University of Alberta

Physicalism, Compositionality, and Parthood: A Perspective from the Physical Sciences

by

Patrick Hugh John McGivern

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Department of Philosophy

Edmonton, Alberta
Fall 2005

0-494-08695-5

Your file Votre référence
ISBN:
Our file Notre reterence
ISBN:

for ANN

Abstract

Contemporary philosophy is dominated by the doctrine of physicalism, the claim that in some sense the physical facts about the world determine all other facts about the world. In this thesis, I examine the connection between physicalism and the related but distinct claim of compositionality. Compositionality is the claim that the facts about *parts* determine the facts about the things those parts compose. I argue that a proper understanding of physicalism requires a proper understanding of compositionality, and that this in turn requires an understanding of the variety of types of parthood found in the physical sciences. Most significantly, I argue that some explanations in physics should be understood as appealing to 'non-spatial' parts of systems and processes. Unlike ordinary spatial parts, non-spatial parts are not smaller than or spatially contained within the things they compose. Instead, non-spatial parts are as spatially extensive as the complexes they form. Such non-spatial parts play a crucial role in our understanding of the behavior of many physical systems. For instance, the behavior of many systems can be best understood by decomposing that behavior into distinct components operating on different temporal or spatial scales: these component systems form non-spatial parts of the system being described.

Developing this idea and examining its consequences for our understanding of physicalism is the main task of this thesis.

One of these consequences concerns our understanding of 'non-reductive' physicalism. I argue that the existence of non-spatial parts undermines the main philosophical argument against non-reductive physicalism, the so-called 'supervenience argument'. Current discussions of the supervenience argument all suppose that physical parts are exclusively spatial, and hence that the question of reduction concerns the relationship between entities and their spatial components. I argue that non-spatial parts stand outside of this received view, and are one example of how explanations in physics itself can be 'non-reductive'.

Acknowledgements

When I started this project a century ago, I never imagined it would take so long to complete. Of course, *then* it was a comparative study of the concepts of 'holism' and 'compositionality' in semantics and in quantum mechanics. While the interest in compositionality remained, everything else about the thesis changed: the interest in semantics echoes only very faintly here, and I was convinced long ago to avoid the philosophical perils of quantum theory. So perhaps one reason for the time this project has taken is my own inability to decide on what I was really interested in. But, to be fair, this is my first degree in philosophy.

There are many people I'd like to thank for helping me along the way. Thanks to Alex Rueger for his patient supervision and philosophical guidance over the years, to Rob Wilson for his demanding questions and comments (and for the use of his office this past winter, which was where much of the real work got done), and to the other members of my examining committee – Vadim Bulitko, Philip Hanson, and Bernard Linsky – for actually reading and apparently liking the thing. And thanks to Anita Theroux and Wendy Minns, not only for keeping me registered, paid, and organized, but also for making the Philosophy department an enjoyable place to work.

Many friends – especially Doug, Gina, Will, Barb, Andrew, Sheree, Jim, and Lucia – have offered me encouragement and support in this project, and I'm very grateful for that. Along with my family – Brenda, Trevan, Brian, Valeria, John, Aileen, Catherine, Tina, and my parents Jack and Sheila in particular – they have never made me question the wisdom of pursuing abstract studies in philosophy (well, except for Trevan maybe). Many thanks to all of you.

Part of this research was supported by a fellowship from the Social Sciences and Humanities Research Council of Canada, which was very nice of them.

Far, far above all, I'm grateful to Ann for all her support and encouragement through what has sometimes been an almost overwhelming project. At times, I was convinced I could not do it, but you always had faith in me. Thank you.

Mark, Rachel, Daddy's big paper is finally done.

Contents

List of Figures

Introduction

Roughly put, physicalism is the claim that 'the physical facts determine all of the facts': this claim plays an important role in contemporary philosophy, in particular in the philosophy of mind, but also in metaphysics and the philosophy of science more generally. In this thesis, I argue that properly understood, physicalism is really a claim about *compositionality*: the properties of 'wholes' are determined by those of their parts. Consequently, a proper understanding of physicalism requires a proper understanding of *parts*. In particular, a proper understanding of physicalism requires a proper understanding of the sorts of parts described by the 'physical' sciences.

Philosophers tend to assume that physical parts are *spatial* parts: if one thing is part of another, then the part is smaller than and contained within the spatial boundaries of the whole. Together with the assumption that the properties of parts determine the properties of wholes, this leads to the conclusion that the properties of wholes are determined by those of their smallest constituents. Since the properties of the very small – the subatomic particles and the like composing everything there is – are plausibly 'physical' by default, compositionality entails physicalism.

I don't question the existence of parts of this sort; instead, I defend the existence of parts of another sort as well. Accurately describing physical systems often involves decomposing those systems into distinct parts that are not smaller than nor spatially contained within the systems they compose. For example, patterns of fluid flow, chemical reaction rates, and oscillatory behavior of all sorts can often only be described by decomposing the systems involved into distinct systems operating on different *scales*. But these component systems are not smaller than or spatially contained within the system they compose. Instead, they overlap with the composite system in space and time. They are parts in the sense that the behavior of the composite system is the *result* of the combination of the behavior of the components: call such components *non-spatial parts*.

There are two reasons for wanting to investigate these sorts of parts. The first is that any adequate characterization of physicalism ought to accommodate the variety of entities and properties found in physics itself. The general focus on *spatial* parts in the existing literature on physicalism gives the impression that explanations in physics are characteristically or even paradigmatically *micro*-level explanations. But physics is a diverse field involving descriptions of the world at a wide variety of levels, and many of these are not immediately connected with the micro-level at all. *Macro*-level branches of physics, such as fluid mechanics, offer views of the world quite distinct from those found in particle physics. Understanding these branches of physics and the sorts of entities and properties they appeal to in their explanations is an important part of understanding physicalism.

The second reason for investing non-spatial parts follows from this first point about the diversity in physics itself. One of the central debates about physicalism concerns whether or not 'non-reductive' physicalism can be defended. The claim of non-reductive physicalism is that while the physical facts determine all other facts, it is not true that all facts reduce to or are identical with physical facts: facts from non-physical sciences can enjoy a degree of autonomy from those of physics. Given the usual understanding of physicalism as a claim about the micro-level, the question of whether or not non-reductive physicalism is possible becomes relevant for these macro branches of physics as well. Fluid mechanical properties, for example, surely *depend* upon micro-level properties of fluids, but the point of defending non-reductive physicalism is to defend the idea that the macro-level descriptions from fluid mechanics offer genuine insight into the structure of the world and are not merely convenient abbreviations of a true 'micro' reality too complex for us to comprehend directly. The spirit of non-reductive physicalism is found in the following remark from the physicist Steven Weinberg:

> Suppose that at some time in the future we came to know everything there is to know about water molecules, and that we had become so good at computing that we had computers that could follow the trajectory of every molecule in a glass of water...Even though we could predict how every molecule in a glass of water would behave, nowhere in the mountain of computer printout would we find the properties of water that really interest

us, properties like temperature and entropy. These properties have to be dealt with in their own terms...(Weinberg 1987, pp. 64-65)[1]

Appreciating non-spatial parts helps us to understand how and why these macro-level properties must be dealt with in their own terms. Spatial decomposition and non-spatial decomposition are different ways of 'carving up' the same reality: on the 'micro' decomposition, a given system consists of sub-atomic particles and their properties, and the facts about those parts determine the facts about the entire system. But on a 'non-spatial' decomposition, the same system consists of two or more components of a different sort. These are equally well parts of the system, but they are not 'micro-level' parts. Facts about these components *also* determine the facts about the entire system, but while these non-spatial components might depend on the features of the micro decomposition, they cannot be identified with any features on the micro-level.

The thesis is organized as follows. In chapter one, I outline the general interest in physicalism and argue that rather than the usual understanding of physicalism in terms of 'microphysicalism', physicalism is best understood in terms of compositionality. I then argue for a 'top-down' view of compositionality designed to accommodate the various ways in which individuals or systems can be decomposed into parts. My main goal here is argue that an appreciation of a variety of kinds of parts is essential to physicalism, and that our attention should

[1] Note that Weinberg's exact views on reductionism are unclear, given the surrounding discussion, and his purpose in this text is not specifically to discuss reductionism; however, the view he expresses here nicely captures the fundamental idea of non-reductive physicalism.

not be limited to 'spatial' parts. I also discuss the distinction between reductive and non-reductive physicalism and the motivations for defending a non-reductive view.

In the second chapter, I examine Jaegwon Kim's 'supervenience argument' against non-reductive physicalism, along with some prominent criticism of that argument from Ned Block[2]. Kim argues that since higher-level properties are determined by the distribution of micro-level properties, the causal efficacy of micro-level properties *excludes* that of any higher-level properties, since any higher-level causation would 'compete' with causally efficacious properties at the micro-level. This leaves non-reductive physicalism in the unattractive position of defending the irreducibility of causally inert properties: these might well be 'autonomous' from their micro-physical realizers, but being causally inert, they can barely be regarded as real properties at all. Thus non-reductive physicalism is untenable, and we are forced to identify higher-level properties with their micro-level realizers.

Block tries to show that Kim's argument must be flawed since it apparently has far reaching consequences: if Kim's argument is not flawed, then all causation takes place at the micro-level and no non-microphysical properties can be causally efficacious. Since it seems obvious that not all causation takes place at the micro-level, Block concludes that Kim's argument is flawed, and that non-reductive physicalism is defensible after all. To make matters worse, Block suggests that it is an open question whether or not there is *any* fundamental micro-

[2] Kim 1998, 2003; Block 1997, 2003.

level, or whether entities and properties are endless divisible. If *that* were true, then Kim's argument seems to show that there would be *no* causation in the world, since putative causation at leach 'lower' level would be pre-empted by causation at yet lower-levels. Since the conclusion that there is no causation in the world is even more clearly absurd than the suggestion that there is only micro-level causation, again we are to conclude that Kim's argument must be flawed. Kim counters Block's argument by suggesting that in most cases, we readily accept the reduction of higher-level properties to lower-level ones, and so the fact that higher-level properties can be causally efficacious is not surprising: those higher-level causes just *are* lower-level ones described using different concepts. This response raises an important question: *can* higher-level properties be identified with their lower-level 'realizers' in the way Kim suggests?

In the third chapter, I consider this last question in detail, and argue that the existence of 'non-spatial' forms of decomposition show that Kim's claim about property identities cannot be maintained. I begin by explaining the distinction between descriptions at different *scales*: this is often a more common way of characterizing what philosophers refer to as descriptions at different 'levels'. I then examine cases where systems are understood in terms of their behavior on multiple-scales. Descriptions at different scales are often necessary for finding an accurate account of a physical system, and sometimes the behavior of a system can only be described by decomposing it into distinct component systems operating on different scales. Such components form 'non-spatial' parts of the system they compose, since they are (at least typically) neither smaller than nor spatially contained within that composite system. More specifically, such components are characterized as distinct *processes*, where a process can be understood in terms of a generalized form of 'structural' properties. Structural properties are typically conceived of as properties relating to the spatial structure of an entity: having a particular structural property entails having particular parts with particular properties, standing in particular relations to one another. When the distinct scales involved in multi-scale analysis are purely spatial, it seems most natural to interpret these as claims about ordinary structural properties. Extending this idea, we can interpret distinct *temporal* scales in terms of 'spatio-temporally' structural properties: properties of having particular parts with particular properties standing in particular relations *at particular times*. I end the chapter by discussing the relevance of such parts to arguments such as Kim's, and by examining a question about the possibility of giving a 'reductive' account of these component systems.

In chapter four, I turn to other some other issues concerning the nature of non-spatial parts. I begin by examining the connection between spatial parthood and *mereology*, the formal theory of parts and wholes. I argue that there is no formal reason to exclude the possibility of non-spatial parts, and give an account of how we can understand non-spatial parthood in standard mereological terms. I then discuss an important objection to the claim that the examples I use in the third chapter reveal genuine parts of the systems of interest. One might object that the component processes described on different scales are not *real* but merely the reflection of artificial devices used to solve otherwise intractable mathematical problems. This connects the question of the reality of component processes to the more general problem of *component realism*, and I discuss this problem in detail in chapter four. This problem is quite general and applies not only to components on multiple scales, but also to components of vectors, forces, and the like. I argue that this opposition to realism about components is unfounded and that a plausible account can be given of the distinction between real or 'natural' decompositions and 'artificial' ones.

In the final chapter, I examine the connection between the idea of non-spatial parts and a recent suggestion about how the sort of causal competition worries raised by Kim might be addressed. One novel suggestion is that higher-level 'realized' properties might in some sense be parts of their lower-level realizers[3]: realized properties can be causally efficacious without being identical with their realizers, because the causal efficacy of the realized property is only part of that of its realizer. Obviously, this would be a non-standard form of parthood, and so I examine how well suited non-spatial parts are to making sense of this claim. I review some current formulations of this suggestion and various criticisms raised against it. I then describe an example of a relationship between systems described at different *spatial* scales that seems to support the claim of 'realization as parthood'. I conclude by considering the wider applicability of this idea and the need for a more detailed understanding of parthood in physics.

[3] See Shoemaker 2001, Clapp 2001, Rueger 2004.

I

Physicalism and Compositionality

> *The extension, hardness, impenetrability, mobility, and inertia of the whole result from the extension, hardness, impenetrability, mobility, and inertia of the parts; and hence we conclude the least particles of all bodies to be also extended and hard and impenetrable and movable and endowed with their proper inertia. And this is the foundation of all philosophy.*
>
> *Isaac Newton,*
> *Mathematical Principles of Natural Philosophy*

Introduction

Physicalism is the claim that the physical facts about the world determine all of the facts about the world. This is the dominant view in contemporary metaphysics. Here I'll examine two common assumptions related to physicalism. The first is the widespread suggestion that the content of physicalism can be adequately captured by the claim of *microphysicalism*. While contemporary physicalists no longer expect that 'macro-physical' properties such as extension, hardness, impenetrability, mobility and inertia necessarily result from microphysical properties of the same sort, it is widely assumed that physicalism is a claim specifically about the relationship between the micro-physical features of the world and all the rest. For example, Barry Loewer (2001) describes physicalism as the claim that "all facts obtain *in virtue of* the distribution of the fundamental entities and properties – whatever they turn out to be – of completed *fundamental physics*" (Loewer 2001: 37, emphasis in original). Similarly, Andrew Melnyk characterizes physicalism as the claim that "whenever any non-microphysical property is instantiated, it is in fact realized...by some more or less complex spatio-temporal configuration of instances of microphysical properties" (Melnyk 1994: 225). And David Lewis defends his own version of physicalism as the claim that all fundamental facts are facts about point-sized entities, upon which all other facts 'supervene':

> We have geometry: a system of external relations of spatio-temporal distance between points. Maybe points of spacetime itself, maybe point-sized bits of matter or aether or fields, maybe both. And at those points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. For short: we have an arrangement of qualities. And that is all. There is no difference without difference in the arrangement of qualities. All else supervenes on that. (Lewis 1986: *x*)[4]

Against this first assumption, I'll argue that while these microphysicalist claims might be *true*, they aren't adequate for the sort of work we expect physicalism to

[4] And the list goes on. Other prominent examples include Paul Humphreys' (2000) discussion of 'generative atomism' (though Humphreys does himself endorse this view, but instead 'diagnoses' it as prevalent as a vision of 'limit science'), and Philip Pettit's (1993, 1995) defense of microphysicalism (which I discuss in more detail later). Jeffrey Poland (1994) presents a comprehensive study of the concept of physicalism, in particular as it relates to the philosophy of mind; however his discussion says little about the adequacy of various accounts of physicalism for physics itself. This is the sort of gap in the understanding of physicalism the present project is meant to help fill.

do in our understanding of the world[5]. Instead, I'll argue that physicalism should be understood in terms of the other commitment expressed in the above quote from Newton, the commitment to the more general claim of *compositionality*: physicalism is best seen as a claim about the relationship between properties of wholes and those of their parts.

The second assumption I'll challenge is connected with the first: it is the assumption that the parts of physical systems are always 'spatial' parts – parts which are smaller than and spatially contained within the things they compose. I'll suggest that the physical sciences themselves appeal to a variety of sorts of parts in their explanations, including parts that are *not* smaller than the things they compose: call these *non-spatial* parts. The nature and importance of such parts will be the focus of the rest of the thesis.

I'll discuss the general idea of non-spatial parts towards the end of this chapter, and in more detail in chapters three and four; in the meantime, my main goal in this chapter is to argue that physicalism is best understood as compositionality.

Physicalism, Microphysicalism, and Compositionality

The 'microphysicalist' formulations of physicalism I mentioned above aren't the only ways of thinking about physicalism. When Otto Neurath and Rudolph

---

[5] Lewis' claim is an exception: it seems questionable whether Humean supervenience is even true, let alone adequate as an account of physicalism. The main issue concerns its compatibility with standard interpretation of quantum mechanics. Lewis himself (Lewis 1986) admits the possible conflict, but regards this is a measure of his theory's genuine empirical content.

Carnap coined the term 'physicalism', they apparently didn't have any particular claim about microphysics in mind: instead, the claim was that all phenomena could be described using only statements about things in space and time. For example, Neurath often refers to physicalism as "the formulation of order", and insists that for physicalism "it is essential that *one* kind of *order* is the foundation of all laws, whichever science is concerned, geology, chemistry or sociology" (Neurath 1983 [1931]: 54). In particular, physicalists are those who hold "that everything we can sensibly talk about is spatially and temporally ordered" (Neurath 1973: 325). Carnap (1936) describes physicalism as a commitment to the thesis of 'physicalistic confirmability', which is the claim that "every descriptive predicate of the language of science is confirmable on the basis of observable thing-predicates" (Carnap 1936: 468). A predicate is to count as an 'observable thing-predicate' just in case "for an organism (e.g., a person) N, if, for suitable arguments, e.g. 'b', N is able under suitable circumstances to come to a decision with the help of few observations about a full sentence, say 'P(b)', i.e. to confirmation of either 'P(b)' or '~P(b)' of such a high degree that he will either accept or reject 'P(b)'" (Carnap 1936: 454-455). Thus predicates such as 'red' will count as observable thing-predicates, and thus 'physicalistically acceptable, while predicates such as 'an electric field of such and such an amount' are not observable thing-predicates, since "although we know how to test a full sentence of this predicate, we cannot do it directly, i.e., by a few observations; we have to apply certain instruments and hence to make a great many preliminary observations in order to find out whether the things before us are instruments of the kind required." (Carnap 1936: 455). Clearly, predicates applying to *micro* entities would count even less as 'observable thing-predicates'.

Paul Needham (1998) attributes a view of physicalism similar in spirit to Neurath's 'formulation of order' to Pierre Duhem. For Duhem, the commitment to physicalism was intended as a commitment to a general claim of the universality of physics, "borne of a general motivation to accommodate the whole world within a single theoretical perspective, and to deal with all phenomena consistently" (Needham 1998: 42). Needham's interpretation of Duhem is particularly interesting because he is explicit in characterizing Duhem as committed to a form of 'macrophysicalism' rather than the microphysicalism now predominant in philosophy.

Contemporary views of physicalism, however, are for the most part put in terms of microphysics. One curious feature of these discussions is that often, after physicalism has been defined in terms of microphysicalism, it is often defended on the basis of examples that don't involve 'micro' parts at all. For example, Andrew Melnyk (1994) argues that 'realization physicalism' offers the best answer to what he calls 'the problem of the many sciences'. Melnyk describes this problem as follows:

> Alongside [the] profusion of sciences...there is a vaguer, more philosophical, and more intuitive thought, the thought that somehow it must be possible for the many sciences to fit together so as to constitute what is in some sense a single and unified account of the world; and possible also for the respective *domains* of the many sciences – the entities

and properties they seem to be about – to fit together so as to constitute a single and unified world. (Melnyk, 1994: 222)

Thus the primary role for physicalism is supposed to be to give some metaphysical account of the relationship between the views of the world offered by the various different sciences. According to realization physicalism, microphysical entities and properties form a common denominator linking the entities and properties of all sciences, since all entities and properties either are microphysical entities or properties or are 'functional' properties that are realized by microphysical entities and properties. Melnyk's idea is that the fact that the microphysical description of the world is the *only* description that applies globally is what matters to physicalism.

Clearly there is some sense in which the microphysical forms a common denominator of the sort Melnyk describes. However, what isn't clear is that the microphysical can play the sort of explanatory role of connecting the many sciences Melnyk imagines. Consider Melnyk's account of the importance of realization physicalism for understanding cases of 'cross-scientific causation', where "causes and effects are events described in the language of different sciences which seem…to be at roughly the same 'level' as one another" (Melnyk 1994: 235). Melnyk's example is the causal relationship between taking a pill of certain type and the resulting change in psychological states: without the understanding of inter-scientific relationships physicalism is supposed to allow us to find, "we will…possess some causal generalizations to the effect that, for instance, ingestion of aspirin-tablets often relieves pain, but we will not be able to specify in any more detail the exact conditions under which it does so, or the ways in which it might be prevented from doing so or encouraged to do so" (Melnyk 1994: 235-6). Appreciating physicalism is supposed to allow us to envision a 'microreduction' of this causal relationship, where "we can think of both aspirin-tablets *and human beings* as composed of certain chemical substances…we can then apply our knowledge of *chemistry* to the observed interactions and try to determine more precisely the effects of aspirin on humans" (Melnyk 1994: 236). Thus the commitment to physicalism enjoins us to seek common denominators and "to analyse two apparently disparate kinds of thing (aspirin-tablets and human bodies) into lower-level elements (chemical substances) whose interactions we understand reasonably well." (Melnyk 1994: 236).

Melnyk's example is surely correct in some sense, but what is striking is his own admission about its significance. Realization physicalism is supposed to assert relationship between facts and specifically *micro*-level facts, yet:

> This example, admittedly, is perhaps not quite of a microreduction within the meaning of the act, since it does not involve analyzing superficially disparate phenomena into lower-level *microphysical* elements. But it nevertheless illustrates the point – or at least a point – of microreductions, which is to hypothesise some lower-level commonality between different kinds of event or thing, with a view to increasing and precisifying our knowledge of the causal interactions between them. (Melnyk 1994: 236)

Rather than a specific claim about the fundamentality of the microphysical, what Melnyk really seems to have in mind is a more general claim of *compositionality*: the properties of wholes depend upon those of their parts, and thus explanations of 'cross-scientific causation' appeal to common denominators by finding common decompositions. But nothing requires that these decompositions be *micro*-physical.

This equivocation between a commitment to microphysicalism and a more general claim of compositionality is not new to accounts of physicalism. For example, nearly half a century ago, J. J. C. Smart described physicalism as the claim that "so far as science is concerned, [there is] nothing in the world but increasingly complex arrangements of physical constituents" (Smart 1959: 142) and consequently that one day we can expect even explanations of human behavior to be based on the future fundamental laws of nature applying to "whatever ultimate particles are then in vogue" (Smart 1959: 143). Clearly this is a microphysicalist claim. Consider, however, how Smart presents the apparent problem for physicalism and 'the mental' (a problem he goes on to dispel):

> There does seem to be, so far as science is concerned, nothing in the world but increasingly complex arrangements of physical constituents. All except one place: in consciousness. That is, for a full description of what is going on in a man you would have to mention not only the physical processes in his tissue, glands, nervous system, and so forth, but also his states of consciousness" (Smart 1959: 142)

The difference between the avowed commitment to physicalism (to the fundamentality of 'ultimate particles') and examples intended to illustrate the unproblematic cases for the physicalist – physical processes in tissues, glands, nervous systems, etc – is striking: talk of tissues, glands, nervous systems, and so forth is a long way from talk of elementary *particles*. It seems much more reasonable to understand Smart's real interest as a commitment to compositionality: of course, we expect that tissues, glands, and nervous systems *have* elementary particles as parts, but their relationship with *those* parts is not all there is to 'physicalism'.

A similar idea is found in Hartry Field's (1992) argument for a methodological role for physicalism in science. While Field insists that it is "beyond serious doubt that there is an important sense in which all facts depend on physical facts and all good causal explanations depend on good physical explanations" (Field 1992: 271), his account of 'reductive sketches' suggests that it is ultimately a relationship between explanation in terms of 'wholes' and explanations in terms of parts that is at issue. For Field, physicalism is supposed to guide us to find physical foundations for any acceptable causal explanation, and to lead us to reject otherwise acceptable explanations if no physical foundation is reasonably conceivable. Understood as a claim about *microphysicalism*, this becomes quite implausible: it would be quite surprising for most scientists – biologists studying frog populations in central Alberta, for example – had any idea of how their theories relate to the theories of particle physics. Field concedes as much, admitting that while we might expect all good facts and explanations to depend on physical ones, "if the facts and explanations are sufficiently 'high-level,' we will not look directly for a *physical* foundation:

we will simply look for a foundation in terms of 'lower-level' facts and explanations" (Field 1992: 271).

Understanding physicalism as a general claim about compositionality makes much more sense of these examples than does microphysicalism. For instance, in Melnyk's example of trying to understand the effect of taking some drug, it clearly *is* a good strategy to try to analyze apparently disparate kinds of things into lower-level commonalities. But that strategy isn't the result of a commitment to microphysicalism – instead it is the product of a commitment to compositionality. Compositionality says that the properties of wholes are determined by those of their parts, but it doesn't say *which* parts might be relevant for understanding a given property.

The idea of understanding physicalism as compositionality is developed most prominently in Oppenheim and Putnam's seminal paper "Unity of Science as a Working Hypothesis". With concern about 'the many sciences' similar to Melnyk's, Oppenheim and Putnam defend the view that science can and should be seen as a unified project, where the different sciences aim to describe reality at different 'levels' of analysis. Such levels are to correspond to distinct domains of scientific inquiry, where a given level is characterized by a set of types of entities and a group of predicates drawn from scientific theories describing those entities. The role of compositionality lies in characterizing these different levels: the levels themselves are supposed to be related by the part-whole relation, so that objects found at one level are always composed of objects found at the next lowest level. But not just any decomposition will count as a distinct level. In particular, Oppenheim and Putnam stipulate that levels must be in some sense 'natural' and "justifiable from the standpoint of present-day empirical science" (Oppenheim and Putnam 1958: 9). The division in to levels they suggest is:

6...Social groups

5...(Multicellular) living things

4...Cells

3...Molecules

2...Atoms

1...Elementary particles

The claim of scientific unity derives from the thought that explanations at a given level might 'reduce' to explanations at the next lowest level, in the sense that lower level explanations of the same phenomena might allow one to *replace* the higher-level account with the lower-level one[6]. Note that unlike other prominent sorts of reduction, such as that develop by Ernest Nagel (1961), the account suggested by Oppenheim and Putnam does not require any *logical* connections between laws of different theories. For instance, their claim is not that higher-level laws must be *derivable* from lower-level ones. Instead, one theory reduces to another just in case the 'observational data' explainable by the one (the reduced theory) is explainable by the other (the reducing theory), and the reducing theory is 'better systematized' than the reduced theory. In a sense, one

[6] This general account of reduction is developed in Kemeny and Oppenheim 1956. The account of reduction defended by Oppenheim and Putnam differs significantly from the more standard concept of 'Nagelian' reduction, as developed by Ernest Nagel in (Nagel 1961). I'll return to the topic of reduction in chapter two.

theory reduces to another just in case the reducing theory is a *better theory* of the phenomena to be explained.

