composed of parts $T_0(x)$, etc. The difference between this and the singular case is that the parts and their whole exist here at one and the same scale.

These considerations should be sufficient to respond to an objection that Batterman (2002, 116) has leveled against a part-whole interpretation of the kind we propose. The case he discusses is different from ours but the basic features of his argument apply generally. He considers two theories that are supposed to account for the various features of rainbows, an optical theory operating with rays (geometric optics) and one based on a wave theory of light (wave optics). Can the ray-theoretic structure of a rainbow be understood as part of an overall wave-theoretic structure? That they could is suggested by the fact that solutions to the wave-theoretic equations can be constructed with a ray-theoretic component as their leading term. But Batterman argues that this attempt to read a part-whole relation into the equations would be misguided. While everybody agrees that the wave theory is more ‘fundamental’ than the ray theory, the suggested reading would have us regard ray-theoretic features as parts of wave-theoretic ones. This, Batterman suggests, violates the general principle that ‘wholes’ cannot be more fundamental than their parts.

Such a principle seems unobjectionable enough, provided that we are clear about (i) the sense of parthood being considered, and (ii) the sense of fundamentality involved. In our case, we are considering a hierarchy ordered by a non-spatial parthood relation connecting micro and macro behaviors so that the macro behavior is part of the micro behavior. Obviously, the micro is in some sense ‘spatially’ more fundamental than the macro, and if we were insisting that the macro was somehow a spatial part of the micro, we would be in trouble. But there are other ways of understanding fundamentality. For instance, if we understand the claim about parthood between behaviours in terms of the subset relation between their associated causal powers – so that the causal powers of the

macro property or behaviour form a subset of those of the micro property or behaviour – we can accept the principle that parts are more fundamental than wholes in the very reasonable sense that if X is a subset of Y, then X is more fundamental than Y.

4. Conclusion

The hierarchy of levels we have described differs in several important ways from the usual Oppenheim-Putnam hierarchy. It is *behaviours*, and not entities, that populate levels; and the success or failure of reduction between levels has ontological consequences for those levels. Furthermore, if we take attitude (AR2) towards $T(x, y)$, we cannot say – as we pointed out before – that this behaviour belongs to either the micro or the macro level. This is surprising if we think in terms of the Oppenheim-Putnam hierarchy. While that view assigns a level to every object one might think of, the hierarchy of behaviours at different scales is less inclusive: there are behaviours that do not find a place in the micro-macro ordering. Thus it is not true that for any pair of behaviours we can say that one is ‘higher’ than the other in the spatial sense of ordering.¹¹ But although the spatial scales ordering may not be applicable anymore, the procedure for calculating $T(x, y)$ and $T_0(x)$ defines an alternative ordering in terms of parts and wholes. This alternative hierarchy, however, has still very little to do with the Oppenheim-Putnam model because it is an ordering according to a non-spatial sense of parts and wholes.

References

R. Batterman 1995: “Theories Between Theories.” *Synthese* 103, 171-201

¹¹ But see Kim 2002 for analogous observations about the Oppenheim-Putnam model.

- R. Batterman 2002: *The Devil in the Details. Asymptotic Reasoning in Reduction, Explanation and Emergence*. Oxford UP
- M.H. Holmes 1995: *Introduction to Perturbation Methods*. New York: Springer
- J. Kemeny and P. Oppenheim 1952: "On Reduction." *Philosophical Studies* 7, 6-19
- J. Kim 1998: *Mind in a Physical World*. MIT Press
- J. Kim 2002: "The Layered Model: Metaphysical Considerations." *Philosophical Explorations* 5, 2-20
- M. Morrison 2000, *Unifying Scientific Theories*. Cambridge: Cambridge UP
- T. Nickles 1973: "Two Concepts of Intertheoretic Reduction." *Journal of Philosophy* 70, 181-201
- P. Oppenheim and H. Putnam 1958: "Unity of Science as a Working Hypothesis", in: H. Feigl et al. (eds.), *Concepts, Theories, and the Mind-Body Problem*. (Minnesota Studies in the Philosophy of Science, vol. II). Minneapolis: Uof Minnesota Press, 3-36
- A. Rueger 2006: "Functional Reduction and Emergence in the Physical Sciences." *Synthese* 151, 335-346
- S. Shoemaker 2001: "Realization and Mental Causation", in: C. Gillett et al. (eds.), *Physicalism and Its Discontents*. Cambridge UP, 74-98
- L. Sklar 1967: "Types of Intertheoretic Reduction." *British Journal for the Philosophy of Science* 18, 109-124
- W.C. Wimsatt 1994: "The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets." M. Matthen et al. (eds.), *Biology and Society*. (Canadian Journal of Philosophy, Supplementary volume 20.), Calgary: Uof Calgary Press, 207-274
- S. Yablo 1992: "Cause and Essence." *Synthese* 93, 403-449