More specifically, Oppenheim and Putnam hypothesize the existence of what they call *micro*-reductions between explanations at different levels. The term micro-reduction is somewhat unfortunate: what really matters in the account of reduction they give is the relationship between properties of wholes and those of their parts, regardless of whether or not those parts are 'micro'. For example, they define the reduction of one 'branch' of science to another as follows:

> [L]et us suppose B_2 is a branch of science that has multicellular living things as its universe of discourse. Let B_1 be a branch with cells as its universe of discourse. Then the things in the universe of discourse of B_2 can be decomposed into proper parts belonging to the universe of discourse of B_1. If, in addition, it is the case that B_1 reduces B_2 at the time t, we shall say that B_1 micro-reduces B_2 at time t. (Oppenheim and Putnam 1958: 6)

Oppenheim and Putnam's 'working hypothesis' is that the sciences can be unified by such reductions, and the purpose of calling attention to different levels is both to suggest where reductions ought to be sought, and to show how the availability of such reductions would lead to overall unity in science. Such reductions are supposed to form steps towards 'unitary science' since if micro-reductions are available, then the terms and laws of the *reduced* branches of science can in principle be discarded since, by definition, a reducing theory explains all phenomena explained by the reduced theory. Of course, there might still be practical reasons for retaining 'reduced' branches or theories, such as the prohibitive complexity of trying to construct all explanations at the level of the reducing science, but in principle those higher-level accounts would no longer be necessary.

The idea that each level represents a 'crucial step' in an over-all reduction is that while it is clearly unreasonable to expect even a sketchy account of how observational data explainable in terms of, say, cells could be replaced by explanations in terms of elementary particles (which would be required for an reduction of level 4 to level 1), it is feasible – or so Oppenheim and Putnam suggest – that explanations in terms of cells could be replaced by explanations in terms of molecules, and explanations in terms of molecules be replaced by explanations in terms of atoms, and so on. Since parthood and reduction are both assumed to be transitive, micro-reductions themselves are thought to be transitive[7]:

---

[7] There is a curious point to note here about the part-whole relation. The concept of parthood, which is crucial to the account presented by Oppenheim and Putnam, is not developed by them in any detail. Instead, they cite favourably the work done on the parthood relation by Nicholas Rescher in two papers, one in collaboration with Oppenheim. While the Rescher and Oppenheim paper contains valuable work on the nature of decompositions and the relationship between properties of parts and those of wholes, it contains little on the parthood relation itself. Rescher's own short paper on this topic – Rescher (1955), "Axioms for the Parthood Relation" – cited by Oppenheim and Putnam – sets as its goal an axiomatization of the parthood relation more in line with the concept of parthood found in science (as opposed to standard axiomizations of 'mereology' developed by Lesniewski). However, one of the key features of Rescher's account of parthood is that he (partly) rejects transitivity. This is striking, since transitivity plays such a key role in the account of Unity of Science defended by Oppenheim and Putnam. I return to Rescher's argument briefly in chapter four.

> [This] *formal* property of the relation 'micro-reduces' – its transitivity – is of great importance for the program of Unity of Science. It means that micro-reductions have a *cumulative* character. That is, if a branch B_3 is micro-reduced to B_2, and B_2 is in turn micro-reduced to B_1, then B_3 is automatically micro-reduced to B_1. (Oppenheim and Putnam 1958: 7)

Oppenheim and Putnam argue that this transitivity is what foils those arguments against the idea of 'unitary science' that depend upon the implausibility of explaining higher-level phenomena (such as the psychological) in terms of elementary physics. Finding 'level-wise' reductions across the various intermediate levels will mean that "psychological laws will have, in *principle*, been reduced to laws of atomic physics, although it would nevertheless be hopelessly impractical to try to derive the behavior of a single human being directly from his constitution in terms of elementary particles" (Oppenheim and Putnam 1958: 7).

It's true that Oppenheim and Putnam regard the 'lowest' level – that of 'fundamental physics' as particularly important for their account. It is the indirect reducibility of all phenomena to this bottom level that is supposed to make their account a 'physicalistic' one, conducive to their goal of the an "overall physicalistic reduction" of the sciences. However, the nature of this bottom level plays no real role in their argument: what matters for their account of scientific unity is not the 'physicality' of the bottom level, but the fact that the domains of different branches of science are related by the part-whole relation. What matters, that is, is compositionality.

Aren't Microphysicalism and Compositionality the Same?

One reason for not distinguishing between microphysicalism and compositionality is the thought that they are equivalent, since 'parts' of physical systems are always *spatial* parts. If the properties of wholes are determined by those of their parts, then considering smaller and smaller parts will eventually lead us to consider 'microphysical' parts; by compositionality, the facts about these will determine the facts about everything in the world. That is, microphysicalism will be true.

Later in the chapter, I'll suggest why we might challenge the assumption that the parts of physical systems *are* always spatial parts, but for now I'll show that even if that assumption about physical parts *were* true, microphysicalism and compositionality would not be equivalent.

To see this, suppose there were two brains, and – putting aside any worries about the extreme unlikelihood of the situation – those brains were identical at the micro-level. Microphysicalism tells us that those brains would also be identical in all of their intrinsic properties: they would be chemical identical, biologically identical, psychologically identical, and so on. This is the result the physicalist should expect. However, now suppose we have two brains that *differ* at the micro-level, and yet are identical at some intermediate level, for instance in terms of their neural features: each is an instance of exactly the same 'neural network' with the same patterns of stimulation and (neural) activity. What can we say about the 'higher-level' properties of such brains, such as their psychological properties? It seems quite plausible to expect that such brains would be identical in terms of their higher-level properties. Certainly, much of cognitive science and neuroscience in particular depends upon such an assumption. Yet in this case – because the two brains *differ* at the micro-level – microphysicalism imposes no constraint on their higher-level properties at all. Microphysicalism would be consistent with such brains differing entirely in their psychological properties. Psychological properties could still supervene on microphysical properties without supervening on neural properties. But without the supervenience of the psychological on the neural, what methodological principle could direct us to consider explaining higher-level properties and behavior at the neural level in the first place? If our commitment were simply to microphysicalism, it seems we could only be justified in expecting to find explanatory commonalities at the *micro*-level. But clearly we expect more than that. Let's call this the problem of *higher-level variation*, since the worry is that microphysicalism allows for higher-level properties to vary in situations where we expect them to be similar.

This argument depends on two related assumptions which are not irrefutable but which are widely accepted. The first is that higher-level properties are not *identical* with micro-level properties: if we adopt a classical 'type-type' identity theory of reduction, then the sort of situation described above will not be possible. If higher-level properties were *identical* with micro-level properties, then nothing could be isomorphic with respect to one sort of property without also be isomorphic with respect to the other. I've already noted that I won't argue *against* reduction here, but note that clearly many people defending microphysicalism assume that this in some way captures the idea that the world can be described at different levels. For instance Philip Pettit, commenting on his own account of microphysicalism, claims that

> the physicalism defined here is a concrete but fairly cautious version of an abstract and plausible claim...that the various kinds of things in the world are composed of less varied items, that this composition establishes a hierarchy of different levels of thing, and that there is some less-than-highest level of composition such that if we fix how things are governed from there down, then we shall have fixed how things at every level are governed. (Pettit 1995: 405)

And this despite the fact that "a physicalist need not imagine that it is going to be possible for human beings, at least under some idealisation of their capacities, to reduce macro-level laws to microphysical conditions and regularities" (Pettit 1993: 219)[8].

One reason Pettit gives for thinking that such reductions might fail is because macro-level properties might be "multiply and wildly realizable...at the micro-physical level" (Pettit 1993: 219). This multiple realizability is the second

---

[8] Pettit's views are slightly ambiguous. Alongside this apparent suggestion of a non-reductive view of physicalism, he characterizes the lesson taught by his definition of physicalism as follows:

> What [Pettit's account of physicalism] means is that when we involve ourselves in the special sciences then we track the laws and controls – the real and important laws and controls – that an ideal microphysics would address in a different idiom. What it means is that research in the special sciences – research in chemistry and biology, psychology and social science, even research of a common-sense type – can be seen as microphysics by other means. (Pettit 1995: 420)

I don't see how this claim can be meaningfully understood other than as exactly what the non-reductive physicalist wants to deny.

assumption made in the higher-level variation argument. The idea of multiple realizability is a prominent reason for rejecting accounts of reduction involving identity: in most cases it seems clear that higher-level properties *can* be multiply realized and thus they *cannot* be identified with any one particular realizing property. The classic examples of multiple realizability concern mental states (as my example reflects): for instance, while it might seem plausible that there is some particular neural state or pattern of activity that precisely coincides with my *being-in-pain*, it seems less plausible that that same state or pattern of neural activity also coincide with another person's *being-in-pain* – different people, after all, have different brains, and it's possible that there is some variety in ways in which different people's brains 'behave' when they are in pain. Even more convincingly, it seems entirely reasonable to think that animals from other species (a dog, for instance) can also experience pain, despite the fact that their brains are definitely distinct in their physical structure from mine. And it is even conceivable that an entirely distinct sort of creature (a life form on some other planet, for instance) could also experience pain, despite having a brain of a completely distinct sort from mine. Thus the same mental state, being-in-pain, is thought to be 'multiply realizable' as a variety of distinct physical states. This multiple realizability is supposed to contradict the reductionist claim: a higher-level property cannot be *identical* with a particular lower-level property if there are a variety of lower-level properties that can realize the higher-level one.

The multiple realizability of mental states was long taken for granted, and the idea has recently come under increased scrutiny (see Shapiro 2000, for instance). However, multiple realizability is not unique to psychological properties, and many macro-*physical* properties also seem to be multiply realizable. For example, consider the relationship between fluid-mechanical properties of bodies of fluids and the micro-level properties of the molecules that compose those bodies. Like thermodynamics, fluid mechanics was initially developed in the nineteenth century, before the establishment of what we now call the atomic view of matter. Whether or not the laws and principles of fluid-mechanics apply to a body depends on certain 'macro' features of that body: for instance, a fluid is assumed to be continuously deformable. Fluids come in an enormous variety: water, whiskey, milk, etc. But the micro-level structure of those substances is in general irrelevant to their fluid mechanical properties. The relevant properties for fluid mechanics – such as viscosity, for example – might be explainable in terms of the micro-level structure of a fluid, but as far as fluid mechanics is concerned, any other micro-level structure that gave rise to those some macro-level characteristics would be equivalent. So the same fluid mechanical properties could be 'multiply realized' by distinct micro-level structures.

Granting these two assumptions, the problem of higher-level variation makes it clear that microphysicalism is insufficient for capturing the sort of dependencies we can capture with a more general commitment to compositionality: so the two concepts are distinct. Appreciating this difference also allows us to address a prominent worry about 'supervenience' formulations of physicalism, which claim only that no two individuals can differ with respect to any property without differing with respect to some physical property. Hartry Field (1992) argues that supervenience claims are too weak to play the sort of methodological role we expect physicalism to play, and hence that physicalism must involve a 'neo-classical' account of type-type reduction. His main reason for thinking this is that since supervenience only constrains the characteristics of individuals identical in terms of their 'subvenient' properties, it is exceptionally unlikely that any individuals in the actual world *are* identical in this way. This is clearly right if one assumes microphysicalism: microphysical identity is a rare thing. As in the case of higher-level variation, a belief in the supervenience on the *micro* wouldn't give us any methodological direction in evaluating our theories. But while it's unlikely that many things are identical in terms of their micro parts, such isomorphism becomes much more plausible the more 'coarse grained' the decomposition we consider. Hence it's much more reasonable to expect that many things *are* isomorphic on *some* decomposition, and thus fall within the purview of compositionality.

Supervenience claims represent the minimal claim of physicalism: any claim of physicalism must *at least* insist on supervenience. The claim of supervenience doesn't insist on lawful connections between supervenient properties and their subvenient base properties, and it is not supposed to entail that explanations in terms of supervenient properties can be *replaced* by explanations in terms of base properties, but only that the supervenient properties depend on subvenient ones. We can formulate a claim of compositionality analogous to the claim of micro-reduction suggested by Oppenheim and Putnam by modifying what Jaegwon Kim calls 'mereological supervenience', which is a dependence relation between parts and wholes. Kim's version of mereological supervenience is as follows:

> *Mereological Supervenience*: For any x and y, belonging to level L (other than the lowest level), if x and y are indiscernible in relation to properties at all levels lower than L (or, as we may say, x and y are *mereoindiscernible*), then x and y are indiscernible with respect to all properties at level L.
>
> *Mereoindiscernibility*: x and y, belonging to level L, are *mereoindiscernible* if and only if for every decomposition D of x into proper parts belonging to lower levels, y has an isomorphic decomposition C in the sense that there is a one-one function I from D to C such that for any n-adic property or relation P at levels lower that L, $P(\mathrm{d}_n)$ iff $P(I(\mathrm{d}_n))$, where d_n is any n-tuple of elements in D and $I(\mathrm{d}_n)$ is the image of d_n under I, and conversely from y to x. (Kim 1998: 17)

Note that Kim's formulation tacitly assumes a restriction to the *intrinsic* properties of an individual – the properties an individual may instantiate independent of its relations to any other individuals. Clearly, an individual's *extrinsic* properties will not depend solely on the properties of *its* parts. In the discussion that follows – and in fact throughout the thesis – I follow Kim in focusing on intrinsic properties, though for convenience I will omit explicitly repeating this qualification.

I've made a slight terminological change in Kim's definition: while Kim calls his claim one of *mereological* supervenience (Kim 1998: 18), suggesting that it is specifically a claim about supervenience between parts and wholes, the definition he actually gives is put in terms of *microphysical* supervenience and *micro*-indiscernibility. Oppenheim and Putnam made the same choice in terminology (and perhaps Kim's choice simply reflects theirs): despite the fact that their definitions of supervenience or reduction refer only to parthood, their choice of names reflects the assumption that it is specifically the *micro*-physical parts that matter for physicalism. In fact, given Kim's choice of quantifiers, these *are* the only parts that matter. On his definition two things should be indiscernible just in case they are isomorphic on decomposition to *all* lower levels. This means that Kim's supervenience claim effectively only constrains things that are indiscernible at the lowest level to be themselves indiscernible – similarities at higher-levels are irrelevant. This means that despite appearances, Kim's formulation doesn't really accommodate the idea of there being different levels of decomposition that might be relevant to a particular type of property: Kim's formulation is really the more specific claim of microphysicalism than a general one of compositionality.

We can make Kim's claim more appropriate for compositionality by introducing a few modifications. The fundamental idea of compositionality is that the properties of wholes are determined by those of their parts. Kim's 'mereological supervenience' expresses this in terms of decompositions at different 'levels', but since the classic account of levels is given in terms of parthood, we can simply talk directly of *decompositions*. A decomposition of an individual or a system of individuals is a specification of the parts of that system, their properties, and the relations between them. Our claim is that for any (intrinsic) property of an individual, there is a decomposition of that individual such that anything with an isomorphic decomposition also has that property. In terms of supervenience we can define:

> *Compositionality*: for any individual x and property P such that Px, there is a decomposition D of x such that for any individual y, if y is mereologically indiscernible from x on D, then Py.

Where mereological indiscernibility is given by modifying Kim's earlier definition to give:

> *Mereological Indiscernibility*: x is *mereologically indiscernible* from y on a decomposition D just in case D is a decomposition of x and y has an isomorphic natural decomposition C in the sense that there is a one-one function I from D to C such that for any n-adic property or relation P, $P(d_n)$ iff $P(I(d_n))$, where d_n is any n-tuple of elements in D and $I(d_n)$ is the image of d_n under I, and conversely from y to x.

Since we don't expect just any decomposition to be relevant to the supervenience claim, we need to distinguish the relevant ones: call this 'natural' decompositions. Oppenheim and Putnam make a similar distinction with regard to levels: they require every level to be 'natural' in some sense and "justifiable from the standpoint of present-day empirical science." (Oppenheim and Putnam 1958: 9) The intuitive idea of a natural decomposition is one where the components of that decomposition fall under natural laws: natural decompositions 'carve nature at the joints' and display the underlying commonalities that are described by the laws of nature[9]. I'll return to this idea of naturalness in chapter four.

Before saying more about the idea of compositionality and the nature of physical parts, there are two prominent arguments in favor of microphysicalism that I should address. These are arguments for thinking that physicalism *must* be understood in terms of microphysicalism, but I will argue that this is false.

The Meaning of 'the Physical'

Oppenheim and Putnam's focus on the 'physicality' of the lowest level reveals a common assumption about physicalism. This assumption is that it is at the micro-level that entities and properties are *paradigmatically* physical. This is one reason why many philosophers think that physicalism can only be meaningfully defined in terms of 'microphysicalism'. Microphysicalism, it is widely assumed, is the best – and perhaps the only – way to answer an obvious question about what physicalism *means*. Since this problem – and the assumption that

---

[9] One specific constraint we would expect to be associated with the idea of a natural decomposition concerns different sorts of 'structural' properties. In their account of levels, Oppenheim and Putnam prohibit the mention of properties concerning the 'super-structure' (the term is mine, not theirs) of entities at a level. For example, if we were allowed to attribute atoms with the property 'Tran' ("the property of being an atom of a transparent substance"), then finding a microreduction of the phenomenon of transparency would be trivial (see p. 10). This needn't mean that there *aren't* such 'super-structural' properties, but only that they can't be included in natural decompositions, since otherwise compositionality will be trivial.

microphysicalism provides the best solution – is quite prominent, I'll discuss it in some detail.

If the general claim of physicalism is that 'the physical facts determine all of the facts', then this invites the question "what distinguishes *physical* facts from all the rest?" Some philosophers reject physicalism on the grounds that talk of 'physical' facts is meaningless unless some account can be given of what distinguishes the 'physical' from the 'non-physical'. The problem is supposed to be that specifying what is meant by 'physical' is not just difficult, but is in fact impossible to do without trivializing the claim of physicalism. This sort of criticism of physicalism is most vigorously pursued in Tim Crane and D. H. Mellor's 'anti-physicalist manifesto', "There is No Question of Physicalism" (1990). Crane and Mellor argue that claims of physicalism are either obviously false or trivially true, and thus that a commitment to physicalism makes no demands on our understanding of the 'non-physical' sciences such as psychology or sociology. Their argument laid down a tacit challenge to other philosophers to come up with a suitable 'definition' of physicalism.

The heart of the challenge Crane and Mellor present is the thought that it is either inappropriate or perhaps even impossible to give any sort of *conceptual* criteria of what it means for something to count as a 'physical' entity, property, or fact. This unwillingness to go out on a conceptual limb is one distinction between the contemporary claim of physicalism and its ancestor, seventeenth-century *materialism* or 'mechanical philosophy'. As Crane and Mellor put it, physicalism is the less ambitious descendent of materialism, which held not only that all of the facts about the world (the material world, at any rate) could be explained by physics, but also dared to say just what sorts of explanations we could expect to find in physics:

> In its seventeenth-century form of mechanism...materialism was a metaphysical doctrine: it attempted to limit physics a priori by requiring matter to be solid, inert, impenetrable and conserved, and to interact deterministically and only on contact. But as it has subsequently developed, physics has shown this conception of matter to be wrong in almost every respect...Faced with these discoveries, materialism's modern descendents have – understandably – lost their metaphysical nerve. No longer trying to limit the matter of physics a priori, they now take a more subservient attitude: the empirical world, they claim, contains just what a true complete physical science would say it contains (Crane and Mellor 1990: 186)

Thus materialism was a substantial, though ultimately false, claim, and the challenge for physicalism is to say something of similar substance without being similarly false.

We might think we could define physicalism simply by referring to the entities and concepts used by actual theories in *physics*. However, this sort of approach is also problematic. The difficulty concerns saying *which* theories of physics should be used as a reference point for defining 'the physical'. This problem has become known as "Hempel's Dilemma", Carl Hempel being one of the most prominent philosophers to have raised it[10]. Here is Hempel's original statement of the dilemma:

> [T]he physicalistic claim that the language of physics can serve as a unitary language of science is inherently obscure: The language of *what* physics is meant? Surely not that of, say, 18th century physics; for it contains terms like 'caloric fluid', whose use is governed by theoretical assumptions now thought false. Nor can the language of contemporary physics claim the role of unitary language, since it will no doubt undergo further changes, too." (Hempel 1980: 194-5).

Hempel's version doesn't actually express the contemporary concern, largely because physicalism is no longer seen as a claim about the *language* of science. We can recast Hempel's concern in a less linguistic form by asking the same question about what facts are to count as physical. Are they the facts typical of contemporary physics – facts about entities such as electrons, photons, neutrinos, quarks, together with properties such as *mass*, *charge*, *spin*, and 'laws' such as the law of universal gravitation (relativistic or not)? If so, then it appears unlikely that physicalism is true, though not for the reason Hempel suggests.

---

[10] The concern expressed in the dilemma can be found elsewhere. For instance, of 'mechanical explanation', the physicist Ludwig Boltzmann writes: "If by a mechanical explanation of nature we understand one that rests on the laws of current mechanics, we must declare it as quite uncertain whether the atomism of the future will be a mechanical explanation of nature. Only insofar as it will always have to state the simplest possible laws for temporal change of many individual objects in a manifold of probably three dimensions, can it be called a mechanical theory, at least in a metaphorical sense. If it should for example turn out to be impossible to find a simpler description of electromagnetic phenomena, one would have to retain the vector atoms discussed in the text above... *Whether the laws according to which these changes with time are to be called mechanical or not will be entirely a matter of taste.*" (Boltzmann 1974: 52).

True, contemporary physics no longer contains the term 'caloric', nor makes reference to facts about caloric, but *that* isn't the reason eighteenth century physics seems inadequate to express all of the physical facts about the world. Instead, the problem is that contemporary physics contains terms and assumptions not expressible in (and perhaps inconsistent with) the language of eighteenth century physics. The language of 'past physics' is thus insufficient to express the facts according to contemporary physics. Thus if by 'physical' we *did* mean "expressible in terms of 18^{th} century physics", physicalism would pretty clearly be false: in fact, the evidence for its failure would come from physics itself.

The same problem is thought to plague contemporary physics: based on past experiences in the development of science, we have good reasons to believe that the physics of the future will contain terms not found in contemporary physics. Thus if by 'physical' we meant "expressible in terms of contemporary physics", then it again seems likely (though not certain) that physicalism would be false, since we expect that the physics of the future will not be entirely expressible in terms of contemporary physics. The 'physical' facts in terms of present-day physics wouldn't determine all of the facts about the world.

On the other horn of the dilemma, if we associate the claim of physicalism with some 'future' physics that isn't susceptible to the sorts of changes we expect in contemporary physics, physicalism becomes unacceptably vague. Who knows what sorts of entities, properties, or laws future physics might include? Most dramatically, it is *conceivable* that future physics could include irreducibly *mental* entities, perhaps in something like the way that some philosophers assume contemporary quantum physics does[11]. We needn't take this possibility very seriously in order for the dilemma to have its bite: the problem does not revolve around the likelihood of mental entities proving to be irreducible to physics, but instead concerns the thought that any reasonable doctrine of physicalism ought to *reject* such a possibility. After all, physicalism might well turn out to be false, and it seems that one of the best ways for it to be shown to be false would be for such developments in physics actually to occur. But on a 'futurist' definition of physicalism, such developments would *not* make physicalism false. With David Lewis, we might share the vague confidence that "we may reasonably hope that future physics can finish the job [of providing an inventory of the fundamental entities, properties, and laws of nature] in the same distinctive style" (Lewis 1994: 412), but unless we have a firmer grasp of just what distinguishes that style, we cannot have confidence that 'future physicalism' is a non-trivial doctrine[12].

---

[11] For example, David Lewis (1986a) argues that though there is an apparent conflict between his 'Humean supervenience' and orthodox quantum mechanics, this need not count against his view, despite his general commitment to the primacy of physics. Before he is "ready to take lessons in ontology from quantum physics", Lewis insists that quantum physics be (among other things) "purified of supernatural tales about the power of the observant mind to make things jump" (Lewis 1986a: xi).

[12] Noam Chomsky expresses the worry that future physics might simply appropriate whatever concepts it needs as 'physical', noting that "...the material world is whatever we discover it to be, with whatever properties it must be assumed to have for the purposes of explanatory theory. Any intelligible theory that offers genuine explanations and that can be assimilated to the core notions of physics becomes part of the theory of the material world, part of our account of body." (Chomsky 1987) I should note that Chomsky himself isn't *worried* by this at all: his point is precisely that one can't reject work in the cognitive sciences on the grounds that it doesn't provide 'physical' explanations: just what counts as a 'physical' explanation may well have to be modified to *include* the types of explanations offered in cognitive science.

Thus the physicalist is left stranded between these two horns: the fact that we expect physics to evolve in the future prevents us from identifying the 'physical' with the terms of contemporary physics, and the fact that we cannot say just *how* it will evolve prevents us from identifying the 'physical' with the terms of some future physics. The challenge of Hempel's dilemma is to provide a characterization of a science we are at present unable to formulate exactly.

There have been various attempts to overcome Hempel's dilemma. For example, Andrew Melnyk (1997) suggests simply biting the bullet over the dilemma and insisting that physicalism *is* a claim about contemporary physics, despite our suspicion that the concepts of contemporary physics will be revised in the future. Melnyk argues that in general we do not only endorse theories we think are probably true, but also endorse theories we think are more likely to be true than any of their ('relevant') rivals. No matter how likely we think 'currentist' physicalism is to be wrong, it is still more likely to be true than rival claims such as vitalism or dualism – or at least, so says the physicalist. This is undeniably a clever answer (and Melnyk's account is a good deal more nuanced than I have recounted); however, it is unlikely to satisfy anyone genuinely interested in physicalism As Crook and Gillett (2001) point out, we could argue that currentism fails to be a credible hypothesis even on Melnyk's own standards. Let P represent the claim of Currentism – that all facts are determined by the facts of current physics – and let $P+$ be currentism modified to the include the claim that there is one as yet undiscovered particle with familiar properties such as mass, charge, spin, etc. Crook and Gillett insist that given the history of physics in the past century, $P+$ is *more likely* to be true than P. $P+$ is more likely to be true given the fact that the population of the 'particle zoo' – the variety of entities studied by particle physicists – has tended to expand over the past century, and it may be thought probable that it will continue expanding. But then by Melnyk's own standards it is $P+$ we should endorse rather than P, meaning that we should *reject* physicalism as Melnyk formulates it[13].

Crook and Gillett's criticism of Melnyk is that by associating physicalism too closely with a particular list of inhabitants, he makes *any* developments in physics falsify physicalism. But as physics develops, it should be an open question (of some sort) whether physicalism remains a viable theory. One way to highlight this is to consider not just the ways we might expect physics to develop (by the discovery of new particles, for instance), but instead focus on a broader 'metaphysical' claim about what sorts of 'possible worlds' should count as

---

[13] Another 'currentist' approach is suggested by Alvin Plantinga in his commentary on Bas van Fraassen's (1996) arguments about the vacuity of the idea of 'materialism'. Van Fraassen takes the fundamental claim of 'materialism' to be that *all is matter*: the dilemma then ensues over just what it takes for something to count as 'matter'. Plantinga's suggestion is that we amend this definition of materialism to read: "all is matter, and matter is what current science says there is, *together with anything sufficiently similar* to what current science says there is" (Plantinga 1996: 350).

J. J. C. Smart also defended a version of currentism: in discussing his view of physicalism, Smart acknowledges that "to prevent [the claim of physicalism] from degenerating into a triviality, I must rule out the suggestion that psychic forces or entities, or forces that apply only to complex configurations such as brains, might one day be absorbed into physics. For this purpose, I take 'physics' to mean 'present-day physics'" (Smart 1981: 109). Note that Melnyk and Smart's concerns are somewhat different: Melnyk's worry appears to be physicalism is *vacuous* unless it involves claims about specific entities, while Smart's concern seems more specifically directed at the possibility, as suggested by Chomsky, that future physics might not just invoke *different* concepts, but intuitively 'mental' ones.

physicalistic. Suppose ours does. Surely we would want to count as physicalistic other worlds that might differ quite drastically in terms of their fundamental particles. For example, consider a world pretty much like ours except that what we commonly call 'atoms' really are indivisible atoms: in this other possible world, hydrogen, helium, and so on truly are the simplest forms of matter. I presume that it is *possible* that such a world exist[14]. There will be significant differences between such a world and ours: in particular, whatever phenomena provoked physicists to speculate that atoms were *not* truly atomic will presumably not occur in such a world. But I presume – perhaps wrongly – that there could also be substantial overlap between that world and ours. According to the spirit of physicalism, such a world should still count as 'physicalistic', but on Melnyk's account it would not[15].

This worry about Hempel's dilemma has led others to insist that rather than tying the definition of physicalism to the changing tastes of physics, we *can* find a conceptual distinction sufficient to give the claim of physicalism some bite. David Papineau (2001) has argued that we can avoid worries about what counts as 'physical' by focusing instead on what counts as 'non-mental', which he assumes is a much more stable concept than that of the physical: Carl Gillett and Gene

[14] Perhaps there are some 'Kripkean' difficulties in calling these 'true' elements 'hydrogen', 'helium', etc. on the grounds that the referents of those terms *have* atomic structure; no matter, call the other worldly elements 'trans-hydrogen', 'trans-helium', etc. Perhaps hydrogen *couldn't* have truly been atomic, given that it is in fact not; still, a world with atomic stuff very much like hydrogen is surely possible, and such a world should equally count as a physicalistic one.

[15] To make this point even more emphatically, consider such an 'atomic' possible world that is so simple – perhaps by only containing one or two atoms – that there are no non-physical facts. Surely physicalism should be trivially true, but on Melnyk's account it would not be.

Witmer (2001) dub this approach to defining physicalism the 'via negativa'. This approach tries to meet Crane and Mellor's challenge by noting that while it is true that it turned out that seventeenth-century materialists were wrong about matter being inert, impenetrable, conserved, and so on, surely they were not *very* wrong. Physics has had to give up some of these ideas about the nature of matter, but denying that matter is inert and impenetrable is a long way from claiming that there are irreducible psychological entities, biological forces, entelechies, or souls in the world. The 'via negativa' allows that any list of physical entities is likely to be revised in the future, but insists that such entities as 'minds' (or whatever) will never appear on such a list.

This approach might be appropriate if our general concern is simply with the relationship between mental or psychological properties and non-mental properties. However, it is less obvious that it will apply more broadly if our concern with physicalism encompasses the relationship between the chemical, the biological, the geological, etc. and the physical. One problem is that in some of these areas of science, it seems much less obvious that we can determine a set of proprietary facts unique to that science: for example, if we are interested in geology, the relevant geological facts are not so obviously 'non-physical'. Does *mass* count as a physical feature or a geological feature? Intuitively, it should count as a physical feature: if geological facts (such as a type of rock's tendency to fracture) were determined by features including *mass* (but excluding fracture), we should be satisfied that geology was 'physicalistic'. But, in the absence of a characterization of physical facts, how would we know that this was a victory for physicalism?

There is another problem in generalizing this approach from the case of the mental and the physical to physicalism more broadly construed. Papineau claims that the generalization is straight-forward: "The same point applies if we want to apply the causal argument to chemical, biological, or economic states. As long as we can be confidant that all nonchemical effects are fully caused by nonchemical (nonbiological, noneconomic...) states, then we can conclude that all chemical (biological, economic) states must be identical with something non-chemical (nonbiological, noneconomic...)" (Papineau 2001: 13). But not only does it seem unlikely that each of these non-physical categories is stable in a way that 'physical' is not, capturing physicalism in its broadest sense would require a complete list of all 'non-physical' sciences. For instance, the claim that all non-chemical effects are fully caused by non-chemical states may well be true, but that in itself doesn't rule out the possibility that non-chemical states are caused by (say) mental states. Unless we formulate the 'via negativa' as a claim about the *conjunction* of the non-chemical, the non-biological, etc., we can't really capture the view we intuitively associate with physicalism. And deciding what to add to that conjunction amounts to deciding what counts as 'physical', which was the problem we were trying to avoid in the first place[16].

[16] And there are many other attempts: for example, one suggestion is that the 'physical' is distinguished by 'spatiality', either in the sense of being spatially *located* (that is, standing in spatial relations to other things) or in the sense of being spatially *extended*. In an early paper on the relationship between the mental and the physical, Jaegwon Kim defines physical properties in this way:

The solution microphysicalism provides to Crane and Mellor's challenge is to appeal to the 'physicality' of the *very small*. Its implicit suggestion is that if we consider facts about sufficiently small entities, the qualifier 'physical' becomes redundant. For the microphysicalist, the micro is 'physical' by default, perhaps because micro-level entities are too small to have any sort of properties *other* than physical ones. Accounts of microphysicalism then insist on some sort of constraint between these micro-level features and all other features of the world. As Philip Pettit puts it in defending his own definition of microphysicalism, microphysicalism imposes a "dictatorship of the proletariat",

---

> Property *P* is a *physical property* if and only if that an object has *P* presupposes that it has extension. (Kim 1971: 335)

And more recently, Ned Markosian (2000) argues that the true mark of the physical is not so much spatial extension as it is spatial *location*: to be physical is to be located in space (and time, one presumes). But while such 'spatialist' approaches have the advantage of being relatively clear in their metaphysical presuppositions, they come dangerously close to making physicalism a trivial doctrine: spatiality is an effective criterion for distinguishing the physical only if we have reason to believe that 'non-physical' entities are *not* spatial (whether spatially located or spatially extended). And while this is *perhaps* a plausible position to take when considering the relationship between physical and mental entities, it is clearly not a plausible position if our concern is with a broader form of physicalism: no one doubts that chemical, biological, geological, etc. entities are situated in space and time. These therefore count as physical entities, according to the spatialist approach. That may strike many as the appropriate result, at least as far as *entities* are concerned; but the spatialist criterion leaves us unable to distinguish between physical and non-physical *facts* or *properties* of any sort, insofar as those facts relate to spatially situated entities. Physicalism could only be violated if there were facts about entities *not* situated in space and time that were also not determined by facts about entities that were situated in space and time. Such a form of physicalism may say something interesting about Cartesian souls and the like, and there is a tradition of distinguishing between the physical and the mental as 'spatial' and 'non-spatial' (see Feigl 1958, for instance). But it would say nothing about the relationship between physical properties and biological properties or as they are usually conceived. If our interest in physicalism extends beyond the relationship between the mental and the physical, the spatialist form of physicalism will seem quite trivial.

where it is the facts about the very small constituents of the world that ultimately determine how all things behave (Pettit 1993: 220-221).

However, while clearly this is one way to meet Crane and Mellor's challenge, it is not the only way: compositionality makes at least as strong a claim. And it isn't clear that microphysicalism has any *special* claim to capturing the *concept* of 'the physical'. Clearly it is reasonable to think that sub-atomic entities lack higher-level properties: we don't expect muons or quarks to have mental lives, or to obey principles of economics. But this appears just to be a contingent fact: nothing *requires* that micro-level entities be characterized in terms of *mass*, *charge*, and *spin* and not *temperature*, *affinity*, or *fitness* (or whatever). The fact that we don't expect muons to have minds or photons to possess psychological states is beside the point: microphysicalism would be consistent with the claim that they *did*, provided higher-level facts depended on those lower level facts in the right way.

Pettit admits as much when he glosses one of his principles of physicalism, the claim that "different kinds of things in the empirical world share subatomic levels of composition of the kind that...microphysics...posits" (Pettit 1993: 214) by noting that the microphysicalist "can be more or less sanguine about the accuracy of actual physics, or even about the propriety of its methods" and "can even admit that microphysics may be forced to countenance entities that by present intuitions are not of an intuitively 'physical' character" (Pettit 1993: 214). Thus microphysicalism doesn't even *attempt* to give an analysis of the concept of 'the physical'. What microphysicalists such as Pettit really defend might better be called *smallism*[17]. Smallism asserts the fundamentality of the very small, and it will be false provided that facts about 'big' things are not determined by the facts about the very small.

While this is not a trivial claim, it is unclear why there is anything particularly 'physical' about Pettit's definition. Instead, the constraint imposed by microphysicalism is one of *scale*: the fundamental facts about the world are facts about the very small constituents of the world, whatever they may be.

This rather absurd possibility doesn't mean that microphysicalism is meaningless: microphysicalism still defines a non-trivial doctrine that is plausibly true and so meets Crane and Mellor's challenge to provide a meaningful definition of physicalism. However, it suggests that microphysicalism doesn't have a necessary connection with our conception of the physical: we can conceive of 'smallism' being true in a world that violates our intuitions about physicalism in some other way. But equally, the more general claim of compositionality is non-trivial and plausibly true. So we can't argue for microphysicalism on the grounds that it is the only way of defining physicalism.

The Argument for Physicalism

One common reason for defending microphysicalism that I haven't addressed has to do with the argument *for* physicalism, in particular for versions of physicalism that are put in terms of supervenience, as we've just done. In their *reductive* account of physicalism, Oppenheim and Putnam argued for their view by

[17] The term is Robert Wilson's, and seems particular apt for Pettit's view.

presenting what they thought were compelling examples of successful micro-reductions from various sciences. Arguments for the weaker claim of *supervenience* are usually less direct: what we need to do is simply to show that a failure of supervenience would lead to some contradictory result, rather than to give any account of *how* one specific theory or branch of science or group of properties or whatever depends on another. It is typically assumed that a crucial premise in the standard argument *for* this type of physicalism – the so-called 'completeness' principle – is only satisfied at the micro-level[18]. If that were true, then only microphysicalism would be supported by an argument, and we might think that that shows it is the appropriate way to understanding physicalism after all.

The idea of completeness is that a given domain is 'causally complete' just in case every event in that domain that has any cause has a cause also in that domain. The argument for supervenience then runs as follows. First, suppose that causation is deterministic in the sense that two identical entities or groups of entities can never give rise to distinct effects. This assumption is not necessary, as the argument can be recreated provided that indeterministic causation is regular (i.e., identical entities give rise to effects with the same probability), but it simplifies the argument greatly. Suppose that two entities differed with respect to

[18] See Kim 2003, for example. Scott Sturgeon (1998), appears to take completeness not only to be particular to the micro-level, but more specifically a feature associated with the quantum mechanical description of the micro-level: he regards the so-called 'collapse postulate', which says that the probability of observing a system in a particular state is a function of the super-posed state of the system before observation. The connection between this postulate and the concept of causal completeness is interesting, but beyond the scope of this present work.

some property P and yet were identical in terms of their micro-level constituents. Furthermore, suppose that any property must be detectable in some way, so that there is some context in which an individuals' having or lacking a property makes a specific difference in the world. Perhaps not all properties are detectable in this sense, and if there *are* any undetectable properties, this argument will not show that *they* supervene on microphysical properties; but since they aren't detectable, it's difficult to be upset by this. Suppose that the specific difference a property makes consists in moving some sort of pointing device: in otherwise identical contexts, an entity's having property P causes the pointer to move, while an entity's lacking P does not. Whether or not the pointer moves makes a difference for its micro-level constituents: if the pointer moves, they move, too. Now, if the micro-level is 'complete', then the movement of those micro-level constituents of the pointer must have a micro-level cause. But by hypothesis, there is no micro-level difference between the two entities that differ with respect to P. Thus (given our assumption of determinism) those two entities can never give rise to distinct effects in the same context and we have a contradiction: if P is detectable, then the entities could not have been identical at level L after all.

Completeness plays a vital role in this argument. This is why many have assumed that we can only *argue* for microphysical supervenience: while it might not be *certain* that the micro-level is complete in this sense, it seems clear that any 'higher-level' domains will *not* be complete. It is not true, for instance, that all chemical events have purely chemical causes: a chemical event might be instigated by a release of electrical charge, for example.

However, if we pay attention to the structure of the argument, we can see that the assumption of completeness is actually stronger than necessary. For the argument to work, we need to be able to say that the presence or absence of some property – call it the 'higher-level' property – makes a difference in terms of some other sorts of properties, such as the movements of the micro-level constituents of a pointer – call these the 'lower-level' properties. Then all we need to establish is that *that difference* must have some cause at the lower 'level'. This is a more limited claim than true completeness: perhaps many other events at the lower-level only have causes from yet other levels. But as long as we are confident that the consequence of the characteristic effect of the higher-level property has a lower-level cause, the argument for supervenience goes through.

This seems to provide a much more plausible account of why we think that we *can* find explanations at lower (but not 'micro') levels than does the more demanding argument from causal closure. For example, why do we think that psychological properties and states must supervene on 'neural' properties and states? Surely it is because the characteristic effects of those psychological properties (desires, wishes, beliefs, etc.) are things such as bodily movements: we exhibit our beliefs and desires by acting on them (or by being able to act on them). Obviously, such actions are produced by the actions of muscles which are constituted by cells and fibers (whatever) and it is the thought that the behavior of *these* can be account for in terms of neural activity that leads us to conclude that the psychological must supervene upon the neural. Obviously not *all* cellular events are caused by other cellular events, but if enough relevant ones are, we have enough to argue for supervenience. So it's wrong to think that physicalism must be formulated as microphysicalism because there is no way to *argue* for physicalism otherwise.

Composition and Parthood

If physicalism is understood as compositionality, then understanding physicalism ought to involving understanding the nature of *parts*. 'The properties of wholes are determined by those of their parts', but what sorts are parts are there in the world?

This brings us to the second assumption about physicalism I want to challenge: the assumption that it is only *spatial* decompositions that are explanatorily relevant in the physical sciences. This sort of assumption is clearly implicit in Oppenheim and Putnam's account of reductive levels consisting of entities related by parthood: lower-level entities are not just parts, but specifically spatial parts of the higher-level entities they compose. The idea that physical parts are *only* spatial parts has even been suggested as a criterion of 'the physical' in response to the sort of challenge from Crane and Mellor discussed earlier. For example, Ariel Meirav (2000) offers a 'mereological' definition of what it is to be a physical object: he argues that while physical objects and non-physical objects might both exist in space, a distinction can be drawn between the ways that physical objects and non-physical objects relate to the regions of space they occupy. Physical objects 'mereologically correspond' to those regions in a way that non-physical objects do not. Given a primitive relation of 'occupation' holding between objects and regions of space, Meirav defines a relation of 'strong mereological correspondence' as:

> Object O bears strong mereological correspondence to region of space D $=_{\text{def}}$
>
> (1) for all u, for all v, such that u is a part of O and v is a part of u, there exist x and y, such that x is a part of D and y is a part of x, and u occupies x and v occupies y.
>
> (2) for all x, for all y, such that x is a part of D and y is a part of x, there exist u and v, such that u is a part of O and v is a part of u, and u occupies x and v occupies y.

So an object bears this relation to a region of space just in case any parts of that object occupy parts of that region of space, and any parts of the region are occupied by some part of the object. *Physical* objects are distinguished as those objects that bear this relation to some region of space. Meirav observes: "The structure of physical objects is thoroughly spatial; the parts of a physical object always bear more or less definite spatial relations to one another" (Meirav 2000: 627) This is in contrast with non-physical objects. For example, a non-physical objects such as an 'economy' might have a part such as 'the public sector', but there will not necessarily be a corresponding relationship of parthood between the region of space occupied by the economy and that occupied by the public sector. Physical objects are thus distinguished by spatial parthood.

However, while we might not doubt that physical things *have* spatial parts, we might wonder whether they could have other sorts of parts as well. For instance, in discussing the idea of reduction in chemistry, Krishna Vemulapalli and Henry Byerly (1999) argue that the explanatorily relevant parts of a system are often *not* spatial parts:

> A major of source of confusion about the significance of reductions, as we shall illustrate, is that the components involved may not be proper spatial parts of the system. Attempts to interpret new data commonly proceed from what has been understood in a simpler system. The simple system, however, is not necessarily a micro system. (Vemulapalli and Byerly 1999: 19)

The illustrations Vemulapalli and Byerly give are primarily drawn from thermodynamics. For example, they argue that while the equilibrium volume of a mixture of substances is a function of the equilibrium volumes of the components mixed, those components are not spatial components, since they are spatially as extensive as the mixture they compose. What's more, the relevant characteristics of those components – their 'partial' equilibrium volumes – depend upon the sorts of mixtures they are parts of. So the explanation of the volume of a mixture is not given in terms of characteristics of independent, spatially segregated components.

A rich source of examples involving non-spatial decomposition is found in the field of multi-scale analysis. Multi-scale analysis encompasses a variety of mathematical techniques for describing the behaviour of complex physical systems by analyzing their behaviour on different temporal or spatial scales[19]. A contemporary collection of papers addressing mathematical issues associated with

[19] The relevant scales may also include energy scales, though I don't discuss such cases in subsequent chapters.

multi-scale analysis summarizes the widespread applicability of these techniques as follows:

> Physical problems are often posed with many time or space scales. Reactors containing different chemical reactions occurring simultaneously may be modelled by equations with many time scales, each corresponding to a different reaction. In electro-optics, the electrical and optical components will exhibit rates that are vastly different; in neurobiology, different biochemical mechanisms and electrical effects result in a variety of time scales occurring the models. (Jones and Khibnik 2001: vi)[20]

The authors go on to emphasize the special utility in applying multi-scale analysis to complex problems, in particular the reduction in complexity it affords:

> The physical significance [of multi-scale analysis] is fortuitous as the multiplicity of scales allows us to reduce the order of the system. This reduction of systems can bring a seemingly complicated system to a surprisingly manageable form. The reduction of multiple-time scale systems offers more than other types of reduction. Reductions onto different sub-systems can be made, namely those living on fast and slow time-scales respectively, and each of these may be analyzed by virtue of their lower dimensional character. There is an added extra in this case,

[20] The classic mathematical text in this field is Nayfeh 1973, which includes an extensive bibliography of applications of multi-scale methods. Applications to fluid mechanics in particular are discussed in van Dyke 1975. While discussions of multi-scale techniques are spread throughout various areas of mathematics and other sciences, there is now a SIAM journal, *Multiscale Modeling and Simulation*, as well as numerous collections of papers devoted to both theoretical and applied problems associated with describing systems on multiple scales, such as Brackbill and Cohen 1985.

however, which is that intrinsically higher-dimensional behaviour may result from piecing the different reductions together. (Jones and Khibnik 2001: vi).

As I mentioned earlier, there is a variety of senses of 'reduction' in philosophy, and in some ways the sense of reduction Jones and Khibnik seem to have in mind here is not *so* different from the 'micro-reductions' of Oppenheim and Putnam (for example). The point of these multi-scale reductions is not, however, to analyze complex systems into *smaller* components, but instead to analyze them into *simpler* ones; those simpler components in these cases are component system 'living' on different time-scales. They are *parts* of the system, but not spatial parts; they are *non-spatial* parts[21].

Multi-scale techniques are extremely useful in modelling systems in a wide variety of disciplines, including fluid mechanics, aerodynamics, chemistry, and geophysics. I'll discuss examples of decomposition on distinct fast and slow time scales in chapter three, and decompositions on distinct *spatial* scales in chapter five, but for now I only want to point out that there *are* examples of decomposition in the physical sciences that are not of the spatial sort. As an informal illustration, consider the behaviour of a system modelled as a van der Pol

[21] It *is* common to distinguish between spatial and *temporal* parts, but that isn't the distinction I'm going to suggest here. I avoid using the more accurate terminology of 'spatio-temporal' parthood purely for stylistic reasons: as will become clear from the discussion, 'non-spatial' parts are not merely temporal parts. Perhaps the term 'overlapping parts' would best describe the relation between the parts I have in mind, since each is assumed to be as extensive as the system they compose. But the term 'overlap' already has a use in mereology: two things overlap when they have some third thing as a common part, and that clearly isn't what we have in this case. So instead I call these non-spatial parts.

'relaxation oscillator', described by the differential equation $x'' + \mu(x^2 - 1)x' + x = 0$. Figure 1.1 illustrates the typical behavior of such an oscillator, showing the amplitude of oscillations over time:

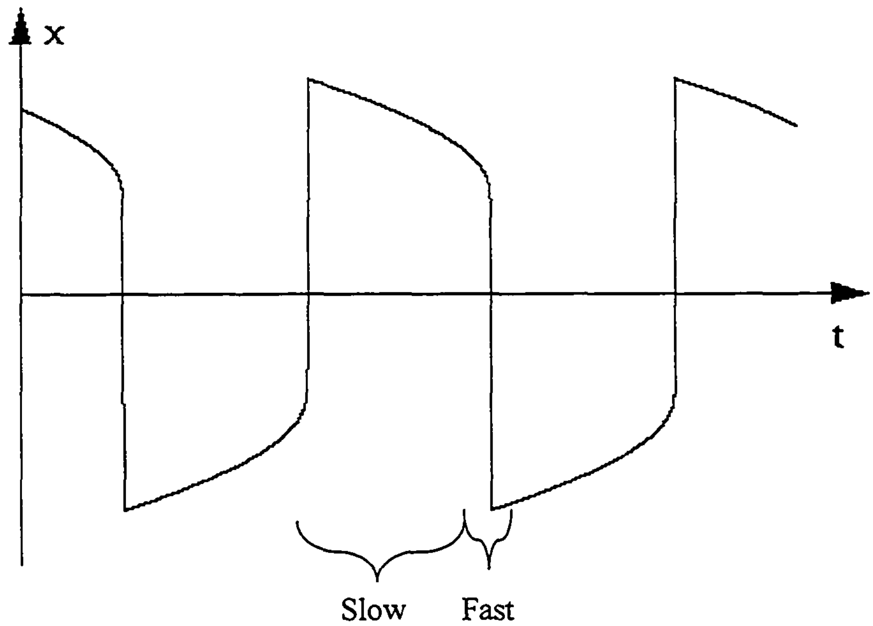

Figure 1.1 Representative behavior of a Van der Pol relaxation oscillator.

The characteristic behavior of such systems is the combination of a gradual build up in amplitude followed by a sudden 'relaxation' involving a sharp plunge or climb. Solving the equations describing such behavior is often impossible to do directly. On the multi-scale approach, the behavior of the system as a whole is attributed to the combined behavior of two or more component systems operating on different scales. It is important to note that this is not simply a matter of partitioning the behavior illustrated in figure 1.1 into 'slow' and 'fast' regions and then solving these independently: that would be an example of a *temporal* decomposition. On the multi-scale approach, the component systems are assumed to operate *simultaneously*, so that the behavior of the composite system is the result of the interaction of those components. In the case of the relaxation oscillator, the 'slow' scale system dominates in the wide regions characterized by the gradual change in amplitude, while the 'fast' scale system dominates in the narrow regions characterized by a sudden rise or plunge. However, both systems have *some* influence on the behavior of the composite system throughout the entire interval of its 'life'.

The suggestion I want to take seriously is that such decompositions reveal genuine parts of the system: multi-scale analysis works precisely because it focuses on genuine parts of a system's behavior and gives a suitable account of these. The behavior of the parts determines the behavior of the system as a whole. In these cases, the parts are not spatial parts, but they must still be included in any characterization of physicalism adequate to physics. So 'microphysicalism' is doubly wrong: it misses out on the importance of higher-level *spatial* decompositions, and it misses out on the importance of *non-spatial* decompositions as well.

Non-Spatial Parts and Reduction

To see the importance of recognizing non-spatial decompositions, let's return to the account of physicalism as compositionality defended by Oppenheim and Putnam. The most contentious aspect of their account is their claim about the *reducibility* of theories of one level to those of another[22]. The idea of reduction is widely regarded as suggesting a sort of second-class status for the theories reduced, and while it is readily acknowledged that there are often interesting explanatory relations between different theories, there is also a widespread resistance to claims about reductionism. In particular, within the philosophy of mind and psychology, there has been a long-standing resistance to the idea that psychological properties or theories can be 'reduced' in any useful sense to physical ones. Similar views are found in the life-sciences and social sciences.

There are a variety of philosophical accounts of what counts as a reduction, and so whether or not one finds the idea of reductionism in general plausible depends largely on what one counts as a 'reduction'. We've already seen the account of micro-reduction defended by Oppenheim and Putnam. The other classic account of inter-theoretic reduction is found in Nagel 1961. For Nagel, reduction is a matter of logical derivability between the laws comprising distinct theories: from the 'reducing' theory, together with various 'bridge laws' connecting the terms of the distinct theories, we are supposed to be able to *derive* the laws of the reduced theory. Again, while there can be undeniable utility in continuing to use a theory after it has been shown to reduce to another, *in principle* the reduced theory is no longer necessary.

[22] For example, Jerry Fodor's (1974) classic defense of a view of psychology as an 'autonomous' science, "Special Sciences: Or, the Disunity of Science as a Working Hypothesis", is clearly directed at claims of the sort Oppenheim and Putnam advance; for a critical appraisal of micro-reductionism contemporary with Oppenheim and Putnam's paper, see Schlesinger 1961.

It is often suggested that reductionism in these areas offends our dignity as humans by suggesting that what we take to be our characteristically human features are in fact 'simply' the products of our physical structure. This is the sort of view expressed by Marcelo Sabates, for example, when he writes of non-reductive physicalism that it appears to promise "loyalty to a broadly naturalist or physicalist metaphysics appropriate to the times, while keeping *what seems an ineliminable part of our dignity as human beings; the autonomy of our minds*" (Sabates 2001: 14). However, while this sort of objection may play an important role in attitudes toward reductionism, there are other concerns about reductionism that don't depend on the idea that it is somehow undignified. The question of reductionism can be raised in chemistry, biology, ecology, economics, and the like, even though it is seems that there would be nothing aesthetically unpleasant or undignified about such reductions. It hardly seems right to object to reductionism in chemistry, for example, on the grounds that suggestions of reduction offend the dignity of our conception of the chemical world. Instead of a concern about dignity, worries about reduction in these areas are better understood as worries about *eliminativism*. The worry is that showing that a particular set of theories reduce to another shows that the reduced theories are unnecessary, and that the entities and properties they describe are in some sense unreal. If theories about higher-level entities and properties are in fact reducible to lower-level theories, then in principle there is no need to posit the existence of those higher-level entities and properties. This is the sort of view rejected by

Peter Smith as he develops an account of 'explanatory interfacings' as part of a more moderate view of reduction:

> If radical physicalism [that is, what I've called 'eliminativism'] is the doctrine that in some sense there are only atoms in the void – that the only genuine entities, properties, and facts are the entities, properties, and facts recognized by fundamental physics – then it is precisely a *denial* of such a physicalism that gives rise to the pressure for explanatory interfacings...If there is a driving prejudice at work here, it is not radical physicalism but a principle, *P*, to the effect that the behavior of wholes is in generally causally produced by the behavior of parts, so that our explanatory stories about wholes must be consonant with our stories about the causal mechanisms constituted by their parts. (Smith 1992: 25)

Note that rather than any specific commitment to the fundamentality of the physical, Smith's 'driving prejudice', which he gives the rather unlovely title 'Principle P' is a claim of compositionality: facts about wholes are determined by facts about parts.

While we can't doubt the practical utility of talking about higher-level entities and properties, the worry about eliminativism is that such things are only needed because of our cognitive limitations: *in principle*, we could get by without them, if only our cognitive abilities were stronger. This view of reduction might seem to be implausibly extreme, and some have tried to defend reductionism by suggesting that the worry about eliminativism is simply misguided. For instance, Carl Hempel writes:

> The kinetic theory of gases plainly does not show that there are no such things as macroscopic bodies of different gases that changes volumes under changing pressure...and that there "really" are only swarms of randomly buzzing molecules. On the contrary, the theory takes for granted that there are those macroscopic events and uniformities. (Hempel 1966: 78)

However, the worry over eliminativism is not groundless, given the standard accounts of reduction. For instance, Oppenheim and Putnam's account makes it clear that "even if we cannot define in B_1 analogues for some of the theoretical terms of B_2 [that is, if 'Nagel-reduction' fails] we can *use* B_1 *in place of* B_2" (Oppenheim and Putnam 1958: 6, emphasis added). Finding a reduction from one theory (or 'branch' of science in this case) to another means that the former is no longer needed, and that the reducing theory can be taken to give the *true* picture of reality[23].

Rosen and Dorr (2002) defend just this sort of view of reality, though their motivation and argument is concerned more with the metaphysics of parthood than with any particular account of reduction. They defend a view of 'compositional nihilism', according to which composition is a 'fiction' and there are *only* non-composite entities and properties. For the compositional nihilist, all

[23] Perhaps the most extreme statement of the sort of view opposed by non-reductive physicalists is that *attributed* to Albert Einstein by the physicist Hans Dehmelt. Einstein is supposed to have once commented "You know, it would be sufficient to really understand the electron" (Dehmelt 1989: 8618). However, I have been unable to find any source other than Dehmelt for this, nor any other comments in Einstein's writings that suggest such a view.

talk of composite 'wholes' and their properties is merely a matter of convenience, and in principle could be reformulated as talk of (ultimate) parts[24]. For instance:

> When the chemist says that a water molecule is made of two atoms of hydrogen and one of oxygen, he does not take himself to be speaking figuratively. If you ask him whether his claim is meant to express the sober truth, he may well say, "Yes, of course; this is serious business." Nonetheless, apprised of the considerations we have rehearsed in this essay [that is, the argument for nihilism], he may be inclined to back off from his confident claim about composition. If he is canny he may say, "I'm not sure whether what I said is strictly true. But what I am sure of is this: what I said was *true on the assumption that composite things such as molecules exist.*" (Rosen and Dorr 2002: 169-70)

Note that Rosen and Dorr's argument is a purely 'metaphysical' one, not based on any account of reduction, but instead based on their view of parthood. But the conclusion they reach is clearly the sort of conclusion that 'non-reductionists' think follows from reductionism. The *objection* to reductionism is that this consequence of reductionism is absurd: *of course* there are 'wholes' and these have characteristic properties. Hence, reductionism must be false. Instead, the view non-reductive physicalists try to defend is something like that I associated with the Steven Weinberg in the introduction[25], where he insisted that

[24] In conversation, Mark McGivern – who is five – explained compositional nihilism to me as follows: "Sand is really just bits of rocks and shell. So 'sand' isn't a real word. Instead of saying 'sand' we should really just say…'bits of rock and shell'".

[25] Here I mean the introduction to the thesis, not the introduction to this chapter.

macro-level properties such as temperature and entropy cannot be treated merely in micro-level terms, but must be dealt with in their own terms instead (again, see Weinberg 1987: 64-65)

This view shows why Hempel's response to the worry about eliminativism is unsatisfying. The worry is not that physics gives us no reason to believe that 'such things as macroscopic bodies of different gases' really exist: clearly we have good reasons to think that such things exist. For instance, as Hempel says, some of our theories of physics refer to such things, and by one prominent principle of ontological commitment – the Quinean principle that we are committed to those entities posited by our best scientific theories – this gives us good reason to believe in the reality of such things. The worry about eliminativism, however is not that physics makes no mention of such macroscopic bodies, but that they play no *essential* role on our physical theories. To show that the macro is *really* real, we need to show that our best theories of the world *must* refer to them. We need to show – as Weinberg says – that omitting such entities and properties from our theories would mean missing out on some genuine feature of the world, that talk of macro entities and properties is not simply a convenient way of representing a reality that is too complex for *us* to deal with directly. Since non-spatial decomposition is distinct from the sort of spatial decomposition that leads worries about the 'elimination' of the macro, understanding non-spatial parts could give us some grip on why some macro-level properties must be dealt with on their own terms.

Conclusion

I've argued that physicalism is best understood as compositionality, and that to better understand physicalism, we need to investigate parthood better. In particular, I've suggested that physics itself appeals to sorts of parts that are not usually mentioned in philosophical discussions, 'non-spatial' parts. This isn't meant to challenge the existence or importance of parts in the ordinary spatial sense, nor the importance of specifically *micro*-level decompositions: my point is that these sorts of parts are not the whole story on physical explanations that should be of interest to philosophers.

Since one reason for being interested in non-spatial parts is the thought that they don't fit the usual mold of 'reductive' accounts of physicalism, I'll now examine one of the main arguments against 'non-reductive' physicalism – Jaegwon Kim's 'supervenience argument – before going on to discuss non-spatial parts in more detail.

II

The Supervenience Argument

Introduction

In this chapter, I'll examine one of the central problems for physicalism – the problem of whether or not physicalism can be 'non-reductive' in the sense that while all facts are determined by physical facts, they are not all necessarily identical with, nor reducible to physical facts. The central argument that non-reductive physicalism is *not* possible is the 'causal exclusion' argument, which claims to show that if a property is not a physical property, it cannot play a causal role in the world. Non-physical properties are thus 'epiphenomenal'. To the extent that epiphenomenalism is unacceptable, the argument goes, we are forced to accept reductionism. Since Jaegwon Kim is the most prominent defender of this type of argument, I will focus on his version here, which is also known as the 'supervenience argument'.

Kim's argument is directed at the supposed non-reducibility of *mental* or psychological properties. Kim claims that this argument shows that on a non-reductive account, mental or psychological properties cannot play a causal role in the world. One the central issues about the supervenience argument is whether or not it 'generalizes' to generate similar consequences for other sorts (i.e., non-mental) of properties. This worry is particularly acute if we think of physicalism in terms of compositionality: if the properties of wholes supervene on those of their parts, and if the causal efficacy of *subvenient* properties means that supervenient properties are epiphenomenal, then it seems that the supervenience argument leads to the same sort of eliminativist position the non-reductive physicalist was trying to avoid by rejecting reductionism. Higher-level properties might be 'real' in some diminished sense, but lacking a causal role seems to deprive them of any true importance in our understanding of the world.

The Supervenience Argument

Kim's 'supervenience argument' runs as follows[26]. Note that since Kim develops his argument explicitly for the case of mental-physical supervenience, I will do likewise, though my subsequent interest will be in its significance in purely physical contexts. The applicability of the argument to supervenient properties more generally is the question of 'generalization', which I will deal with in the next section.

Kim's argument is constructed as a dilemma based on the acceptance or rejection of mental/physical supervenience, which he defines as:

> if something instantiates any mental property M at [time] t there is a physical base property P such that the thing has P at t and necessarily anything with P at a time has M at that time. (Kim 1998: 39)

[26] Kim has developed his argument in a number of places, and I take the version considered here (from Kim 1998: 38-47) to be more or less 'canonical form' of the argument, though a somewhat revised version appears in Kim 2003. A virtually identical account of the argument appears in Kim 1997.

If one *rejects* such supervenience, then Kim thinks one cannot hope to give a physicalist account of mental causation. Kim's argument for this horn of the dilemma is brief and I will simply accept the conclusion, since I've already given an argument for supervenience (the 'M-argument') in chapter one. Kim's main argument aims to show that accepting the supervenience of the mental is no more promising.

Kim considers a putative case of mental causation between two instances of mental properties M and M*. We would like to understand how an instance of property M can be causally related to an instance of property M*. Since we accept the supervenience of the mental on the physical, we must suppose that there are also instances of physical properties P and P* that necessitate the presence of M and M* respectively. We can represent this diagrammatically as:

$$\begin{tikzcd} M & & & M^* \\ P \arrow[u, "\text{supervenience}"] & & & P^* \arrow[u, "\text{supervenience}"'] \end{tikzcd}$$

Figure 2.1 Mental-physical supervenience.

where the arrows represent the fact that P and P* are 'subvenient' properties that necessitate M and M*. The problem is now to find a plausible causal role for M to play in the production of M*, given that physical properties P and P* are already present. This is problematic since the supervenience base property P* is itself sufficient (by hypothesis) for the presence of M*. There seems to be no work left for M to do in the production of M*.

Perhaps, however, M causes M* by causing P* (the 'base' of M*): that is, perhaps the putative case of 'mental-to-mental' causation is really a case of 'mental-to-physical' causation. We can grant this possibility, but we must keep in mind that M itself has a supervenience base property P (a physical property). And since by the definition of supervenience P will be sufficient for M, any instance of P will necessitate an instance of M. Kim concludes that P will 'pre-empt' M as a cause of P*, since on the standard accounts of causation, P will count as a cause of P* if M does. For instance:

> if you take causation as grounded in nomological sufficiency, P qualifies as a cause of P*, for, since P is sufficient for M and M is sufficient for P*, P is sufficient for P*. If you choose to understand causation in terms of counterfactuals, again there is good reason to think that P qualifies: if P hadn't occurred M would not have occurred (we may assume, without prejudice, that no alternative physical base of M would have been available on this occasion), and given that if M had not occurred P* would not have occurred, we may reasonably conclude that if P had not occurred, P* would not have either. (Kim 1998: 43)

So the core of the argument is that since supervenient properties are necessitated by their 'bases', they will compete with those bases for causal efficacy. If the supervenient property M was itself causally sufficient to produce P*, then P* would be *over-determined*, having more than one sufficient cause. As a matter of principle, wide-spread over-determination is to be ruled out, and since we are considering a generic case of 'mental causation', deciding that M and P were both causes of P* would violate that principle. The only sensible conclusion, says Kim, is "P caused P*, and M supervenes on P and M* supervenes on P*" (1998: 45). Thus the *true* causal activity lies between P and P* (the physical property instances): the causal activity between M and M* is only apparent. The mental property M (and thus mental properties in general) cannot be causally efficacious. Our revised diagram is now:

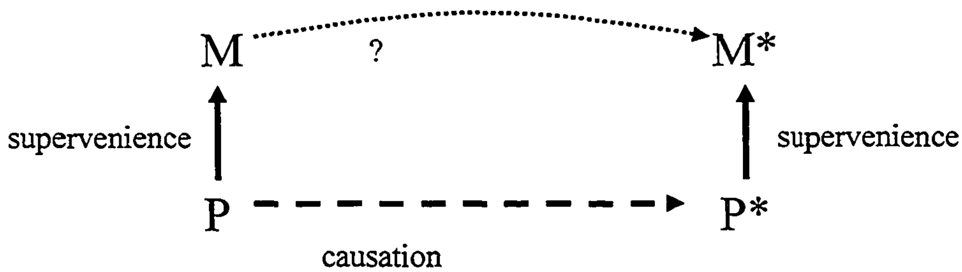

Figure 2.2 Causal competition.

where the arrow connecting M and M* indicates the causal relation we'd like to think exists between mental states, but which the supervenience argument calls into question. Note that Kim considers – and rejects – construing the relationship between P, M, and P* as a 'causal chain', and thus that P and M are not competitors in the causation of P* but simply both 'causal predecessors': "In general", says Kim, "the relation between base properties and supervenient properties is not happily construed as causal." (1998: 44). He gives two reasons:

(i) causes are standardly thought to precede their effects, whereas the instantiations of P and M are "wholly simultaneous"

(ii) it is difficult to imagine a causal chain from subvenient to supervenient properties, partly because it is difficult to see what sort of 'intermediary events' could connect the two

While it does seem odd to think of the relationship between P, M, and P* as a causal chain, these don't seem to be very compelling reasons for not doing so. The obvious response to (i) is 'so much the worse for the standard thoughts about causes preceding their effects'. For (ii), it isn't clear why there would need to be any 'intermediary events' between base and supervenient properties – particularly since these are assumed to be instantiated simultaneously. Kim's explanation of the second complaint includes the claim that such chains would violate the 'causal closure of the physical domain'. This is the principle mentioned in chapter one, also known as the claim of 'the completeness of physics'. This principle is supposed to capture the idea that physical events can always be accounted for in terms of purely physical causes. However, unless we define causal closure to *exclude* mental causation then viewing mental causes as links in a causal chain connecting one physical event to another won't obviously violate causal closure. Admittedly, Kim does seem to define causal closure in this way. He defines the principle as the claim that "if you pick out any physical event and trace out its causal ancestry or posterity, that will never take you outside the physical domain. That is, no causal chain will ever cross the boundary between the physical and the nonphysical." (1998: 40) A more appropriate definition of causal closure would ensure the existence of a physical cause for any physical event without explicitly ruling out the possibility of a non-physical cause. Rather than insisting that all physical events have *only* physical causes as a matter of principle, we can insist that physical events *have* physical causes (or have no cause at all). David Papineau explains the principle in this way:

> I take it that physics, unlike the other special sciences, is *complete*, in the sense that all physical events are determined, or have their chances determined, by prior *physical* events according to *physical* laws. In other words we need never look beyond the realm of the physical in order to identify a set of antecedents which fixes the chances of every physical occurrence. A purely physical specification, plus physical laws, will always suffice to tell us what is physically going to happen, insofar as that can be foretold at all. (Papineau 1993: 16)

Whatever we think of Kim's reasons for rejecting this 'causal chain' response to the worry about over-determination, the assumption of completeness gives rise to a second way of arguing for over-determination that *cannot* be resolved by appealing to causal chains. Rather than arguing that P should count as a cause of P* because of the supervenience of M upon P (and the assumption that M causes P*), we can directly assume that P* has a prior physical cause because of the completeness of physics. This now appears to be Kim's preferred formulation of the argument (see Kim 2003)[27]. By completeness, if P* has any cause at all, it must have a physical cause. The completeness principle itself doesn't guarantee that this cause will be the supervenience base of M, but without

[27] Note that there is an important difference between these two formulations of the supervenience argument. The original formulation remains effective even if the physical domain is *not* complete, provided we have reason to think that the mental (or whatever) does supervene on the physical. For instance, our evidence for the dependence of one type of property on another might come from an empirical generalization, and we might have no good reason to think that the domain characterized by the 'base' properties *is* complete. In such cases, supervenient causation would remain problematic given the original argument, while the completeness formulation of the argument simply would not apply.

loss of generality for the current argument we can assume that P is this physical cause. By assumption, though M supervenes on P, the two are distinct properties, and hence if each is a sufficient cause of P*, P* will be causally over-determined. Again, since wide-spread over-determination is thought to be unacceptable, we must reject the causal sufficiency of one of either M or P. As physicalists, the causal significance of P in producing P* is not to be questioned. Hence it is the 'mental' cause M that must lose out. What does this mean for 'non-reductive' physicalism? Kim:

> it is clear that the tacit assumption that gets the supervenience argument going is mind-body antireductionism; if the mental properties are viewed as reducible to physical properties in an appropriate way, we should expect to be able to disarm the argument (although of course the details will need to be worked out). (Kim 1998: 46)

What Kim has in mind here, I take it, is that if we accept the 'reduction' of mental properties to physical properties, then the causal competition argument will not raise problems, since we will in effect find a way of 'identifying' properties M and P: M causes P* (and thus M*) because P causes P* and M *is* P.

This is the supervenience argument: if mental properties supervene on physical properties, then mental properties cannot be causally efficacious unless they are in fact identical with physical properties. If we insist that mental states cannot be identified with or 'reduced' to physical ones – that is, if we stand by non-reductive physicalism – we must reject the assumption that mental states are causally efficacious in the first place. Non-reductive physicalism leads to epiphenomenalism about the mental, and to the extent that epiphenomenalism about the mental is unacceptable, so, too, is non-reductive physicalism. And to the extent that *reductionism* is unacceptable, we need to find a way of countering the supervenience argument.

Does the Supervenience Argument Generalize?

The supervenience argument is usually put in terms of 'mental' and 'physical' properties, but its assumptions are few and its conclusion should be correspondingly general. If properties of one type supervene on properties of another type, then, if the reasoning in the supervenience argument is correct, there should be a danger of causal competition between those properties. If the supervenience argument raises a real problem for the causal efficacy of the mental, then there seems to be a real danger that that same problem will face all supervenient properties, mental and otherwise. There are plenty of other types of properties (biological, geological, chemical) that we might also think supervene upon the physical. This means that the supervenience argument should apply equally to these properties, depriving them of any true causal efficacy. The question of whether or not the supervenience argument generates a similar difficulty for *all* supervenient properties is called the problem of 'generalization'.

The real worry about generalization is the idea that it leads to a 'draining down' of causal efficacy, leaving us with a view much like eliminativism. If we take non-physical and physical properties to apply to relatively 'macro' and 'micro' levels of reality – as most philosophers do – then the supervenience argument seems to show that all causation *really* takes place only at the micro-level: only micro-level entities are endowed with 'causal powers' in virtue of the properties they instantiate[28]. There are two ways of viewing this conclusion. We can either regard generalization as a corollary of the main argument, extending the implications of the supervenience argument to a wide variety of putatively non-physical (yet supervenient) properties, or we can regard generalization as a *reductio* of the supervenience argument, on the grounds that to question the causal standing of chemical, biological, geological properties and the like is absurd. A number of philosophers appear to take the latter view. For instance, Robert van Gulick writes:

> ...reserving causal status for strictly physical properties...would make not only intentional [i.e., psychological] properties epiphenomenal, it would also make the properties of chemistry, biology, neurophysiology and every theory outside microphysics epiphenomenal...If the only sense in which intentional properties are epiphenomenal is a sense in which chemical and geological properties are also epiphenomenal, need we have any real concern about their status: they seem to be in the best of company and no one seems worried about the causal status of chemical properties.[29] (van Gulick 1992: 325)

---

[28] See, for instance, van Gulick 1992, Baker 1993, and Burge 1993

[29] Actually, there are two senses in which generalization can be seen as a *reductio*. The most common sense is that defended recently by Block (i.e., Block 1997, 2003) according to which generalization shows that there must be some flaw in the supervenience argument. But van Gulick's comment in particular might be read as a *reductio* of a different sort: we accept the conclusion of epiphenomenalism, but conclude that – due to the company it keeps –

In response, Kim notes that "This is a little like being told that we shouldn't worry about, say, being depressed because everyone else has the same problem" (Kim 1998: 78). While chemists themselves might not worry about the efficacy chemical properties, this doesn't by itself mean that there is nothing there to worry about. Kim points out:

> Perhaps no one is worried about the causal efficacy of chemical properties or biological properties, but then not many people are really worried about mental causation either. What some of us are worried about is finding an intelligible *account* of mental causation. (Kim 1998: 78-9)[30]

Here Kim seems to be viewing the worry about causal powers 'draining down' as a corollary of the supervenience argument, rather than a sufficiently absurd consequence to form a *reductio*. If we are unconvinced that the draining down of all causal powers is absurd enough, Block 2003 offers the following more extreme version of the *reductio*. If the supervenience argument is correct,

---

epiphenomenalism about the mental is not as bad as we might have thought. On this interpretation, it is the unacceptability of epiphenomenalism that is the subject of the *reductio*.

[30] Note that Kim's comments on how we should regard generalization are ambiguous. For example, Kim notes that though we surely think geological properties supervene upon physical ones:

> ...no one seems to worry about geological causation, and there evidently seems no reason to start worrying. If so, shouldn't we conclude that there must be something wrong with the argument of the preceding section [that is, the supervenience argument]? (Kim 1998, p. 46)

And in responding to criticism from Paul Noordhof, Kim writes:

> Noordhof argues that the supervenience argument, if it works against supervenient properties, works equally well against micro-based properties, and, further, that mental properties can perfectly well be construed as micro-based properties. If these two points are correct, that would put me in a very tight spot: Noordhof's first claim would show that my supervenience argument for mental epiphenomenalism generalizes beyond supervenience properties after all, *indicating that the argument must be flawed...*" (Kim 1999: 115, emphasis added)

and if it generalizes, then causation occurs only at the lowest level of reality. But Block considers it a real possibility that there is no *lowest* level of reality: that matter is endless decomposable into smaller and smaller different constituents, and that the properties of each supervene on those at the some lower level[31]. If this were the case, then causation might 'drain away' entirely, meaning that *all* properties are epiphenomenal. Block argues that while the question of whether or not there is a bottom level is an open one (even if we regard the answer as very likely), the question of whether not there is causation *at all* is not.

> It is an open question whether there is or is not a bottom level, but it is not an open question whether there is any causation. It may be an open question whether cigarette smoking causes cancer but it is not an open question whether anything ever causes anything. So something is wrong with Kim's Causal Exclusion Argument. (Block 2003: 139)

Block draws the following principle as a consequence of Kim's argument:

> The Anti-Reductionist Conditional: If there is no bottom level, then cancer never causes suffering or death. (Block 2003: 139)

But it would be absurd to suppose that questions such as whether or not smoking causes lung cancer depend upon whether or not there is a bottom level to reality. For instance, it would be absurd to argue that cigarette manufacturers were not liable for the suffering of their customers on the grounds that it is open question whether or not there is a bottom level to reality. Conversely, it is absurd

[31] Block's intent here is not to argue that this view of matter is likely to be true, but instead only that its truth is a genuinely open question; here he cites the work of Nobel laureate Hans Dehmelt (1989), who apparently does regard this as a genuine possibility.

to suppose that our conviction that smoking *does* cause cancer (or any other convincing case of causation) can lead us to affirm that there *must* be a bottom level of reality. So the Anti-Reductionist Conditional, and hence Kim's argument, must be rejected: Kim's reasoning about supervenience and causation must in general be wrong. And if it's wrong in general, then there is no reason to think it is correct about psychology in particular. Thus non-reductive physicalism is compatible with the causal efficacy of the mental, and the autonomy of the special sciences is saved[32].

The fault, says Block, must lie in the assumption that causes at different levels compete with each other in the first place: that genuinely causal facts about atoms and molecules cannot co-exist with genuinely causal facts about cells and organs or causal psychological facts about complex organisms. He calls this assumption the *Exclusion Principle* and argues that what the supervenience argument really shows is that the Exclusion Principle should be rejected:

---

[32] One way of responding to Block's argument would be to argue that the fact that there is no bottom level doesn't undermine the existence of causation entirely, because of some intricacies in the nature of 'infinitary reasoning' (as Block calls it). The idea would be that although causation at each level is in fact pre-empted by causation at the next lowest level (and so on endlessly), this infinite chain of causal pre-emption is somehow compatible with the existence of real causation, perhaps in something like the way real movement is compatible with the sort of infinite chains of partial movement found in Zeno's paradoxes. In response to this, Block offers a 'Conservative' Anti-Reductionist Conditional: "If there is no bottom level, and if the issue in philosophical logic concerning infinitary reasoning turns out as I [Block] suggested, then cancer never causes death." (Block 2003: 139). Now the claim about the existence of causation takes into account the possibility that infinitary logic might show that causation could exist without there being a 'bottom level'. Yet it still seems absurd to suggest that a person's liability for a role in any putatively causal sequence (i.e., manufacturing cigarettes that 'cause' cancer) should depend on results from particle physics *and* philosophical logic.

> If there is no bottom level and if there is endless supervenience, then Kim's Causal Exclusion Argument would yield absurd results. The Exclusion Principle (that causally sufficient properties at one level exclude causally sufficient properties at another level) is to blame and should be abandoned. (Block 2003: 140)

In response, Kim argues that the Exclusion Principle can be maintained because the supervenience argument does not 'generalize' in the way suggested by Block. Partly this is because Block's line of reasoning only pays attention to one of the two premises in the supervenience argument: the supposition that higher-level properties are causally efficacious. The other crucial premise is that higher-level properties are not reducible to lower-level ones: and Kim insists that it is *this* premise he is interested in attacking[33].

The challenge here for Kim is to give an account of the relationship between properties that clearly does not lead to eliminativism, either by the original argument that reduction leads to eliminativism, or by Block's extension of the supervenience argument that the 'causal exclusion principle' also leads to eliminativism. Central to Kim's response is a distinction he draws between the *order* of a property and the *level* at which it applies. The *level* at which a property applies is some indication of the size and sort of entity instantiating that property in terms of the usual division of entities into levels of organisms, cells, molecules, atoms, and so on; Kim calls this the 'macro-macro hierarchy'. The *order* of a

[33] Of course, there is a third premise in the argument: the completeness of physics. However, neither Block nor Kim considers challenging this claim, and in general its detractors are few; however, see Hendry 1999 for a dissenting view.

property, on the other hand, is supposed to be a relative indication of whether that property is *realized* by another property or not. Higher-*order* properties are properties that are instantiated in virtue of some other (lower-order) property being instantiated by the same individual: higher-order properties are *realized* by lower-order ones. Kim's example of this is the relationship between the 'second-order' property *dormitivity* (the property of inducing sleep) and the various physico-chemical realizers of this property found in various drugs and substances that induce sleep. For instance, *dormitivity* is realized both by the properties characterizing *secobarbital* and by those characterizing *diazepam* (among others). In this case, properties of different *order* are instantiated by the same individual (a pill, say), and so any causal competition between properties and their realizers occurs at the same 'level' in the macro-micro hierarchy. The causal exclusion principle applies to these sorts of properties: no individual can have two distinct properties each causally sufficient (on a particular occasion) for a particular effect. The causal competition in the supervenience argument is thus *intra*-level causal competition: *dormitivity* is *not* causally efficacious, since a pill's being dormitive depends on its being secobarbital (for example). It is the pill's being secobarbital that *really* causes someone who ingests it to fall asleep (along with various contributing contextual factors). But, Kim emphasizes, these properties apply at the same level: it is the *pill* (say) that is both *dormitive* and *secobarbital*. The case of mental properties is similar: if my thinking of Vienna supervenes on my being in brain state V, then it is the brain state that is truly causal (in causing me to reply 'Vienna' when asked what I'm thinking of, for instance), but both states are properties of *me*, a 'mid-level' entity in the macro-micro hierarchy.

Block's reductio of the supervenient argument, on the other hand, depends on the assumption of *inter*-level competition. He assumes that it is the causal powers of lower-*level* entities that compete with those of higher-level entities, thus raising the threat that causal powers could 'drain down' to the micro-level or drain away entirely if there is no lowest level. But, says Kim, this is simply a misunderstanding of the supervenience argument: the argument concerns competition between properties *at* a given level and is consistent with causal efficacy at different levels.

> In general, supervenient properties and their base properties are instantiated by the same objects and hence are on the same level. This again is a simple consequence of the concept of supervenience: Socrates' goodness supervenes on his honesty, generosity, courage and wisdom, and it is the same person, Socrates, who instantiates both these subvenient virtues and the supervenient goodness...This means that the supervenience argument, which exploits the supervenience relation, does not have the effect of emptying macro-levels of causal powers and rendering familiar macro-objects and their properties causally impotent. (Kim 1998: 86)

What creates the illusion of inter-level competition, says Kim, is the fact that many lower-order properties are 'structural' or 'micro-based' properties pertaining to the parts of a higher-level individual. For instance,

> Having a mass of ten kilograms is a property of certain aggregates of molecules, like my coffee table. And it is a micro-based property of the table in the following sense: for my table to have this property is for it to consist of two parts, its top and its pedestal, such that the first has a mass of six kilograms and the second a mass of four kilograms. (Kim 1998: 84)

Kim defines 'structural' or 'micro-based' (Kim doesn't distinguish between the two) properties as follows:

> P is a *micro-based property* just in case *P* is the property of being completely decomposable into nonoverlapping proper parts, $a_1, a_2, \ldots, a_n$, such that $P_1(a_1), P_2(a_2), \ldots, P_n(a_n)$, *and* $R(a_1, \ldots, a_n)$[34]. (Kim 1998: 84)

So a micro-based property is a property something has just in case it has parts which themselves have certain properties and stand in a certain relation to one another. And while a micro-based property is intimately connected with the lower-level parts of an individual, it is still a property of that higher-level individual. So while micro-based properties might causally compete with other properties, this competition will be 'intra-level' rather than 'inter-level', as Block assumes.

However, we might still worry that causal powers might drain away *at* a given level. For instance, given the supervenience argument, my 'chemical-order' properties pre-empt my 'biological-order' properties and are in turn pre-empted by my 'physical-order' properties. Admittedly, these are all properties of *me*, but in the end it is only my 'physical-order' properties that are efficacious.

[34] Here Kim cites David Armstrong's account of structural properties with approval; see Armstrong 1978, 1997.

And this could still rob the chemical, biological, etc., properties at a given level of their apparent causal powers. We might see this result as just as absurd as the possibility of causal efficacy draining down to the micro-level. Worse, it isn't clear that this solution to the worry about inter-level competition avoids Block's worry about causal efficacy draining away entirely: we could still argue that if there is no end to the 'sub-orders' of properties, then there is still after all *no* causal efficacy, which is absurd. Rather than a worry about *inter*-level competition, this is a worry about *intra*-level competition.

Kim's response to this new worry is to claim that chemical, biological etc. properties can in many cases be *identified* with lower-order physical properties: "with properties like geological and biological properties, we are much more willing, intuitively, to accept a reductionist picture in relation to basic physical properties" (Kim 1998: 46). If a 'chemical order' property simply *is* a physical property, then there is no danger of the one pre-empting the other.

To argue for these identity claims, Kim examines the question of what it means for something to count as a 'physical' entity or property. Clearly, says Kim, "the basic particles and their properties and relations" are part of the physical domain, but so, too are "aggregates of basic particles, aggregates of these aggregates, and so, without end; atoms, molecules, cells, tables, planets, computers, biological organisms, and all the rest must be, without question, part of the physical domain" (Kim 1998: 113). Physical entities are thus the entities of 'basic' (i.e., particle) physics, together with 'aggregates' of those entities. For *properties*, Kim argues that they too are built from properties of fundamental physics. For instance, though the property of having a mass of one kilogram is not a property of any entity from particle physics, "it is a micro-based property whose constituents are physical properties and relations":

> We can think of this property as the property of being made up of proper parts, a_i, each with a mass of m_i, where the *m*s sum to one kilogram. And it seems appropriate to assume that the physical domain is closed under formation of micro-based properties: if P is a micro-based property of having parts $a_1, \ldots, a_n$ such that $P_1(a_1), \ldots, P_n(a_n)$, and $R(a_1, \ldots, a_n)$, then P is a physical property provided that $P_1, \ldots, P_n$ and R are physical properties (and relations), and each a_I is a basic particle or an aggregate of basic particles. (Kim 1998: 114)

To make Kim's suggestion here clear, consider the property of mass along with two simplifying assumptions. First, assume that 'atoms' are truly atomic, so that a carbon atom counts directly as a physical entity and its properties count directly as physical properties. Second, assume that Kim is right in treating mass as 'additive' so that the mass of a collection of bodies is simply the sum of the masses of those individual bodies[35]. Among the basic physical entities will be carbon atoms, and these have a mass of about 1.99×10^{-23} grams. So *mass of* 1.99×10^{-23} *grams* counts as a physical property. Intuitively, other masses, such as *mass of 12.01 grams* should also count as physical properties, but no fundamental physical particle (as far as I am aware) has a mass of 12.01 grams, so these can't

[35] This assumption is not exactly correct, due to what is known as the 'mass-defect' resulting from the equivalence of mass and energy, but I assume that that defect alone doesn't seriously affect Kim's argument.

be *basic* physical properties. Instead, *mass of 12.01 grams* must be identified with some micro-based property, based on the properties that fundamental particles *do* have. As it happens, 12.01 grams is the molar mass of carbon, so the appropriate micro-based property will be something like the property *being made up of proper parts* $a_1,\ldots a_{6.02 * 10^{23}}$ *such that the mass of each* a_i *is* $1.99 * 10^{-23}$ *grams*. Since this micro-based property is based on properties of fundamental physical particles, it too counts as a physical property.

This is a somewhat surprising account of what it means to be a 'physical' property, partly because Kim seems to intend it to apply equally to 'determinates' of properties (such as *mass of ten kilograms*) as well as to 'determinables' (such as simply *mass*). One consequence of this is that uninstantiated determinates of a given property cannot count as 'physical'. Suppose that the micro-world were simpler than it is, and that all fundamental particles had the same mass. Assuming again – as Kim does – that mass is additive, only positive integer multiples of that 'basic' mass would count as physical properties, since there could be no micro-based property where the component masses summed to any fractional value. If the basic mass unit were one 'kimogram', a property such as *having a mass of two kimograms* (the property of having proper parts x and y such that the mass of x is one kimogram and the mass of y is one kimogram) would count as 'physical', but the property of *having a mass of one and half kimograms* would not since the appropriate structure would not be possible. The same sort of situation could equally well arise in the actual world, though the details are obviously more complicated.

With this account of what it means to be a physical entity or property in hand, Kim argues that the worry about intra-level competition between higher and lower order properties can be met: chemical, biological, geological properties and the like can in general be identified with such micro-based properties, or with second-order 'functional' properties quantifying over first-order micro-based properties. For instance, "being a water molecule is a physical property, and being composed of water molecules (that is, being water) is also a physical property" (Kim 1998: 114). Once the higher-order property is *identified* with a particular micro-based property, the worry about causal competition disappears, since there is no competition between a property and itself.

This account of micro-based properties, says Kim, is adequate for a variety of physico-chemical properties:

> If transparency is taken as the property of passing light beams through without altering them, it counts as a second-order functional property. If transparency is identified with some microstructure, it will qualify as a micro-based property. The same can be said of such properties as water-solubility, ductility, thermal conductivity, inflammability, and the like." (Kim 1998: 115)[36].

[36] Note that Kim's talk of *functional* properties as physical properties is somewhat misleading: on his account, being a physical property is a matter either of being a fundamental physical property – i.e., a property of elementary particles – or of being a property derived from physical properties by the two 'closure conditions' on physical properties. The first of these says that structural properties based on fundamental – i.e., 'micro' – physical properties also count as 'physical', while the second says that 'second-order' properties which range over physical properties count as 'physical', too. But since it is *micro*-properties that are the fundamental physical properties in Kim's account, any second-order property can only range over micro-properties, or structural

Biological properties are to be accommodated in a similar way: "Being a cell may be a micro-based property; being a heart may be a second-order functional property (i.e., being a heart is plausibly viewed as being an organ/device with powers to pump blood)" (Kim 1998: 115). Thus, contrary to Block's argument, most properties in the sciences do not face any difficulties from the supervenience argument: the argument does not 'generalize' and the hence the Causal Exclusion principle does not lead to any absurd or unpalatable *general* conclusions about causal powers 'draining away'.

Let's summarize Kim's responses to the two worries about causal competition. On Kim's account, the Causal Exclusion principle does not lead to *inter*-level competition, because supervenience is an *intra*-level relation, and hence the only possible competitors a property could face would be properties from the same level. Thus the supervenience argument does not imply that all causation takes place only at the lowest level. And while the Causal Exclusion

---

properties based on micro-properties, or perhaps previously defined second-order properties already counted as 'physical'. Talk of transparency as the functional property of "passing light beams through without altering them" *sounds* like a property that is defined purely in terms of a *macro*-physical functional specification. But the only account of transparency truly available to Kim, given his definition of what it means to count as 'physical', would be to say that transparency is the property of having one of a number of specific micro-structures each based on the whatever properties characterize fundamental particles. The problem for this is that we might plausibly think that transparency should be genuinely *indifferent* to micro-structure, rather than simply being the property of having one of a number of different possible micro-structures. Defining transparency just as 'the property of passing light beams through without altering them' does seem like the appropriate functional definition, but it isn't really the definition Kim can give without abandoning his commitment to the idea that it is the connection to micro properties that makes any non-micro property 'physical'. I'll return to this characterization of macro-physical properties as *indifferent* to the micro-level in chapter three.

principle *does* lead to *intra*-level competition, this is not usually troubling because we are typically quite willing to accept reductionism in the case of non-mental properties. And contrary to the usual worry about reduction and eliminativism, it is clear that the identities involved in these reductions are between properties of *higher*-level entities. We have to accept the reduction of all causal properties to 'micro-based' properties, but *these* can still be properties of higher-level things.

Inter-level Competition Revisited

I'll examine the acceptability of Kim's claim about the reducibility of all (physical) properties to micro-based properties in a moment. First, however, there is another aspect of the supervenience argument that needs to be discussed. One consequence of Kim's appeal to micro-based properties in solving the worry about *intra*-level competition is that it revitalizes the original worry about *inter*-level competition. That was the worry that the supervenience argument shows that the properties of the fundamental particles composing me, for instance, exclude any of my 'higher-level' properties from playing any causal role. For Kim, it seems obvious that this is not the case. Considering his table as an example a higher-level object, Kim writes

> This table has a mass of ten kilograms, and this property, that of having a mass of ten kilograms, represents a well-defined set of causal powers. But no micro-constituent of this table, none of its proper parts, has this property or the causal powers it represents. H_2O molecules have causal powers that no oxygen or hydrogen atoms have...Clearly then

macroproperties can, and in general do, have their own causal powers, powers that go beyond the causal powers of their micro-constituents. (Kim 1998: 85, emphasis in original)

Kim is obviously right in these claims, but we might think he is *too* obviously right. Of course a ten kilogram table has causal powers that none of its micro-constituents have. But surely this isn't in doubt: surely, that is, no one suspects that the supervenience argument at *any* level entails that the causal powers we normally associate with some entity are pre-empted by some *individual* part of that entity, and certainly not a 'micro' part. Call this suggested view 'micro-czarism' (in deference to Philip Pettit's characterization of microphysicalism as 'the dictatorship of the proletariat'; see Pettit 1993: 220). Surely micro-czarism is (in general) false[37].

What *is* at issue is how macro properties relate to the collaborative efforts of their many micro-constituents. How do the causal powers associated with 'having a mass of ten kilograms' relate to the causal powers associated with the properties of the millions of millions of molecules constituting Kim's table? Kim's focus on micro-based properties makes this worry clear: even if we accept the identity between *mass of ten kilograms* and a particular micro-based property (thus avoiding *intra*-level competition), there remains a worry about inter-level competition between that micro-based property and the properties it is based on.

[37] I suppose there *could* be cases where a property of a macro entity and its associated causal powers could be attributable to the characteristics of one of its micro-constituents, but I think it is clear that this will be the exception rather than the rule.

Again we can ask whether these 'base' properties pre-empt the causal powers of the supervenient (micro-based) properties. This is a worry about 'inter-level' competition, and hence our solution to the worry about *intra*-level competition reinvigorates the worry about *inter*-level competition[38].

Kim considers this sort of objection in his response to comments made on the supervenience argument by Paul Noordhof. Noordhof argues that since micro-based properties supervene on their micro-bases, the same sort of causal exclusion argument should apply to micro-based properties as applies to any other supervenient properties. Considering Kim's insistence that it is 'obvious' that properties at different levels obviously don't compete, since micro-*based* properties possess 'novel' causal powers, not had by any properties of constituents, Noordhof writes:

> Water has causal powers that oxygen and hydrogen do not – like being able to douse fires. However, this point is just the familiar fact that anything requires cooperating circumstances to make a particular [causal]

[38] This is an important rhetorical point, because it helps us understand why many people don't accept Kim's solution to the first worry about inter-level competition. That was the claim that there is no worry about causation 'draining-down' to the micro *level* because causal competition takes place between different 'orders' of properties in the same individual. Block argues that the distinction plays no role in the actual argument (Block 2003, p. 144). Similarly, Thomas Bontly (2001) argues that Kim's response to the initial worry about inter-level competition misses the real issue of causal competition by sticking too closely to the letter of a particular definition of supervenience. What these critics have in mind, however, is the point addressed in Kim's *second* argument about inter-level competition. This new worry about inter-level competition concerns the pre-emption of macro-level properties by the micro-level constituents of those properties: if we avoid the worry about intra-level competition by identifying higher-order properties with lower-order 'structural' properties, then we can still worry that the causal efficacy of those structural properties is pre-empted by the entities and properties structuring them.

> contribution...There is no problem with relating the contribution of each microconstituent to the joint effect of the aggregate of the microconstituents...[and hence] there appears no need to postulate a property – in this case a property micro-based in micro-constituents having certain other properties – to capture the causal relationship. (Noordhof 1999: 112)

So while aggregates of microconstituents will have causal powers that no individual constituent has, those causal powers still result from the individual 'contributions' of each constituent. Hence we seem to have two possible accounts of a given causal interaction between higher-level entities: there is the causal interaction between the entities and their micro-based properties (as Kim insists) and the interaction between the lower-level entities and properties constituting those higher-level entities and properties. If the novel causal powers of higher-level micro-based properties can be attributed to the collaborative efforts of the constituents of those micro-bases, then those higher-level properties seem to be unnecessary no matter what 'order' they are.

Kim's immediate response to Noordhof's criticism is disappointing. He questions the idea of supervenience and determination Noordhof has in mind:

> But what is it for a micro-based property to have a 'micro-base' [as Noordhof insists will be the case]. Noordhof does not provide an explanation, and it is not a notion that I use anywhere in the target paper. But until we give a clear enough sense to the concept of micro-base, it is

not possible to assess the plausibility of Noordhof's assumption [that micro-based properties have micro-bases]. (Kim 1999, 116)

However, the concept of a micro-base seems to be at least as clear as that of a micro-based property. A micro-base is what is described by a micro-based property. If something has the micro-based property "having parts *x*, *y*, and *z* such that *x* is an oxygen molecule, *y* is an oxygen molecule, and *z* is a hydrogen molecule and *x*, *y*, and *z* stand in the appropriate relation *R* to one another", the micro-*base* of that property will consist of those properties 'is an oxygen molecule', 'is a hydrogen molecule', and so on. There may indeed be more work to be done in fleshing out this idea, but then there is clearly more work to be done in fleshing out the idea of micro-based properties in the first place.

Kim's more considered response is to *identify* micro-structural properties with their micro-bases:

> Difficulties of this sort [that is, difficulties of inter-level causal exclusion] do not arise for micro-based properties in relation to their constituent properties because the former do not supervene on the latter taken individually or as a group. Rather, they supervene on specific mereological configurations involving these microproperties – for a rather obvious and uninteresting reason: they are identical with these micro-configurations. (Kim 1998, pp. 117-8)

Thus micro-based properties are themselves *identical with* configurations of the properties they are based on, and of course identicals don't compete. So there are actually two identity claims in Kim's account of the supervenience argument and the worry about 'generalization': higher-level properties are to be identified with micro-based properties, and these are in turn to be identified with 'specific mereological configurations' of properties, so that micro-based properties are actually *composed of* the micro-properties they are based on. I'll spend the next section considering some prominent objections to these identity claims before discussing how the possibility of non-spatial parts might affect them.

Problems for Micro-Based Properties

While the idea that structural or 'micro-based' properties are actually composed of other properties is not unique to Kim, it's difficult to see how this suggestion is 'obvious and uninteresting': Kim's own initial definition of a micro-based property is of a property of "being completely decomposable into nonoverlapping proper parts, a_1, a_2,..., a_n, such that $P_1(a_1)$, $P_2(a_2)$,..., $P_n(a_n)$, *and* $R(a_1,..., a_n)$" (Kim 1998: 84). On this definition, micro-based properties sound like *specifications* of structure, akin to talk of having a particular decomposition. A micro-based property in this sense might be better understood as a *relation*: having a micro-based property is a matter of standing in the right sort of relation to a group of individuals and their properties. Kim's 'obvious and uninteresting' claim is that having a structural property should also be understood as having a property that is itself composed of other properties:

> Being a water molecule...is the property of having two hydrogen atoms and one oxygen atom in a such-and-such bonding relationship. A micro-

based property therefore is constituted by micro-constituents – that is, by the micro-parts of the object that has it and the properties and relations characterizing these parts. (Kim 1998: 84)

So the property *being a water molecule* is itself composed of the properties *being a hydrogen atom*, *being an oxygen atom*, etc. Admittedly, this *is* often the view associated with structural properties (see Armstrong 1978, for instance) but it is not an unproblematic or uninteresting one. One challenge to this account is the fact that many properties can be 'multiply based' in different micro-level structures. For instance, Thomas Bontly (2002) offers the following example:

> Consider the property with which we would identify temperature-in-gases: viz., mean statistical kinetic energy (MTKE). Many different distributions of energy-states to a gas's component molecules are sufficient for it to have a given MTKE, but obviously no such distribution is necessary since MTKE is a statistical mean...On the face of it, then, the 'multiple base-ability' of micro-based properties makes it unlikely that they can be identified with collections of microproperties/relations. (Bontly 2002: 87)

Ned Block (2003) presents a similar objection to Kim's claims about micro-based properties. Block accepts the first sort of identity claim Kim makes – the claim that higher-level properties are identical with micro-based properties – and presents his argument as an objection to the *second* claim about the identity between micro-based properties and 'specific mereological configurations' of component properties:

> Why can't micro-based properties be micro-based in alternative ways? Why isn't jade an example of a micro-based property, micro-based in both calcium magnesium silicate (nephrite) and sodium aluminum silicate (jadeite)? Recall that P is a micro-based property $=_{def}$ P is the property of being completely decomposable into non-overlapping parts $a_1 \ldots a_n$ s.t. $P_1(a_1) \ldots P_n(a_n)$ & $R(a_1 \ldots a_n)$. So my question is: why can't a micro-based property have more than one decomposition? (Block 2003: 145)

Despite some differences in presentation, I take this suggestion to be in the same spirit as Bontly's: if a micro-based property *did* have more than one decomposition, that 'multiple decomposability' would presumably prevent our being able to identify the micro-based property with its base, and thus the worry about inter-level competition would be revived. The property *jade* (if we can call this a property – I'll put aside any worries about the distinction between properties and *kinds*) is a micro-based property that is decomposable in at least two different ways: decomposed in one way, the property has a micro-base of calcium magnesium silicate (*nephrite*) while decomposed in another way, the property has a micro-base of sodium aluminum silicate (*jadeite*). These 'alternative decompositions' prevent the sort of identity between *jade* and a particular structural configuration of micro-level properties that Kim is counting on: if *jade* cannot be identified with a particular micro-structure, then the issue of competition between the property and its micro-base arises after all. And as long as we think that the higher-level property supervenes on that micro-base, the supervenience argument will apply: causal powers will 'drain away' and the specter of general epiphenomenalism returns[39]. For Block, this means that the supervenience argument leads to an absurd conclusion, and Kim's Causal Exclusion principle must be flawed[40].

In addition to worries about multiple 'base-ability' or decomposability, there is a second line of criticism raised by Bontly, this time directed at the very idea of 'micro-based' properties. Bontly argues that examples of isomerism in chemistry, where distinct chemical compounds are composed of the same

---

[39] At least I *think* this is the intent of Block's argument. Block's argument is brief, and it is obscured by two factors. The first is the unavoidable confusion that arises when speaking of the identity of properties, since it is almost impossible to say anything meaningful about such identities without speaking as if two things were being 'identified'. Of course, if 'they' are identical, they are not two. The second factor is that it is unclear whether in using chemical formulas or proper chemical names, Block means to refer specifically to molecular-level features or simply to chemical types. Of course, chemical formulas *can* be interpreted as structural descriptions of molecules ('H_2O' can be taken to refer to molecules composed of two hydrogen and one oxygen atoms, bonded together in a particular way), but they needn't be: 'H_2O' can also be taken to refer to a particular chemical *substance* that is composed of two parts *hydrogen* and one part *oxygen*; on this interpretation, the formula identifies what we might call a chemical *kind* without implying anything about the 'molecules' composing that kind, or even that there *are* molecules. For a detailed discussion of this point, see Needham 2000.

[40] Sydney Shoemaker (1980) suggests a similar argument for why properties cannot be *identified* with their associated 'causal powers' (the argument is due to Richard Boyd, and it originally appears as a postscript to Shoemaker's paper). Shoemaker considers the possibility that there could be two compound substance X and Y that are composed of different 'basic physical elements' A, B and C, D respectively, yet "although composed of different elements [these compounds] behave exactly alike under all possible circumstances...the property of being made of X and the property of being made of Y share all of their causal potentialities" (Shoemaker 1980: 254). If properties (which apparently are to include 'being made of X') were *identical* with their associated causal powers, then 'being made of X' and 'being made of Y' would be the same property, even though clearly they should be distinct, since they are made of different elements. Hence properties cannot be identified with their associated causal powers.

elementary constituents, show that Kim's claims about identity between micro-based properties and their micro-bases cannot be maintained:

> Methane (CH_5), propane (C_3H_8), and butane (C_4H_{10}) are different gases involving exactly the same microphysics; therefore, *being methane* can't be the mereological sum of those microproperties, for otherwise *being methane* would be identical to *being propane*, which it is not. (Bontly 2002: 89)

Bontly's argument reflects arguments originally advanced by David Lewis (1986) against the concept of 'structural universals' suggested by David Armstrong (1978) (and Kim explicitly acknowledges this as the source of his own concept of 'micro-based' properties; see Kim 1998: 84). To conceive of properties as universals is to view instances of a property as numerically identical, multiply located things. This is not the only way of conceiving of properties: for instance, we might think of properties as numerically distinct, exactly similar tropes, or we might defend some form of 'nominalism', in which case we would deny that there are any properties and insist that there are only particular things. However, Armstrong's originally account of structural properties explicitly treated them as universals related to each other as part to whole. Lewis' criticism of Armstrong's idea actually begins with a more generic conception of a structural universal simply as "a property such that anything that instantiates it must have proper parts; and there is a necessary connection between the instantiating of the structural universal by the whole and the instantiating of other universals by the parts" (Lewis 1986: 81). To use a neutral term, Lewis says structural universals 'involve' other universals, and the challenge is to explain this 'involvement' in a way that accounts for the fact that things instantiating a structural universal have proper parts that instantiate other universals.

Lewis actually considers, and rejects, three different conceptions of what this involvement might be. The first to be rejected is the 'magical' conception on which we regard the involvement relation as primitive: structural universals simply *do* ('magically') make it necessary that anything instantiating them have a particular structure. Unsurprisingly, Lewis finds this conception unsatisfying, concluding that "I might say that the magical conception carries an unacceptable price in mystery; or perhaps I would do better to deny that there is any *conception* here at all, as opposed to mere words" (Lewis 1986: 102)[41]. Lewis also rejects a second, 'linguistic', conception of structural universals, and his reasons for rejecting it are interesting in the light of the discussion in this chapter. According the linguistic conception, the structure associated with a structural universal is be captured in the same way that the syntactic structure of language can be captured using set-theory (here Lewis cites his (1970) work on 'generalized semantics'). On this account, structural properties are identified with sets built from 'atomic' universals, where the sets are constructed in such a way that they can be

[41] Obviously the magical conception would be vulnerable to the supervenience argument as well, since according to it all properties – including structural properties – are distinct, atomic individuals; see Lewis 1986: 100. For instance, if the property *methane* magically drags around the distinct properties *carbon*, *hydrogen*, and *bonded* (to borrow Lewis' example) then either we have to admit that properties like *methane* bestow novel causal powers upon their instances (and hence – in this case – deny the causal closure of the atomic level) or we have to accept that structural properties are epiphenomenal and carry with them no associated causal powers.

interpreted as predicates in a language which is defined so that something satisfies a predicate just in case it has the appropriate structure. However, these set-theoretic constructions must ultimately depend upon universals that are mereologically simple: on the linguistic conception, structural properties are to be built up from 'atomic' properties. But Lewis considers the strongest reason for supposing that there *are* structural universals to be the thought that there might be *no* atomic properties, because the world might consists of endless structure in something like the way Block imagines:

> I take this...to be the weightiest [reason] by far. Infinite complexity [that is, endless divisibility] does seem, offhand, to be a genuine possibility. I might contemplate treating it as negotiable: if structural universals are trouble, and simple universals retain their charm, so much the worse for the alleged possibility that there are no simples! But that seems objectionably high-handed, if not downright intolerable. (Lewis 1986: 86)

This linguistic conception of structural universals also seems to leave the 'structure' of structural properties unsatisfyingly arbitrary. As Lewis says, on the linguistic conception, atomic properties comprise only *some* of the 'vocabulary' of the language being constructed: "we also need logical words – the usual connectives, quantifiers, and variables – and we need mereological predicates of identity, inclusion, and overlap. *These words can be anything the resources of set-theoretic construction have to offer.*" (Lewis 1986: 87, emphasis added) But it is just this abundance of choice as to what structural properties *are* that leaves the linguistic conception sounding hollow: even if we were to accept the identification of properties with sets, if structural properties are to be real features of the world, we shouldn't be able to identify them with sets constructed out of just *anything*[42].

It is the *third* conception that corresponds to Kim's suggestion that micro-based properties are composed of their base properties. Lewis calls this the 'pictorial' conception: a structural universal actually has other universals as parts, in the same way that an entity instantiating that universal will have entities instantiating those other universals as parts. Thus structural universals 'picture' the structure of their instances. Lewis' main objection to the pictorial conception is similar to that given by Bontly: since by their definition universals are things that are themselves 'repeatable', appearing in their many instances at the same time, there seems to be no way to account for the fact that the same universal might count as more than one part of a given structural universal. For instance, a molecule of methane consists of one carbon atom and four hydrogen atoms in a particular bonding arrangement. But if *methane* is a structural universal, itself

[42] If these objections aren't enough, there is the additional 'peculiarity', apparently noted by John Bigelow (see Lewis 1986: 89) that on the linguistic conception a structural universal will be 'present' wherever the atomic universals constituting it are present. So, admits Lewis, "[the] universal *methane* will be wholly present, because its simples are, not only where there is a methane molecule but also wherever there is any sort of molecule that is made of carbon and hydrogen bonded together". This follows as a result of Lewis' insistence that, in order to accommodate the assumption that a universal is located wherever its instances are (that is an 'in rebus' account of universals), sets must be located wherever their members, or the members of their members, are (see Lewis 1986: 88). Lewis' shrugs off this oddity, noting that "we just have to take care to distinguish instantiation from mere presence" (Lewis 1986: 89).

composed of the universals characterizing its parts, and four of those parts are hydrogen atoms, characterized by the universal *hydrogen* (itself a structural universal, incidentally) then the universal *methane* should have four parts that are each the universal *hydrogen*. But there is only one universal *hydrogen*: how can something (*methane*) have the same thing as a part, four times over? Chemical isomerism only reinforces the point: since different isomers such as *methane*, *propane* and *butane* are distinguished only by the number of 'carbon' and 'hydrogen' parts they have, Lewis concludes that there is no hope of accounting for such properties purely in terms of part-whole relations between universals.

However, it isn't clear that Kim's claim about micro-based properties is intended to be one about *universals* or property *types*, and hence one way of responding would be to insist that the claim is about specific property instances or 'tropes', rather than universals. In that case, the problem of isomerism would be avoided: a specific instance of the structural property *methane* is composed of specific instances of the properties *carbon* and *hydrogen* – one of the first, and five of the second. Kim doesn't discuss this suggestion explicitly, though something like this is suggested by his remark that micro-based properties are "constituted by micro-constituents – that is, *by the micro-parts of the object that has it* and the properties and relations characterizing these parts" (Kim 1998: 84, emphasis added). It isn't clear exactly what it would mean for a *property* of something to be partly constituted by the *parts* of that thing, but we might interpret this as suggesting that the parthood claim is directed at specific instances of properties, rather than at 'universals'.

This sort of response would also fit with the response Kim does give to Block's argument about the multiple decomposability of *jade*. Kim's main response to Block's argument is to point out that in general in causal relations it is *instances* of properties that we are interested in, and particular instances of *jade* will presumably not be multiply decomposable in the way Block suggests:

> In spite of jade's multiple composition, each instance of jade – that is, each individual piece of jade – is either jadeite or nephrite, and I don't see anything wrong about identifying *its* being jade with *its* being nephrite (if it is nephrite) or with *its* being jadeite (if it's jadeite)...All we need is identity at the level of instances, not necessarily at the level of kinds and properties; causation after all is a relation between property- or kind-instances, not between properties or kinds as such. (Kim 2003: 168)

A similar response could be given to Bontly's first example, the 'multiple base-ability' of a property such as mean statistical kinetic energy. While many different distributions of energy-states might be sufficient for a given *type* of mean statistical kinetic energy, that alone doesn't prevent identities between the macro-level mean statistical kinetic energy and micro-level energy distributions in particular instances[43]. Note that Bontly does consider how Kim might respond to this problem in terms of *functional* properties: the suggestion is that a statistical property would be identified with the second-order property of having one of a

[43] I say 'that alone' because I share with Bontly some skepticism about such identity claims; however, pointing out the variability at the micro-level *alone* doesn't seem to be enough to create a problem for Kim. Perhaps a more detailed study of such properties would reveal a genuine problem for identifying those macro-level properties with particular micro-level distributions even on the 'instance for instance' basis Kim suggests.

number of possible first-order property satisfying a particular functional description (i.e., functioning as a particular mean statistical kinetic energy state). However, Bontly assumes that any such account must involve 'species specific' reductions akin to those Kim imagines will take place in cases such as the multiple realizability of mental properties such as pain. For instance, supposing that 'pain' has distinct realizations in different species of animals, Kim supposes that the reduction of pain to its physical realizer will involve different reductive relationships for each species, so that "the reduction consists in identifying M with its realizer Pi relative to the species or structure under consideration…Thus M is P1 in species 1, P2 in species 2, and so on." (Kim 1998: 110). The outcome of these identifications is that "multiply realized properties are sundered into their diverse realizers in different species and structures" (Kim 1998: 111). However, Bontly argues that in the case of temperature-in-gases, multiple 'base-ability' is not so easily accommodated:

> There is no way to reduce MTKE-in-neon to one micro-configuration or MTKE-in-freon to another, for there simply are no kinds of gases throughout which a micro-based property like MTKE can be correlated one-to-one with a specific configuration of microproperties. A determinate value of MTKE [can] be realized by many different distributions of energy-states to molecules in different types of gases; it can also be differently realized by many distributions of energy-states to molecules in the same type of gas. (Bontly 2002: 88, emphasis in original)

However, I don't see any reason to expect reductions across *these* species in the case of MTKE: the point of Kim's suggestion about species of animals is that given that different realizations of the property 'pain' show similarities within species (we assume) and differences between species (we assume), it is reasonable to expect 'species-specific' reductions but not 'universal' reductions. It's true that, like animals, gases come in different 'species', but *that* difference isn't what leads us to expect statistical properties such as MTKE to be 'multiply realized'. The only sort of *identity* claim available to Kim in this sort of case would seem to be specific identity between particular instances of MTKE and micro-states of gases. So while Bontly is right that the sort of species-specific reductions he describes are not available in the case of gases, this doesn't really defeat Kim's suggestion, since there was no reason to expect *those* kinds of reductions in the first place.

Conclusion: Genuine Multiple Decomposability

None of these objections to Kim's claims about property identities and reduction seem to be entirely effective. Even if Block and Bontly are right about the difficulties in identifying *properties* in the way Kim suggests, similar identities between property *instances* seem to be unproblematic. What's more, none of their criticisms seem to get at the heart of the matter in the debate about reductionism: even given the kind of 'multiple *micro*-base-ability' Bontly and Block describe, there still doesn't seem to be any need to postulate 'higher-level' properties or any reason to think, as Weinberg put it, that higher-level properties "have to be dealt with in their own terms".

What we really need to challenge Kim's claim's about micro-based properties are cases where it is particular *instances* that are decomposable in different ways, together with some reason for thinking that this multiple decomposability blocked identities between the composite property and any particular decomposition. Notice that in general, multiple decomposability doesn't necessarily prevent the sorts of identities Block is trying to prevent: a patio, for instance, might be 'decomposable' into a bunch of stones separated by layers of fine sand, and 'alternatively' decomposable into a vast array of particles of rock, *and* decomposable into molecules, etc. These sorts of 'multiple decompositions' are compatible, though, since we have an account of how the elements of one decomposition (the patio stones, say) can themselves be decomposed into elements from another decomposition (the particles of rock, say). To challenge the identity claims in the way Block suggests, we need alternative decompositions that 'cross-cut' each other so that the elements on one decomposition cannot be decomposed into elements on the other (nor vice versa). In the next chapter, I'll argue that decompositions into 'non-spatial' parts are much more compelling examples of the kind of 'multiple decomposability' Block seems to have in mind.

III

Parts on Multiple Scales

Introduction

Kim's response to the generalization argument depended upon the claim that higher-level properties can be identified with 'micro-based' properties, where these in turn are identical with 'specific mereological configurations' of lower-level properties. Block's response to Kim's argument questioned the identity between the 'micro-based' property and a *specific* mereological configuration, but didn't dispute the idea that higher-level properties in general should be identified with *micro*-based properties. In either case, the assumption is that higher-level properties are to be explained by micro-based properties of one sort or another, and the only issue is over whether or not these 'macro' and 'micro' properties can be identified. Clearly this fits with the standard view of physicalism outlined in chapter one: we defend physicalism on the grounds that everything is composed of micro-level entities, and accordingly we expect explanations of higher-level properties of wholes to be given in terms of the micro-level properties of those wholes. That is, we expect – at least in principle – relatively 'micro' decompositions to be the explanatorily relevant ones. However, while it is clear that higher-level entities *have* such 'micro' decompositions, we might wonder whether or not these truly are the relevant decompositions for explaining higher-level properties. A fortiori, we might wonder whether Kim's claims about *identity*, or *reduction*, between higher-level properties and micro-level properties can be maintained.

In this chapter, I'll investigate this claim about properties in the context of physics: can all *physical* properties be identified with 'micro-based' properties? To answer this question, I'll examine how physical systems are often described using 'multi-scale analysis', where the behavior of the system as a whole is explained in terms of the interaction between component systems operating on different scales. As I suggested in chapter one, these decompositions apparently appeal to parts that are not *spatial* parts of a system: hence the relevant properties in these explanations are not always properties of smaller entities contained within the system, but are instead properties of sub-systems that in some sense overlap with the composite system in space and time.

Examples of multi-scale analysis will yield parts that don't fit into the received hierarchical view of the physical world Kim presumes. Recall Kim's analysis of what counts as 'physical': the fundamental physical entities and properties are those found in microphysics, and all other (non-fundamental) entities and properties consist of aggregates or 'specific mereological configurations' of these. This is clearly a 'bottom-up' view of physical composition and parthood: we start with some paradigmatically physical entities and properties and then count anything composed of those things also as 'physical'. The examples from multi-scale analysis, however, suggest a 'top-down' approach: for instance, systems are often decomposed into subsystem operating on different temporal scales, so that that behavior of a system is represented as the result of 'fast' and 'slow' component system; but this sort of decomposition doesn't involve distinguishing between 'slow' and 'fast' *particles*. Instead, decomposing a system in slow and fast components is a different sort of decomposition, which I'll call 'non-spatial' decomposition. The same system will *also* have an 'ordinary' micro-decomposition of the sort Kim imagines (or so I'll assume), but being decomposable in this way doesn't preclude being decomposable in other ways as well.

Levels and Scales

Like most philosophical discussions, Kim's account of the entities and properties found in science is put in terms of 'levels'. However, a more common way of distinguishing between different sorts of descriptions in science is in terms of *scales*. In many ways, the two seem simply to be alternative ways of talking about the same thing. Just as we can ask about macro and micro level entities and their properties, we can ask about macro and micro scale descriptions of things and their properties. Kim's suggestion that higher-level properties can be identified with, or in some way 'reduced to', 'micro-based' properties can naturally be translated into the claim that (relatively) macro scale descriptions can be identified with or in some way reduce to, micro scale ones. And the worry that reductionism undermines the reality of the higher-level properties reduced can be put in terms of a worry that it is only at the 'micro' scale that we find the true description of the world, and that descriptions at other scales are only arbitrary approximations *we* need to make use of because of our cognitive limitations: the scales we use to describe a given system are chosen simply to suit our own purposes, and that the choice is arbitrary in the sense that the system could have been described on any other scale if *that* scale was more convenient for us.

Scientists themselves don't treat the choice of scales as arbitrary. Indeed, they often speak of them as representing natural or characteristic scales for describing a given system. And choosing to describe something on certain scales can often lead to a successful solution to the governing equations, whereas other choices leave the equations not just difficult to solve, but intractable. What's more, as I mentioned in chapter one, it is often necessary to describe a system simultaneously on *multiple* scales in order to get an accurate representation of that system For example, describing the diffusion of water through a composite medium may require a description of *both* the 'macro' level diffusion, where the medium is described as continuous with no 'gaps', and the 'micro' level diffusion, where the medium is described as granular, consisting of a lattice of impermeable boundaries and open gaps where fluid can freely flow[44]. It might be tempting to think that these multi-scale descriptions are just giving different views of the same underlying reality: the *real* system is given by the micro-scale description, and all other descriptions are more or less 'fuzzy' approximations of that true description. But what's interesting about multi-scale analysis is that it is only through the *combination* of such distinct views that we are able to find an accurate description. This seems to call into question the idea that there is *any* single 'fundamental' scale.

[44] See Hornung 1997.

An important difference between philosophical talk of 'levels' and talk of scales in science is that while talk of levels typically only includes distinct *spatial* levels, scientists frequently distinguish between different *temporal* scales as well. For instance, describing chemical reaction rates may involve *both* 'fast' scale interactions between individual molecules, and 'slow' scale changes in thermodynamic properties such as volume and temperature. The absence of talk of temporal scales is a curious feature of the philosophical literature on reduction. Philosophers tend to treat all properties 'statically', as if the relations of interest – whether in terms of 'mere determination', explanation, or reduction – were between instantaneous states of a system. In reality, many properties of interest are not static at all, but characterize the way things behave over time. For instance, 'being a harmonic oscillator' is plausibly a genuine property of a variety of types of systems, and the ability to unify descriptions of these systems under such a property is important for explanations.

Scientists tend to talk of these sorts of properties as temporally extended properties of *processes*. The distinction between entities and processes has a long and somewhat controversial history in philosophy, and trying to defend a full account of processes would take me too far from my main interest here[45]. Like most discussions in contemporary philosophy, Kim's account of physicalism, physical entities, and physical properties makes no mention of processes. However, this *lacuna* alone needn't present any particular difficulty for Kim: I think we can extend his account of entities and properties fairly easily and

[45] See Rescher 2000 for a general discussion of the concept of process; Paul Needham (1999) offers a compelling defense of the need to include processes in our understanding of science.

naturally to include processes by generalizing the idea of a 'structural property'[46]. The importance of structure is clear in the philosophical literature on process. For instance, Nicholas Rescher defines a process variously as "an actual or possible occurrence that consists of an integrated series of connected developments unfolding in programmatic coordination", "a complex of occurrences – a unity of distinct stages or phases", "[a] complex of occurrences [with] a certain temporal coherence and integrity", and as something having "a structure, a formal generic patterning of occurrence through which its temporal phases exhibit a fixed format" (Rescher 2000: 24). From this it seems clear that a central feature of a process is its having a certain *structure*: as Rescher says, a process is not 'just one darn thing after another' (Rescher 2000: 27).

Structural properties are usually conceived of as simply concerned with 'spatial' structure: having a structural property is a matter of having particular parts with particular properties, standing in particular relations. We can extend this idea of structure to include *temporally* structured properties: a 'spatio-temporal' structural property characterizing a process is a property of having particular parts with particular properties standing in particular relations *at particular times*.

With this extended idea of structural properties in mind, I'll now turn to some examples where representing the behavior of a system requires decomposing that system into component system operating on different time scales. It's important to note that though these examples are distinct from the

[46] If this account isn't adequate, then the standard view of physicalism needs an even more serious overhaul than what I'm suggesting here. I'll leave that question for further work.

sorts of examples Kim considers because they involve processes extended in time (whereas Kim considers only things with 'static' properties), that fact *alone* doesn't play a role in the argument here. It *could* be that representing the properties of a system through time involves decomposing that system into ordinary spatial parts and considering *their* behavior through time (as, for instance, we might best explain the behavior of a clock in terms of the behavior of its component gears and levers); but in the cases I'll discuss, it does not.

Multi-scale Analysis

As I mentioned in chapter one, multi-scale analysis is used in a wide variety of problems ranging from fluid mechanics to geophysics to the kinetics of chemical reactions. These applications are typically very technical, but we can gain some appreciation of the basic idea of multi-scale analysis by considering textbook illustrations of how it can be applied to much simpler systems. In chapter one, I mentioned how a system modeled as a 'relaxation oscillator' (illustrated again in figure 3.1 for convenience) can be decomposed into independent fast and slow subsystems, and that doing this allows us to give an accurate description of the system we could not otherwise attain.

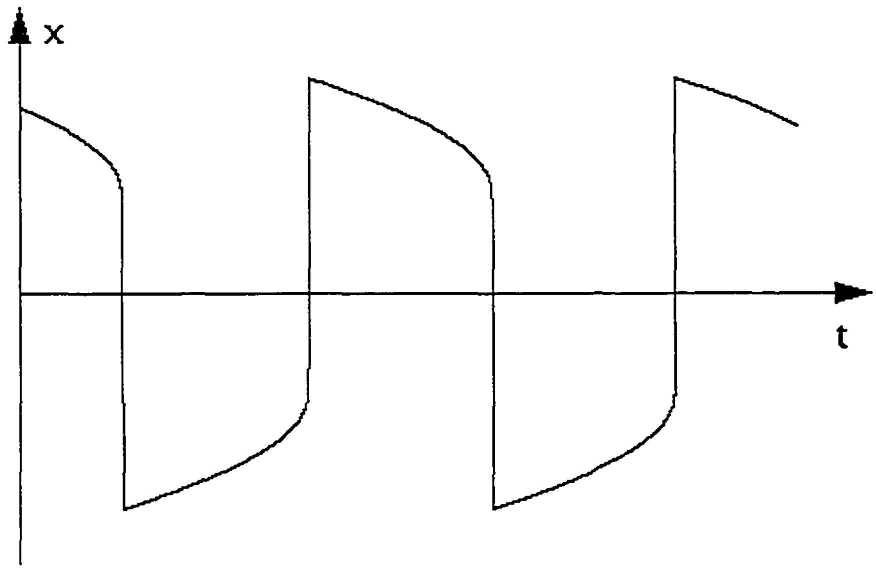

Figure 3.1 Behavior of a Van der Pol relaxation oscillator.

The idea of the fast/slow decomposition is to separate out apparently distinct processes at work in the oscillator: the 'slow' process produces the steady rise or fall in amplitude in the elongated horizontal edges above, while the 'fast' process produces the sudden rises or falls in the near vertical spikes. This sort of multi-scale analysis comes from the field of *perturbation theory*. This is a branch of mathematics that is widely applied to problems where the equations describing a system are unsolvable, but where very accurate approximations can be found by assuming that those real systems are 'perturbations' of more readily describable ideal systems. Unfortunately, the perturbation account of the relaxation oscillator is actually still quite complicated, so I'll illustrate the technique using a more mundane example that will illustrate the same principles.

I'll begin with the basic idea of perturbation theory. Perturbation theory deals with complex equations by exploiting the fact that the governing equations of a system can often be written to contain a small parameter multiplying part of the equation. Usually the equation becomes more readily solvable if the term multiplied by that parameter is ignored (typically, if the value of the parameter is zero). We can call the equation itself the *full* equation, and the limit of the full equation as this parameter approaches a certain limit (usually zero) the *reduced* equation. The perturbation approach to finding a solution to the full equation is to assume that this solution is a perturbation of the solution to the reduced equation. Finding a perturbation solution usually involves representing that solution as a series of terms, the first of which represents a 'core' solution, while subsequent terms represent corrections to this initial approximation.

To illustrate the basic idea of a perturbation solution to an equation, consider the following purely algebraic problem (from Nayfeh 1973: 2-3). Suppose we want to find a function u such that:

(1) $$u = 1 + \varepsilon u^3$$

If ε were zero, then u would be 1, but suppose that ε is not zero, but is instead some small parameter. We can find a solution to (1) by supposing the smallness of ε means that the solution we want – u – is a perturbation of this 'ideal' case where the second term in the equation is not present. Accordingly, we suppose that u has the following expansion:

(2) $$u = 1 + \varepsilon u_1 + \varepsilon^2 u_2 + \varepsilon^3 u_3 + \ldots$$

At this point the u_is are unknown terms: what we're doing is supposing that we can represent u as an expansion of this form, in the hope that this expansion will turn out to be easier to solve that (1). We then substitute this expansion (2) for u in equation (1) to get:

(3) $$\varepsilon u_1 + \varepsilon^2 u_2 + \ldots = \varepsilon(1 + \varepsilon u_1 + \varepsilon^2 u_2 + \ldots)$$

which is algebraically equivalent to

(4) $$\varepsilon(u_1 - 1) + \varepsilon^2(u_2 - 3u_1) + \varepsilon^3(u_3 - 3u_2 - 3u_1^2) + \ldots = 0$$

where we've gathered together all of the terms associated with each power of ε. We then solve for these terms to find our approximation. For instance:

(4a) $$u_1 - 1 = 0$$

(4b) $$u_2 - 3u_1 = 0$$

(4c) $$u_3 - 3u_2 - 3u_1^2 = 0$$

So $u_1 = 1$, $u_2 = 3$, $u_3 = 12$, etc. Now we can return to (2) and we have:

(5) $$u \sim 1 + \varepsilon + 3\varepsilon^2 + 12\varepsilon^3 + \ldots$$

This equation is now our approximation to the solution of (1). As $\varepsilon \rightarrow 0$, $u \rightarrow 1$, and if $\varepsilon > 0$ (but still small, $\varepsilon \ll 1$), (5) gives us a good approximation of u. The symbol '~' is used instead of the 'equals' sign because the expansion is not expected to give us an *exact* solution for u, but instead only to be *asymptotic*. In an asymptotic expansion, each term can be regarded as a minor correction to the approximation given by the preceding terms. While an asymptotic expansion is usually not convergent (that is, considering more and more terms does not usually converge on a finite value in the limit), asymptotic expansions do guarantee boundaries on the error between the approximation and the exact solution. What's more, asymptotic expansions give very good estimates using only a few terms, making them very useful for doing actual calculations. Finding an asymptotic expansion guarantees that the expansion not only well approximates the solution in the short term, but continues to do so in the long term as well (that is, as the independent variable in the equation grows larger).

The approach just described is known as 'regular' perturbation theory. Many problems, however, are not regular, and cannot be solved without altering the perturbation method. These problems must be solved using what's called *singular* perturbation theory. A prominent feature of singular perturbation theory is that otherwise unmanageable equations can often be solved by introducing distinct temporal or spatial scales. In effect, what this does is to model the behavior of a system by treating it as two or more *distinct* systems that combine to produce the observed behavior we are trying to understand. These distinct sub-systems are *parts* of the observed system, but not spatial parts.

Let's begin with an example where the 'regular' perturbation approach fails[47]. A standard example of this is what's known as a damped harmonic oscillator: oscillators can model all sorts of behavior, but a simple case of such an oscillator would be a mass suspended from a spring. If the mass is disturbed, it oscillates up and down. However, because of various 'damping' forces such as friction, the mass-spring system will eventually settle down into a steady state. We are interested in modelling the position of the mass as a function of time. The perturbation approach begins with an equation of the form:

(6) $$my'' + cy' + ky = 0$$

[47] The following example is taken from Holmes 1999: 25-29.

where the position y is a function of time t, m is the mass, c is the damping constant, and k the spring constant, and the strokes on y represent differentiation with respect to time. To find a solution to (6), we begin by introducing a parameter $\varepsilon = c/mk^{1/2}$. This allows us to rewrite (6) so that a single parameter ε represents the various relevant constants, and does so in a way that allows us to treat ε as a small 'perturbing' parameter:

(7) $$y'' + \varepsilon y' + y = 0$$

On the 'regular' perturbation approach – that is, approaching the problem *without* decomposing the behaviour into component systems on different scales – we try to solve (7) by finding a perturbation expansion in powers of ε:

(8) $$y(t) = y_0(t) + \varepsilon y_1(t) + \ldots$$

We now need to solve for the components y_0, y_1, etc. to find the actual series. To do this, we first substitute (8) – the series expansion of y(t) – into (7) and then solve for powers of ε. The resulting equations are:

(9a) $$y_0(t) = \cos(t)$$

(9b) $$y_1(t) = -\tfrac{1}{2}t\cos(t) + \tfrac{1}{2}\sin(t)$$

meaning that (8) becomes

(10) $$y(t) \sim \cos(t) + \tfrac{1}{2}\varepsilon[-t\cos(t) + \sin(t)] + \ldots$$

For the perturbation expansion to work, we want the first term in the series to represent the 'core' solution, and each subsequent term to represent increasing minor correction to the approximation given by all preceding terms: informally, at least, this is what it means for (10) to be 'asymptotic'. However, it turns out that the expansion in (10) is not asymptotic: as t grows larger, the second term, which initially represented small corrections to the core solution given by the first term, start to grow too fast. As t increases, the second term grow uncontrollably, meaning that the right hand side of the expansion (10) can no longer be regarded as representing a good approximation to y(t). And the error of the approximation becomes unbounded no matter how small the perturbation parameter ε is – that is, no matter how small we take the 'corrections' to be in the first place. Mathematically, the problem is that the naïve expansion contains what are known as 'secular terms'. These are terms that grow too fast, relative to the growth of the other terms in the expansion. Figure 3.2 illustrates the difference between (5) and the exact solution:

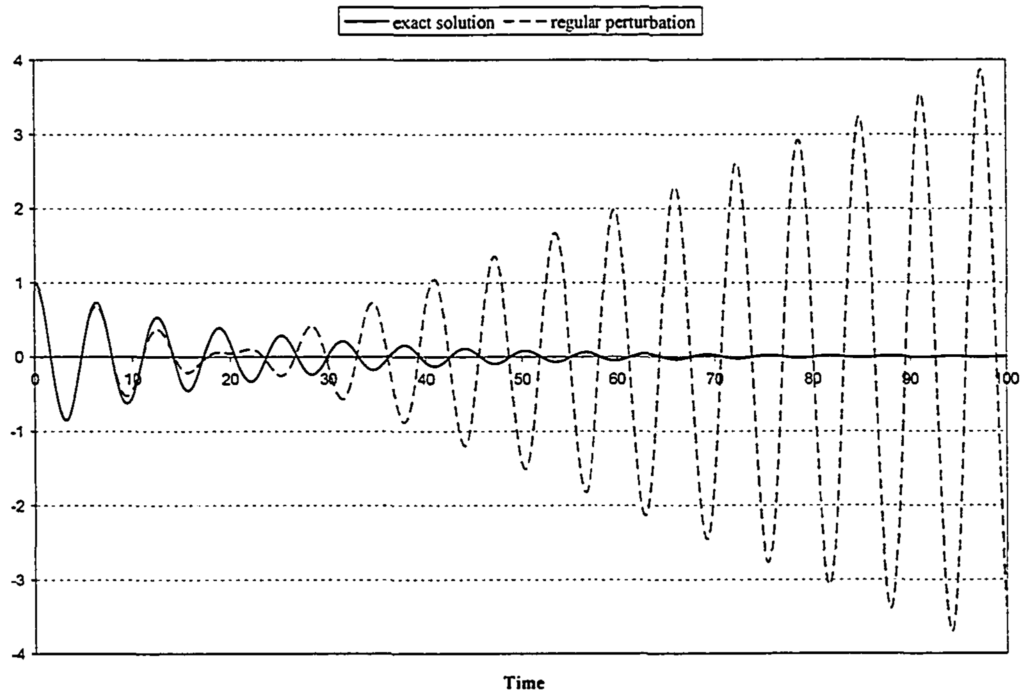

Figure 3.2 The failure of a regular perturbation expansion.

Now consider the solution we get by introducing multiple scales. To eliminate the secular terms in the above solution, we introduce distinct 'fast' and 'slow' time scales, the original scale $t_F = t$, and a 'slower' time scale $t_S = \varepsilon t$ (remember that ε is supposed to be small, so t_S will be some fractional value of t). Though these are both related to our original time scale t, for the purpose of solving the equations they are regarded as independent. The new time scale allows us to control the secular terms which eventually lead our approximation to diverge from the actual solution, while at the same time maintain what was 'right' or 'almost right' in the original approximation. Our multi-scale expansion assumes the form:

$$y(t) = y_0(t_F, t_S) + \varepsilon y_1(t_F, t_S) + \ldots \tag{11}$$

We substitute (11) into (7) and again solve for y_0. This gives us:

$$y_0 = A_0(t_S)\cos(t_F) + B_0(t_S)\sin(t_F) \tag{12}$$

where A_0 and B_0 are undetermined functions of the *slow* scale. We can determine these functions by constraining the expansion (12) to be asymptotic. The result we get is:

$$A_0(t_S) = e^{-t_S/2} \tag{13a}$$

$$B_0(t_S) = 0 \tag{13b}$$

So the solution we get using multiple scales is:

$$y(t) \sim e^{-t_S/2}\cos(t_F) \tag{14}$$

Equation (14) represents the behaviour of the system as a whole as a combination of the behaviour of a fast component operating on the times scale t_F and a slow component on the time scale t_S. Figure 3.3 illustrates the solution that results from the multi-scale approach:

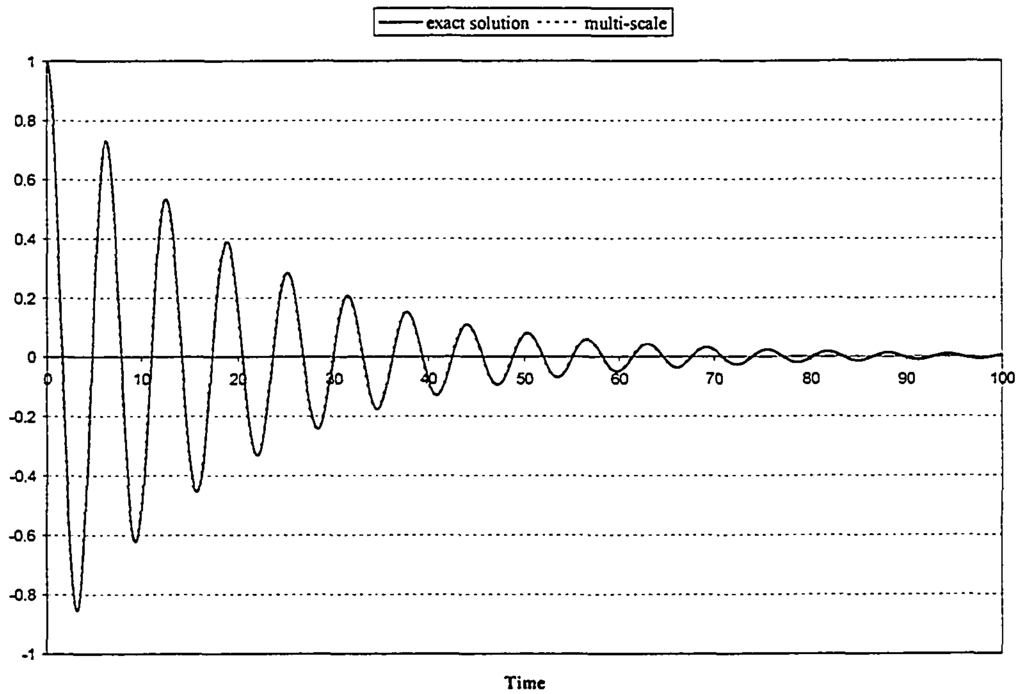

Figure 3.3 The multi-scale solution.

Clearly, the multi-scale approach gives a very good approximation of the exact solution to (6): in fact, the approximation is so close, the two distinct lines can barely be distinguished.

In this case, the component slow and fast systems are represented by the equations $y_S = e^{-t_S/2}$ and $y_F = \cos(t_F)$. The following graph shows how these components represent the distinct processes involved in the oscillation: the rapid fast-scale oscillations and the more gradual slow-scale decay in amplitude:

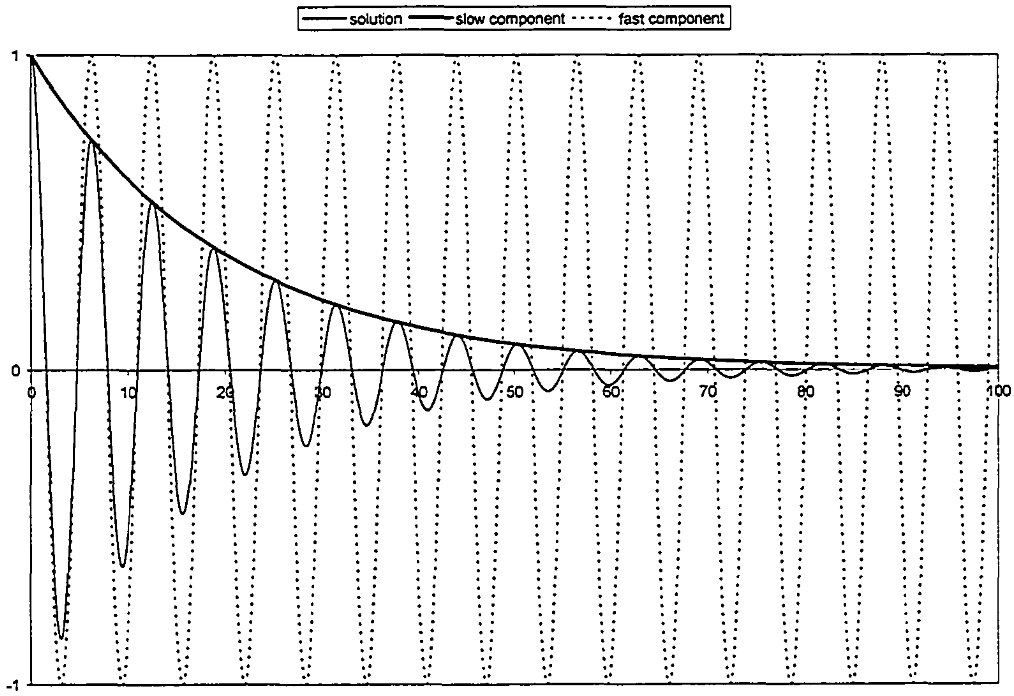

Figure 3.4 Fast and slow components of multi-scale solution.

In this example, the relationship between the component systems and the observed system can be thought of as 'multiplicative' composition, since the observed behaviour is mathematically represented as the *product* of the equations describing the components. This sort of composition is typical of perturbation examples using the method I've just described. However, we can also find perturbation examples where the component systems are represented as 'additive' parts, which might intuitively seem more compelling as examples of genuine decomposition into parts[48]. For example[49], suppose the equation governing a system is:

[48] For a discussion of the mathematics of multiplicative and additive composition, see van Dyke 1975, pp. 94-97. Properly speaking, multiplicative parts are characteristic of the perturbation method known as "the method of multiple scales", while additive parts are more commonly found

(15) $$\varepsilon y'' + y' + y = 0$$

where y is a function (of time t), and y′ represents differentiation with respect to t. This is just a 'toy' example for illustrative purposes, but suppose that this equation is meant to model some oscillatory behavior in a physical system. Figure 3.5 illustrates the type of behavior described by such an equation:

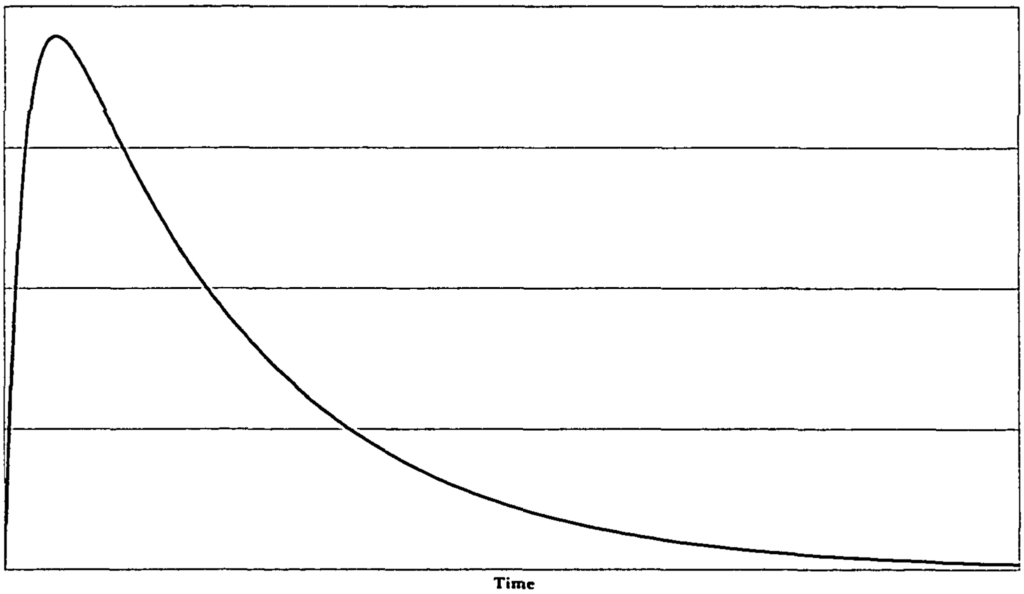

Figure 3.5 Sample oscillatory behavior.

The behavior is characterized by an initial 'fast' cycle of change in amplitude, followed by a longer 'slow' decay. The coefficient of the highest derivative – ε – is a small parameter (in general, governing equations can be rewritten in this type of form by choosing appropriate units for the various

in the method known as "the method matched asymptotic expansions", though neither method makes use *solely* of either form of composition. In either case, the basic principle of representing the behavior of a complex system as the result – additive or multiplicative – of the behavior of component systems operating on distinct scales is the same.

[49] This example is drawn from Nayfeh 1973: 111-112; I've altered the example to apply to temporal scales rather than to spatial ones, though the principles it illustrates are the same.

quantities involved; however, equations where the parameter multiplies the highest derivative represent a special class of equations particularly suited to the technique I'm about to describe). We are looking for a solution for y that satisfies (15) and the boundary conditions:

(16a) $$y(0) = \alpha$$

(16b) $$y(1) = \beta$$

As $\varepsilon \rightarrow 0$, (15) reduces to

(17) $$y' + y = 0$$

However, in general, it is not possible for equations such as (17) to satisfy both boundary conditions given in (16), since those boundary conditions apply to a second-order equation (15), while (17) is a first-order equation (the highest derivative in the equation is a first-order derivative). Maintaining only the second boundary condition, we can give the solution for (15) as $\varepsilon \rightarrow 0$

(18) $$y^S \rightarrow \beta\, e^{1-t}$$

The superscript indicates that this is the *slow* component of the solution for (15): this describes the behavior of the system once the slow component dominates. Figure 3.6 shows the behavior of the slow component and the 'total' system together (note that in figure 3.6 the graph has been slightly rescaled from figure 3.5 to help emphasize the difference between the slow component and the total system's behavior: also note that for much of the graph, the two coincide so closely that they cannot be distinguished).

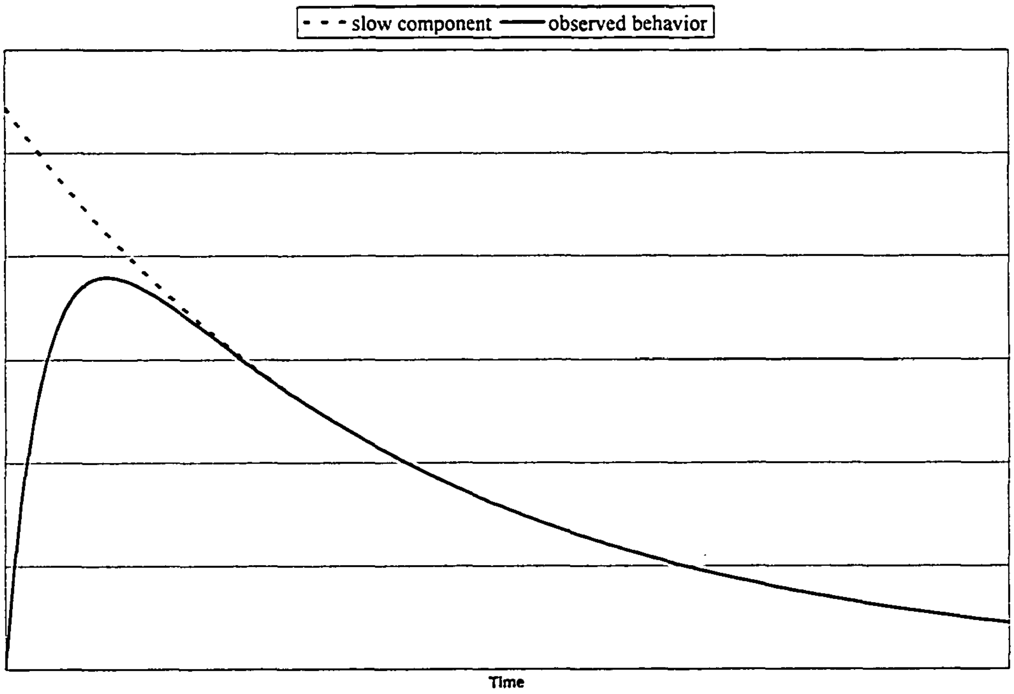

Figure 3.6 The slow component.

However, close to the initial time, when the *fast* component dominates, (18) fails to give an accurate portrait of the system's behavior. To find *that*, we need to change the time scale and 'stretch' the time to consider the system's behavior more closely around t = 0. We do this by transforming the independent time variable near the initial boundary so that:

(19) $$\zeta = t / \varepsilon$$

This transforms the original equations (15) to

(20) $$(d^2y)/d\zeta^2 + dy/d\zeta + \varepsilon y = 0$$

As $\varepsilon \rightarrow 0$, this reduces to

(21) $$(d^2y)/d\zeta^2 + dy/d\zeta = 0$$

The general solution to (21) is

(22) $$y^F = A + Be^{-\zeta}$$

This is the *fast* time-scale component of the solution for (15), and A and B are constants that must be determined. Since the solution must satisfy the boundary condition (16a), $y(0) = \alpha$, and since ζ will be 0 when t is 0 (by (19)), we have

$$y^F(\zeta = 0) = \alpha \tag{23}$$

hence $B = \alpha - A$, and the fast time-scale solution is given by

$$y^F = A + (\alpha - A)e^{-\zeta} \tag{24}$$

To determine the value of A, we apply what's known as a 'matching' condition. To see why this is needed, remember that what we're trying to do is to break the solution to (15) down into two components, a slow one and a fast one. Each is most appropriate in a particular region of the system's overall behavior, but each also describes the behavior of the system throughout the entire interval we're interested in (so the procedure is *not* to cut the time interval into two pieces and then join together the different descriptions for each piece). Since both the fast time scale description and the slow time scale description apply simultaneously, we need to make sure the combination of the two is appropriate by 'matching' them together. In this simple example, matching is a fairly straightforward task of determining the values of previously unknown terms by equating certain limits of the descriptions of the two processes. However, in general, the task is by no means trivial (see van Dyke 1975 for the classic account of the matching procedure). We equate the following limits

$$\lim y^S = \beta e \tag{25}$$

$$t \rightarrow 0$$

and

(26) $$\lim_{\zeta \rightarrow \infty} y^F = A$$

Equating these gives us:

(27) $$A = \beta e$$

What this 'matching' of limits does is to guarantee that the limit of the 'slow' component as its time scale goes to 0 (which is the domain where the *fast* component dominates) is the same as the limit of the 'fast' component as *its* time scale goes to ∞ (the domain where the *slow* component dominates). Once these solutions have been matched, we see that the full fast time scale solution is:

(28) $$y^F = \beta e + (\alpha - \beta e)e^{-\zeta}$$

Now, to find a solution uniformly valid over the entire domain, we combine these two solutions

(29) $$y^c = y^S + y^F$$

(30) $$y^c = \beta e^{1-t} + (\alpha - \beta e)e^{-t/\varepsilon} + O(\varepsilon)$$

Where the final term in the expansion represents an error term telling us that the error in this approximation is on the order of ε, which is the 'small' parameter in the original equation (15). Figure 3.7 shows the slow and fast behavior together:

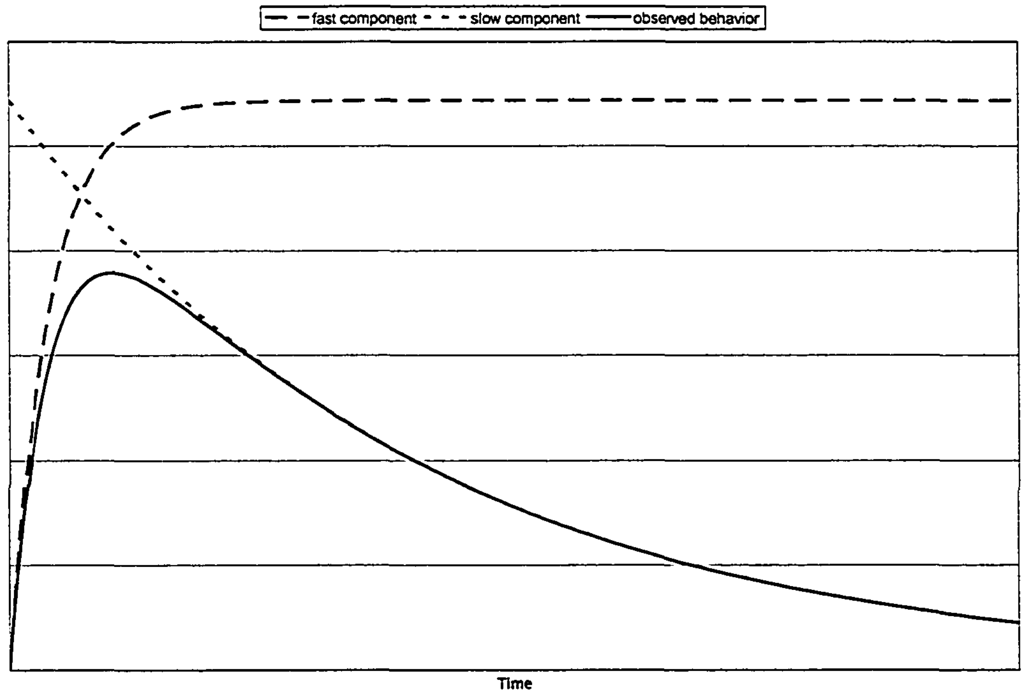

Figure 3.7 Slow and fast components.

These simple examples illustrate the basic ideas of multi-scale analysis: in each case, the system we're interested in describing is decomposed into component systems operating on different scales, and the behavior of the system as a whole is represented as a combination of the behavior of those components. In more complex, real-world cases, systems may be decomposed into a variety of different components operating on different temporal and/or spatial scales. My suggestion is that this sort of decomposition should be taken just as seriously as the sorts of decompositions into 'micro' or otherwise 'spatial' parts more often discussed in philosophy. Multi-scale analysis seems to be as much an instance of *structural explanation*, as described by Ernan McMullin (1978), among others, as the explanation of the hardness of a diamond in terms of the covalent bonding between the atoms in the crystal lattice constituting that diamond (to borrow

McMullin's example). McMullin argues that in an important class of explanations, otherwise unobservable structure is *postulated* to explain the behavior of a system of interest, and that the success of those explanations warrants the belief in the reality of that postulated structure. Explanations of this sort thus not only account for the observed phenomena, but also allow us to investigate entities and properties we cannot directly observe:

> The function of the explanation is not only to help one understand these [observed] features but also to discover the intrinsic structure of the entity. When the astronomer explains the changing shadows on the lunar surface by postulating that they are cast by mountain peaks, he not only explains the shadows, he also tells us something about the moon we might not otherwise know, i.e., that it (probably) has mountains. (McMullin 1978, p. 139)

And anyone doubting that multi-scale analysis does producing successful explanations need only look back to the graph illustrating the difference between the multi-scale analysis of the damped harmonic oscillator and the 'regular' perturbation: the multi-scale approach yielded results virtually indistinguishable from the exact solution[50].

Still, there may be a lingering worry that treating multi-scale analysis as revealing genuine parts of physical systems is the result of an unreasonably literal reading of the mathematics used to describe those systems: one might think that

[50] Perhaps I should emphasize again that in real-world applications, multi-scale techniques are employed precisely because exact solutions are not available. The beauty of the method is that even in these cases it can yield extremely accurate approximations.

multi-scale analysis is a useful 'trick' used to solve otherwise unmanageable equations, but one that is *merely* mathematical and not indicative of genuine structure in the way that explanations of the sort McMullin considers are. Defending the claim that component systems *are* real – though 'non-spatial' – parts will require a more thorough examination of the question of *component realism*: this is the worry about the reality of component forces, motions, and properties, such as the decomposition of the net force acting upon a body into a force due to gravitational attraction and the force due to electromagnetic repulsion. I'll examine this worry in detail in chapter four, but for now I'll simply note that physicists and mathematicians certainly *talk* as if the 'fast' and 'slow' components represent real parts of a system or process. For instance, in describing the need for using two distinct temporal scales to model the oscillations in temperature along a heated rod, D. A. Edwards writes:

> Since the two time variables t and τ represent the time scales for two different but interdependent physical processes, it seems reasonable that trying to solve the system using one time scale alone would fail. (Edwards 2000, p. 330)

And E. J. Hinch (1990) describes the applicability of multi-scaling to problems "characterized by having two physical processes, each with their own scales, and with the two processes acting simultaneously" (Hinch 1990: 116). Of course, this might *just* be talk, but I would suggest that those interested in understanding the ontological commitments of physics and other sciences should at least begin by taking the explanations and accounts found in those sciences seriously, and consider rejecting various claims as 'mere metaphor' (or whatever) only if serious obstacles arise.

Non-Spatial Parts and Micro-Based Properties

Let's return to Kim's supervenience argument, to see what difference recognizing these non-spatial parts from perturbation theory makes for our understanding of physical properties. Recalling Kim's distinction between 'levels' (which distinguish between the properties of relatively large and small entities) and 'orders' (which distinguish between different species of properties *of individuals*), Kim's claim was that since higher-order properties were realized by lower-order properties, those lower-order properties would exclude higher-order properties from playing any causal role in the world: any role a higher-order property might play would already be 'occupied' by its realizer. Kim suggested that while this conclusion had important consequences for our understanding of psychological properties, the argument didn't 'generalize' to other sorts of properties, since it was supposed to be widely admitted that other sorts of properties could be *identified* with 'micro-based' properties.

In the case of fast and slow components, it isn't clear how the properties characterizing those components – the 'spatio-temporal structural properties' described by the equations found in multi-scale analysis – could ever be identified with *any* micro-based properties, even once we extend Kim's own account of entities and properties to include temporally structured ones. Of course, Kim himself doesn't assume that *all* explanatory parts are immediately 'micro-structural' ones: like most philosophers, he assumes that higher-level entities and properties will be analyzed into somewhat lower-level ones, and then these in turn will be analyzed into still lower-level ones and so on, until *eventually* we reach the micro-level (i.e., psychological properties are composed of biological properties, which are composed of chemical properties, an so on). But the problem with the decomposition into fast and slow components is that it represents an entirely different sort of decomposition, and it's hard to see how it could ever be 'cashed out' at the micro-level. That doesn't mean that systems with fast and slow components can't *also* be decomposed into micro-level components. However, it does mean that the micro-level components of the system – whatever they are – are not themselves components of the fast and slow components.

Suppose we explain the properties of an organism by decomposing it into cells and giving an explanation in terms of *their* properties. This sort of decomposition is perfectly compatible with an explanation of the properties of the organism in terms of its decomposition into *molecules* because the cells found in the first sort of decomposition themselves can be decomposed into the molecules found in the second sort of decomposition. Thus pointing out that sciences like biology appeal to cellular decompositions doesn't itself raise any particular problem for Kim's view: those cellular decompositions can themselves (Kim assumes) be further decomposed into the micro-level. However, in the case of perturbation theory, the decomposition into fast and slow components is of an entirely different sort: a system that is decomposable into fast and slow components might also be decomposable into micro-level components, but the fast and slow components themselves (unlike the cells in the first example) won't be decomposable into those same micro-level components. Rather than a hierarchy of levels rising out of the world of elementary particles, the examples from perturbation theory suggest a more ramified structure where macro-level entities and processes can be decomposed in different ways.

This point becomes clearer if we consider the relationship between the slow and fast components and whatever micro-level properties or processes *realize* them. Kim argued that physical properties (and, by extension, processes) could be *identified* with whatever micro-level properties were there realizers. However, one reason for thinking that such identifications cannot be found is that while a given system presumably has only one micro-decomposition or 'micro-based property' in Kim's sense, the fast and slow components of a system are themselves distinct components, and hence *both* cannot be identified with the same micro-structure. Again, in cases involving purely spatial decomposition this isn't a problem. For instance, two distinct cells in a cellular decomposition need not both be identified with the same micro-level structure. Instead, one cell is identified with the micro-level structure of one part of an organism, while the other is identified with the structure of another part. But since non-spatial parts overlap in space and time, this kind of partitioning at the micro-level is not possible. As I said in the introduction, there are not 'slow' and 'fast' particles. If anything, it is the *same* particles that exhibit both 'slow' and 'fast' behavior. Note that none of the examples I've discussed so far explicitly describe micro-level behavior, so it would be improper to interpret slow and fast components as descriptions of micro-level activity. In fact, most of the examples involving this sort of decomposition are found in 'macro' sciences such as fluid mechanics. However, the point is that since slow and fast components of macro-level systems are themselves distinct, they cannot both be identified with the micro-structure of those systems, whatever that micro-structure is.

So examples from perturbation theory apparently show that Kim's claims about micro-based properties in the physical sciences are wrong: we can't deny that physical systems *have* micro decompositions, but they can have other sorts of decompositions as well that aren't easily accommodated by the sort of hierarchical scheme Kim imagines.

Non-Spatial Parts and Reduction

Non-spatial parts raise a problem for Kim's view of properties because we can't expect to identify two distinct component systems with one micro-based property, and yet since the component systems overlap in space and time, they can't each be identified with separate micro-structures of the system they compose. However, an interesting worry about this claim is the thought that just as we might take macro-level descriptions to be coarse-grained characterizations of the 'true' micro-level reality, we might analogously take *slow*-scale descriptions to be coarse-grained approximations of the 'true' fast-scale reality. The suggestion would be that while it is useful in practice to distinguish slow and fast scale processes, *in principle* slow scale processes 'reduce' to fast-scale ones. If that were true, then perhaps there would be no worry about trying to identify distinct non-spatial components with the same micro-based realizer: if the slow-scale component reduced to the fast-scale component, there would only be one true process to be micro-based, and Kim's claims about property reduction would be vindicated.

Of course, Kim himself doesn't make any claim of this sort, since he doesn't discuss temporally extended processes in the first place. But since the question of slow/fast reduction is an interesting one independent of Kim's opinions, I'll now show that such reductions fail.

Prima facie, the thought that such 'temporal' reductions should be available seems as justified as the thought that more familiar 'spatial' reductions ought to be available. As I said, one reason for thinking that the *macro* reduces to the micro is the thought that macro-level descriptions are simply abstractions from a 'coarse-graining' of the excessive detail found at the micro-level. This was part of the worry about reductionism discussed in chapter two. If macro-level descriptions do reduce to micro-level ones, then it seems that macro-level descriptions are in principle unnecessary and only of use because they conveniently gloss over irrelevant details: macro-level descriptions don't themselves capture any features of a system not already found in the micro-level descriptions. Analogously, we might think that relatively *fast* scale descriptions represent a 'truer' picture of reality than slow scale ones and that we appeal to slow scale descriptions only because managing the excessive detail provided by the fast-scale description is impossible for us. Just as it is clear that *we* could never deal with the world entirely at the micro-level, it is equally clear that *we* could never deal with the world entirely on some particularly fast scale (the *femto* scale, for instance, where changes occur on the order of 10^{-15} seconds). But what we want to know is not whether descriptions on relatively slow scales are *useful* (clearly they are), but instead whether they reflect real features of the world. In the case of geophysics, for instance, we might think that our choosing to describe continental drift on a particularly slow scale on the order of tens of thousands of years is simply a matter of convenience: describing it on a faster scale, such as a temporal scale on the order of seconds, say, would result in much more information than necessary. But *prima facie*, there is no reason to think that there is anything particularly natural about the slower time scale: it was arbitrarily chosen by us to make the information more manageable. By analogy with the argument about *spatial* scales, we might think that the proper scale for describing what *really* goes on in geophysics is the 'fast' scale – whatever scale represents the fastest changes in state of the particles constituting the continents and whatever else geophysics studies.

This sort of reduction fails, however, at least in the sense of reduction most appropriate for these descriptions (whether or not it would succeed in some other sense of 'reduction' is always an open question): fast and slow components are genuinely distinct parts of a process. This failure of reduction is interesting in itself, and while it would pure speculation to suggest that Kim should expect to find this sort of slow/fast reductionism alongside the macro/micro reduction he clearly *does* expect, it is difficult to see what other sort of reduction the non-spatial parts could be involved in.

As I mentioned in chapter two, there are various distinct senses of 'reduction' used in philosophy. The most appropriate sense of reduction for these cases is that due to due to Nickles 1973 (this view is also developed and defended by Batterman 1995 and Rueger 2004). Rather than defining reduction in terms of the logical derivability of one set of laws from another via appropriate 'bridge laws', as in the standard Nagelian account of reduction, Nickles-reduction, as I'll call it, relates different descriptions in the limit of some parameter. On this account of reduction, one property or description of properties reduces to another just in case solutions for the reducing description 'go over to' solutions for the reduced description as some parameter tends to a limiting value. We can distinguish these forms of reduction as follows:

(i) One description D *Nagel-reduces* to another D* just in case D can be derived from D* together with some appropriate 'bridge principles'.

(ii) One description D *Nickles-reduces* to another D* just in case D → D* in the limit of some parameter.

In reductions of the first sort, the reducing theory is typically more general and sometimes more 'fundamental' than the reduced theory. Reduction in this sense involves some suggestion of incorporation: finding a Nagel reduction between two theories shows that one theory (the reduced one) is really a special case of the other (the reducing theory): from the general principles of the reducing theory, together with specific bridge principles, we can derive the more specific principles of the reduced theory.

In reductions of the second sort, the reducing theory is typically *less* general than the reduced theory. One of Nickles' examples is the relationship between classical mechanics and the special theory of relativity. The special theory of relativity is more general than classical mechanics, since it omits presumptions about the velocity of light: that is, we can derive classical mechanics from the special theory of relativity if we let a parameter in the special theory of relativity – namely, the velocity of light – go to a particular limit (in this case, infinity). For instance, the formulas for momentum in classical mechanics and the special theory of relativity respectively are:

Classical: $p = m_0 v$ Relativistic: $p = m_0 v / \sqrt{(1 - v^2/c^2)}$

Where p is momentum, m_0 is a body's rest mass, v its velocity, and c the velocity of light. The key feature of the special theory of relativity is that the speed of light is included as a parameter in its equations. It is the fact that the speed of light is finite and constant in all reference frames that gives rise to the sorts of effects that distinguish our world as a relativistic one rather than a classical one. However, if we examine these equations in the limit as this parameter tends to infinity, they smoothly transform to the classical equations of motion. So momentum as described by the special theory of relativity Nickles-reduces to momentum as described by classical mechanics in the limit as $c \rightarrow \infty$. (That is, as $c \rightarrow \infty$, the ratio $v^2/c^2 \rightarrow 0$ for any finite v; hence the denominator in the relativistic equation $\rightarrow 1$, yielding the classical equation; this same relationship also explain the applicability of classical mechanics at relativity low velocities, since as $v \rightarrow 0$, the ratio $v^2/c^2 \rightarrow 0$ for finite c, again resulting in the denominator in the relativistic equations tending to unity.)

Note that the direction of reduction is the reverse of what we might expect: despite the fact that the special theory of relativity is the 'replacing' theory, it is relativity that Nickles-reduces to classical mechanics, rather than the other way around. While Nickles defends the intuitive appeal of this account of reduction (and in particular the idea that it is the more general theory that reduces to the less general one), we could simply reject this terminology – in spite of the fact that "it is the way physicists and mathematicians, in contrast to most philosophers, usually talk" (Nickles 1973: 182) – and maintain the criterion for reduction while reversing the order in the relation. That is, we could just as well define L-reduction ('reduction in the limit') as:

> A description D L-reduces to another D* just in case D* → D in the uniform limit of some parameter *e*.[51]

Now we can apply this definition of reduction to the example from perturbation theory. As I'll discuss in chapter four, in similar perturbation cases involving macro and micro spatial scales, the appropriate limit is the limit of the perturbation parameter ε tending to zero: since the parameter in those cases is defined as the ratio between the micro and the macro scales, this corresponds to letting the micro scale become increasingly fine-grained, becoming continuous in

[51] This seems to be the sense Rueger has in mind: "A theory ρ (the *reduced* theory) reduces to a theory β (the *base* theory) just in case there is a uniform limit in a suitable parameter of β in which the solutions of β go over into the solutions of ρ." (Rueger 2004, p. 6, slightly modified).

the limit. This process is known as 'homogenization' and represents the 'smoothing' of micro-level details as one moves to the macro-level view[52]. If we investigate the corresponding limit in the case of slow/fast components, we see that the slow component system does not tend toward the fast component system as $\varepsilon \rightarrow 0$ (note that, as I'll discuss in the final chapter, L-reduction also fails in the macro/micro cases). From our first example (15), the limit of the slow component as $\varepsilon \rightarrow 0$ is:

(31) $$y^S \rightarrow \beta\, e^{1-t} \text{ [(18) above]}$$

The fast component, on the other hand, is given by

(32) $$y^F = \beta e + (\alpha - \beta e)e^{-\zeta} \text{ [(28) above]}$$

which, since $\zeta = t / \varepsilon$, can also be written as

(33) $$y^F = \beta e + (\alpha - \beta e)e^{-t/\varepsilon}$$

In the limit as $\varepsilon \rightarrow 0$, (33) becomes

(34) $$y^F = \beta e$$

Which is distinct from (31). So the slow component and the fast component do not converge as the perturbation parameter ε goes to zero, and hence there can be no question of 'reducing' the slow system to the fast one. The slow system represents a distinctive view of the behavior of the system that must be dealt with 'in its own terms'.

[52] See Hornung 1997.

Conclusion

In this chapter, I've presented some examples drawn from multi-scale analysis that demonstrate how solving the equations describing dynamical systems can involve 'decomposing' those systems into parts operating on different temporal scales. The parts are not ordinary 'spatial' parts, and I've argued that despite being parts of physical systems, they don't fit into Kim's 'hierarchical view' of physics, according to which all physical properties are built up from properties of fundamental particles. With these examples from physics in mind, I'll now turn to some more metaphysical questions about the idea of such 'non-spatial' parts.

IV

Non-Spatial Parthood

Introduction

In this chapter, I'll defend the idea suggested in the previous chapter: that examples from multi-scale analysis reveal genuine, though 'non-spatial', *parts* of the systems involved. I'll begin by discussing the concept of non-spatial parts, in particular with respect to the understanding of parthood found in the standard conception of *mereology*, the formal theory of parts and wholes. I'll then turn to what I take to be the most significant objection to this interpretation of multi-scale analysis: that while multi-scale analysis is clearly useful, it is better characterized as a sort of mathematical trick rather than a representation of any real structure in the world. I'll argue that we can make sense of the idea of non-spatial parts, and that while there might be no way to convince the ardent anti-realist that such things exist, there are no insurmountable obstacles for the realist in holding that they do exist.

Spatial and Non-Spatial Parts

Recall Meirav's account of physical entities (see Meirav 2000, and the above discussion in chapter one): Meirav claimed that physical entities were distinguished from non-physical entities by the fact that the parts of physical entities corresponded to the parts of space occupied by those entities in a way that the parts of non-physical entities did not. While we are rejecting this as a claim about physical entities, we can still make use of the idea of relating an entity to a particular region of space to define spatial parthood. For the moment, I'll consider spatial (and temporal) parthood to be species of a generalized parthood relation; I'll say more about how *that* relation might be understood below.

> *Spatial Parthood*: x is a *spatial part* of y just in case x is part of y and the region of space occupied by x is contained within the region occupied by y.

Since the examples from the previous chapter involved temporally extended *processes*, and since we want to distinguish the claim that a process is decomposable into 'simultaneous' component processes from the claim that a thing may be decomposable into discrete temporal 'stages', we could define *temporal* parthood as well. Rather than talking of 'regions' of time, I'll follow Paul Needham (1999) in characterizing processes as occupying particular *intervals* of time:

> *Temporal Parthood*: x is a *temporal part* of y just in case x is part of y and the interval of time occupied by x is contained within the interval occupied by y.

I'm taking the idea of 'containment' of regions of space or intervals of time as primitive here: whatever the correct mathematical details are, the intuitive idea is that one region of space is contained within another just in case the boundaries of the first extend no further than the boundaries of the second, and *some* boundary of the second extends beyond the boundaries of the first. Similarly for time: one interval is contained in another just in case the first begins no earlier than the second and ends no later, and the first either begins later than or ends earlier than the second (or does both).

Needham's account of parts of processes comes as a defense of a 'macroscopic' ontology for chemistry similar in spirit to that suggested for physics here[53]. The processes he considers are primarily thermodynamic ones, such as the process of heating a liquid, or the process of a body of gas expanding. Such processes are characterized as involving 'bodies' or objects; in particular, they involve changes in the properties or states of objects over a particular interval of time. For example, the boiling of a quantity of water by heating an electric coil is a process involving certain bodies (the quantity of water and the coil) and their changes in time. Such processes could have spatial parts – for instance, a spatial part of the boiling-by-heating process would be the heating of the electric coil – and they can have temporal parts, such as the process of raising the temperature of the water from 98° to 99° C. But they can also occur as composite processes: for instance, in the heating of a body of gas, the same body of gas might be involved both in a process of heating, and in a process of expanding, and together these two component processes constitute a process characterizing the overall behavior of the gas. Since the process of heating and expansion both primarily involve the same body – the gas – and both occur over the same time, they are not spatial nor temporal parts of that overall process:

[53] Needham has developed accounts both for macroscopic bodies and processes in chemistry (see Needham 1996 for bodies, and Needham 1999 for processes).

instead, Needham refers to such processes as 'component' parts, but calling them 'non-spatial parts' seems equally appropriate.

The distinguishing feature of component processes on Needham's account is that they occur simultaneously and can overlap in space. He argues that this view of parthood can be traced back to an account of parthood between *bodies* defended by the Stoics in their account of chemical mixtures (see Needham 1996; Sorabji 1988). Apparently, the Stoics maintained a view of the composition of mixtures that explicitly denied that all parts of a body need be spatial parts. The Stoic view is best understood in dual contrast with the view defended by 'atomist' philosophers such as Democritus and Leucippes, and with the view defended by Aristotle. According to the atomists, all mixtures simply involve the 'juxtaposition' of elementary particles, as for example, in the mixing of rice and beans. Aristotle rejected this view, however, insisting instead that in a genuine mixture – such as the dissolution of sugar in water, to use a common (ancient) example – the components mixed must in some sense cease to exist, while some new substance – their mixture – comes into being. The components mixed, however, are not entirely destroyed: they maintain *potential* being, since (often) those original components can be recovered (Richard Sorabji's example – which he evidently tested himself – is that of a mixture of wine and water which can be 're-separated' using a sponge soaked in oil; a more mundane example would be the separation of sugar and water through evaporation). The Stoics agreed with Aristotle in rejecting the 'mere juxtaposition' suggested by the atomists, yet apparently maintained that the separability of components could only be explained if they were *actually* present in the mixture, contrary to Aristotle's account. Hence on the Stoic account, the components remained and actually occupied the same region of space as the body they composed. In the terminology suggested here, mixed components are 'non-spatial' parts of the mixture they form.

Keeping in mind that non-spatial parts are not supposed to be just temporal parts (as defined above) we could define non-spatial parthood simply as the denial of both spatial parthood and temporal parthood: non-spatial parts are parts that are neither spatial parts nor temporal parts. However, one possibly undesirable consequence of this would be that non-spatial parts might 'out-strip' the things they compose and actually occupy a *larger* region of space. Whether or not this is possible depends partly on how we understand 'generic' parthood, but we can avoid this sort of worry in advance of an account of generic parthood by defining non-spatial parts not simply as parts that are neither spatial nor temporal, but more specifically as parts that occupy the same region of space and time as the things they compose:

> *Non-Spatial Parthood*: x is a *non-spatial part* of y just in case x is part of y and the region of space and interval of time occupied by x are *the same* as the region of space and interval of time occupied by y.

As I mentioned earlier, the term 'overlapping parts' might be intuitively the most appropriate for the type of parts we are considering, but 'overlap' already has a use in the theory of parthood.

What are Parts?

With this definition of non-spatial parthood in mind, what can we say about 'generic' parthood? Unfortunately, while there is considerable discussion in the literature on the concept of *composition*, there is relatively little discussion of the idea of parthood itself. For example, Peter van Inwagen's paper "When are Objects Parts?", quickly turns from the intriguing title-question to the related but distinct question "when do objects form 'wholes'?". He writes:

> I shall approach the concept of parthood in a somewhat indirect way. There is a mereological concept that I have found it easier to think fruitfully about than I have parthood...I call this notion *composition*. (van Inwagen 1987: 22)

While there is a clear connection between parthood and composition – one thing is part of another just in case the second is composed of a group of things that includes the first – putting the question in terms of composition suggests that we *begin* with distinct entities, compose some 'new' entity from them, and then conclude that those original entities are the parts of the composite entity. However, this sort of approach seems particularly ill-suited to the examples from the previous chapter, where the issue is not over the sorts of things 'non-spatial' parts might compose, but instead over whether or not such parts exist in the first place.

Frequently, the concept of parthood is simply adopted as primitive, the assumption being that the concept is so basic it requires no illumination. In another discussion of parthood, van Inwagen does just this, "Let us say that x is a part of y just in case that x is either a part of y in the ordinary sense of the English word 'part' or is identical with y" (van Inwagen 1994: 207). Rosen and Dorr's (2002) defense of 'compositional nihilism' – the view that there are *only* parts, and no genuine composite entities – does likewise, suggesting only that the parthood relation be understood "in the usual way" (Rosen and Dorr 2002: 153). Similarly, Peter Simons, in introducing the concept 'parts' in his book *Parts*, offers only a list of familiar examples of parthood (the trunk is part of a tree, the roof is part of a house, etc.) considering the idea of parthood itself "the most basic and most intuitive mereological concept" (Simons 1987: 9). While Simons' book contains much valuable discussion of composition and other related notions, little else is said about parthood per se. And Needham, too, despite drawing an explicit contrast between the concept of 'spatial' parthood and a concept of 'non-spatial' parthood – though he doesn't use that term – develops his entire account without saying anything about what makes for parthood in general[54].

[54] Other sources are not much more helpful. For instance, in his well known discussion of parts and wholes, Edmund Husserl begins by explaining his understanding of parthood as follows:

> We interpret the word 'part' in the *widest* sense: we may call anything a 'part' that can be distinguished 'in' an object, or, objectively phrased, that is 'present' in it. Everything is a part that is an object's real possession, not only in the sense of being a real thing, but also in the sense of being something really in something, that truly helps to make it up. (Husserl 1970: 437).

While this may provide some illumination of the idea of parthood, talk of one object being 'present in' another sounds more like a restatement of parthood than an explication of parthood. Of course, this is only an introductory characterization of parthood by Husserl, though like much contemporary work in mereology, the ensuing discussion tends to focus on the idea of 'wholeness' and composition rather than parthood.

In his original presentation of mereology, Lesniewski provides few details about the parthood relation itself, choosing instead to adopt it as a primitive relation and supposing "that this term [i.e., 'is a part of'] will not cause misunderstandings, considering that its intuitive character acquires considerable clarity in the light of Axioms I and II" (Lesniewski 1992: 131). Unfortunately, Axioms I and II in Lesniewski's system simply express the irreflexivity and transitivity of the relation of being a (proper) part. The axioms may both seem formally compelling, but they do little to reveal the distinctive character of parthood. Standard presentations of mereology are now based on three axioms[55]:

Transitivity: If x is part of y and y is part of z, then x is part of z.

Unrestricted Composition: Any group of things has a fusion.

Uniqueness of Composition: Any group of things has only one fusion.

Where a *fusion* – the mereological 'sum' of a group of individuals – can be defined as:

Fusion: x is the *fusion* of the Fs just in case each of the Fs is part of x and no part of x is discrete from all of the Fs.

In addition to parthood and fusion, the other two basic concepts of mereology are *overlap* and *discreteness*. The basic concepts can be defined in terms of one another, so that after adopting one as primitive, the remaining concepts can be treated as a derivative. For example, adopting *parthood* as our primitive relation, we can define:

[55] For ease of exposition, I follow Lewis' (1991) presentation of mereology (see Lewis 1991: 72-74), rather than Lesniewski's original formulation, which involve some unusual assumptions about predication.

Overlap: x and y *overlap* just in case there is some z such that z is part of x and z is part of y.

Discreteness: x and y are *discrete* just in case they do not overlap.

The most serious problem for the idea of non-spatial parts is that concepts like 'overlap' and 'discreteness' suggest that the most obvious candidate for a generic account of parthood is simply *spatial* parthood:

Parthood: x is part of y just in case the region of space occupied by x is contained within the region of space occupied by y.

Obviously that account won't do here, since non-spatial parthood would become contradictory. But perhaps it's simply wrong that the 'spatial' account *is* a good account of generic parthood. In defending mereology, David Lewis (1991: 75) argues that it is wrong to assume that the part-whole relation is inherently spatial (or spatio-temporal). For instance, says Lewis, the part-whole relation can hold between things not plausibly thought of as being situated in space (or time), as when we say that trigonometry is part of mathematics. Even for things that *are* situated in space and time – and this is the sort of thing we are interested in here – Lewis argues that it is wrong to assume that parts must be spatial parts, insisting that "[mereology] is silent about whether spatiotemporal things may have parts that occupy no less of a region than the whole does" (Lewis 1991: 76). For example:

> Suppose it turned out that the three quarks of a proton are exactly superimposed, each one just where the others are and just where the proton is. (And suppose the three quarks each last just as long as the

> proton.) Still the quarks are parts of the proton, but the proton is not part of the quarks and the quarks are not part of each other. (Lewis 1991: 75)

Lewis' example demonstrates the difficulty involved in thinking of some of the basic concepts of mereology 'non-spatially'. For instance, talk of 'exactly superimposed' quarks *sounds* a lot like talk of 'overlapping' quarks, and yet they cannot *exactly* overlap, in the mereological sense, and still be distinct entities: overlapping in the mereological sense is a matter of having the same parts, so exactly overlapping in the mereological sense would mean having exactly the same parts; but Uniqueness of Composition guarantees that distinct things cannot have exactly the same parts, hence 'two' things exactly overlapping in the mereological sense would in fact be one and the same thing. So to make sense of Lewis' suggestion, we need to be able to distinguish between spatial overlap and 'mereological' overlap.

The concept of *discreteness*, on the other hand, may be more amenable to a non-spatial understanding, and hence may help to clarify a non-spatial understanding of parthood and overlap. In their classic presentation of mereology as the 'calculus of individuals', Henry Leonard and Nelson Goodman (1940) suggest the possibility of non-spatial senses of their primitive relation, 'discreteness':

> In our interpretation...parts and common parts need not necessarily be spatial parts. Thus in our applications of the calculus to philosophic problems, two concrete entities, to be taken as discrete, have not only to be

spatially discrete, but also temporally discrete, discrete in color, etc. etc. (Leonard and Goodman 1940: 47)

So spatial discreteness – the lack of spatial overlap – is not a sufficient condition for genuine discreteness. Instead, truly discrete entities must be discrete in other ways as well. While it seems that Leonard and Goodman thought that spatial discreteness was a *necessary* condition for genuine discreteness, we needn't accept this part of their claim: we might instead think that spatial discreteness, temporal discreteness, and discreteness in properties are each *sufficient* conditions for discreteness. Two things can be discrete by being spatially discrete – by occupying different regions of space – *or* by being temporally discrete – by occurring over different intervals of time – *or* by being distinct in some other of their properties – for instance in their *spin*, as in the *up* and *down* quarks composing a proton. In terms of mereological *overlap* those component quarks overlap the proton but not each other because their *properties* – their *mass*, for instance – are not discrete from those of the proton, though they are discrete from each other. For example, presumably, changing the mass of one of those quarks would change the mass of the proton without affecting the masses of the other component quarks. Thus discreteness (and hence parthood) between *things* is intimately connected with discreteness (and hence parthood) between *properties*.

Note that I say 'presumably' because I lack the expertise in physics to say definitively how the mass of component quarks relates to the particles they compose. However, it seems that for there to be a genuine relation of *parthood* between the quark and the proton, there must be *some* sense in which – to borrow

David Lewis' phrase – the properties of the quarks and the proton represent the same 'portion of Reality', carved up in different ways[56]. Hence if a quark *is* part of a proton, there must be some dependence between the properties of the one and those of the other.

The idea that parthood involves a dependence between properties – rather than spatial containment – also emerges in what for our purposes might seem to be a surprising place: the discussion of the parthood relation required for instances of *micro*-reduction. One of the main concerns in genuine cases of micro-reduction – that is, cases involving truly micro-level entities rather than simply the sorts of 'compositional' reductions discussed in chapter one – is over the relationship between the types of property found at the micro-level and those found at higher-levels. For instance, George G. Brittan, Jr. (1970) quotes the physicist Werner Heisenberg commenting on the properties of micro-parts as follows:

> It is impossible to explain...qualities of matter except by tracing these back to the behavior of entities which themselves no longer possess these qualities. If atoms are really to explain the origin of color and smell of visible material bodies, then they cannot possess properties like color and smell...Atomic theory consistently denies the atom any such perceptible properties. (Werner Heisenberg, taken from Brittan Jr. 1970: 455).

[56] The phrase comes from a discussion of the worry over the 'ontological innocence' of mereology. Lewis maintains that despite the fact that mereology holds that whenever there are some things, there is a *fusion* of those things, this 'fusion' is not an 'addition to being' and hence that a belief in mereology does not involve a commitment to the belief that there is anything *more* to the world than the parts composing that fusion. See Lewis 1991: 81-87 for discussion.

The issue here is over whether genuine micro-reductions can involve what Todd Girill (1976a, b) calls 'Empedoclean' explanations, where it is the relative abundance of a given property among the parts that explains the presence of that property in a whole (such as explaining the redness of blood in terms of the redness of its predominant cells), or whether they must – as Heisenberg suggests – be 'Democritean' explanations, where the properties of explanatory parts are distinct from whatever property they are to explain. However, Brittan Jr. argues that micro-reductions of the sort Heisenberg envisages, where 'part' and 'whole' have *no* properties in common, make it impossible to understand the sense in which the entities involved in the micro-reduction are *parts* of a given whole. Instead, he argues, there must be *some* common properties. In particular, Brittan Jr singles out relations between *additive* properties as indicative of parthood between the entities bearing those properties. The idea of additivity applies most naturally to quantitative properties such as *mass*, so that a property is additive if the quantity possessed by a given 'whole' is the arithmetic sum of the quantities possessed by the parts of that whole. Brittan Jr. claims that this sort of relation is central to understanding parthood:

> the additivity of their properties gives us some grip on what it means to call, e.g., the elementary corpuscles "parts" of wholes. Objects are "sums" of particles in just this sense, that their length, mass, etc., is the sum of the lengths, masses, etc., of the particles that compose them. (Brittan Jr. 1970: 457)

But we needn't assume that component properties must be *additive* in order to think that it is the sharing of properties that gives us some 'grip' on parthood. Not only is it difficult to see what the concept of additivity would *mean* for many properties, but there are well known difficulties for this suggestion even for properties – such as mass – where it is relatively clear what additivity *would* mean. In the case of mass, the relativistic equivalence of mass and energy leads to a disparity between the mass of a composite entity and the arithmetic sum of the masses of its parts known as the 'mass defect'. For instance, the mass of a proton – 0.938 GeV/c^2 – is considerably greater than the sum of the masses of the three quarks composing it, as these have a total mass of only 0.02 GeV/c^2, the remaining mass of the composite proton being accounted for in terms of the energies binding those quarks together. Hence expecting the parthood relation between three quarks and a proton to be accounted for in terms of the sum of the component masses would be a mistake.

However, we can retain the idea that dependence relations between properties 'give us a grip' on parthood without insisting that the relation be describable in terms of ordinary addition. This sort of view of parthood is defended by Girill (1976a, b). On Girill's account, one entity *x* is a micro-part of another entity *y* suitable for micro-explanations only if:

1. *x* is spatio-temporally contained in *y*
2. for an other part of *y*, either
    (a) *x* is spatio-temporally discrete from it, *or*

(b) x and it are assigned to different 'levels' such that parts from only one level are used in a given micro-reductive explanation

3. x and y share at least one quantitative, additive property such that

   (a) the amount of this property attributed to x is less than the amount attributed to y, *and*

   (b) the former amount arithmetically contributes to the latter

(Girill 1976a: 77)

Obviously, Girill's interest in formulating parthood for micro-reductionism leads him to adopt (1)[57], and this requirement of spatial containment is exactly the sort of requirement we want to reject for non-spatial parts. Principle (2) is also of questionable use for our purposes, since in the case of the type of decomposition into slow and fast components discussed in the previous chapter, we had two distinct parts that were each used in the same explanation of the overall behavior of the system, and yet were not spatio-temporally discrete. Not being spatio-temporally discrete, such parts cannot satisfy (2a) in Girill's definition. However, they cannot satisfy (2b) either, since there is no plausible sense in which they are assigned to different 'levels' such that parts from only one level are used in a given explanation. We might argue that these components *are* assigned to different levels, since – as I discussed in the previous chapter – talk of 'levels' in philosophy can be associated with talk of 'scales' in science, and clearly the slow and fast components *are* associated with different scales. However, there's no getting around the fact that these components are often used in the *same*

[57] Notice again – as discussed in chapter one – the tendency to characterize 'reductions' involving any sort of spatial part as specifically *micro*-reductions.

explanation, contrary to Girill's second requirement. It's Girill's third requirement that offers the sort of softened version of additivity that seems appropriate for characterizing parthood. Girill insists only that micro-parts and wholes share at least one quantitative property such that the quantity possessed by the part 'arithmetically contributes' to the quantity possessed by the whole, rather than insisting that that arithmetic contribution is specifically one characterized by addition[58]. Aside from giving us some intuitive picture of what parthood involves, this condition is meant to eliminate 'spurious' cases of parthood where an object happens to occupy part of the same region of space as a given 'whole', while intuitively not being among the genuine *parts* of that whole. For instance, this rules out considering various sorts of 'impurities' as genuine parts, such as pollen particles in a sample of gas, or chemical impurities in a crystal (Girill 1976a: 73).

With the idea that it is the dependence relation between properties that makes for 'generic' parthood, we can define:

> *Parthood*: x is part of y just in case there is some property P characterizing both x and y such that the value attributed to x functionally contributes to the value attributed to y.

[58] I take this to be the central feature of Girill's account. I have been unable to find an explanation of why Girill continues to refer to this shared property as an 'additive' one, while at the same time denying that the relation between the quantity attributed to the whole need be the sum of those quantities attributed to the parts. In discussion of these principles in the texts, Girill sometimes omits any reference to additivity, though without any remark on it significance in his explicit definition of micro-reductive parthood.

And the definition of non-spatial parts remains the same (but I'll repeat it for convenience):

Non-Spatial Parthood: x is a *non-spatial part* of y just in case x is part of y and the region of space and interval of time occupied by x are *the same* as the region of space and interval of time occupied by y.

Note two things. First, in the definition of 'generic' parthood, I've generalized Girill's suggestion slightly, removing the suggestion that the 'amount' of the shared property possessed by the part must be less than that possessed by the whole. This condition might seem appropriate if the property we have in mind is *mass*, but there are other sorts of properties that plausibly make for parthood without satisfying this condition. For instance, it's conceivable that a particle – a chargeless, massless neutrino, for example, might have parts that are themselves *charged* though massless. Since the neutrino itself has a neutral charge, we could imagine it having parts that are themselves oppositely charged: in that case, only one of those parts would have 'less' charge than that possessed by the whole. Yet it still seems plausible that it is the dependence between the charge (or lack thereof) of the neutrino and those of the constituents that underwrites the parthood relation between them: as Brittan Jr. said, it is the dependence between these properties that 'gives us some grip' on what parthood *means* in such cases.

Second, note that – perhaps despite appearances – defining generic parthood in terms of dependence relations between properties does not trivialize the claim of compositionality made in chapter one: *that* claim was that *all* properties of 'wholes' are determined by those of their parts, whereas the concept of parthood suggested here only depends on a relationship between *some* property shared by 'part' and 'whole'. This leaves open the possibility (which compositionality denies) that some properties of wholes are independent of those of their parts.

What about the sorts of examples discussed in chapter three? The relevant properties there are the 'spatio-temporal' structural properties describing the behavior of the system and its sub-systems through time. Clearly in these cases the properties of the components 'arithmetically contribute' to those of the composite system, since the whole point of the analysis is to calculate the behavior of the composite system based on the combined behavior of systems operating on different scales. Note that, as I pointed out in chapter three, while in the most perspicuous cases this arithmetic contribution made by the component systems could be characterized as 'additive', the relation between the properties characterizing component systems and the properties of the system they compose can also be 'multiplicative'; hence the more general idea of 'arithmetic' contribution seems appropriate[59].

---

[59] Also note that since multi-scale analysis yields asymptotic expansions, the properties of composites are not *equal* to the sum of those of their parts, even when composition is additive. Hence I would emphasize the idea of *contribution*: the properties of fast and slow components (for example) contribute to those of the whole, while not exhausting the properties of the whole. Finally, it is true that there is a reduction in order between the composite system and its components – that's an important feature of multi-scale analysis since it leads to a simplification of the equations describing the system – and we might take this difference to indicate that the composite and component systems have properties of different *types*; but the suggestion in the

There is, however, a potential fly in the ointment of these claims about non-spatial parts. Even if we can reconcile the idea of non-spatial parts with the standard theory of parthood and give some intuitive account of what makes non-spatial parts *parts*, one still might argue that the sort of 'decomposition' associated with multi-scale analysis is merely mathematical and doesn't reflect any real structure in the system or properties it describes. Consequently, one might argue that the only *real* parts of a physical system are its ordinary 'spatial' parts, and that the only real *properties* characterizing a system are the sorts of 'micro-based' properties Kim talked about in his discussion of reductionism. If multi-scale analysis were viewed as merely 'instrumental' in this way, then it wouldn't have any consequences for the arguments about physicalism after all. I'll spend the rest of this chapter discussing this objection.

Component Realism

The objection to a realistic reading of the sort of decomposition found in multi-scale analysis is part of a broader opposition to claims of *component realism*. It is a familiar fact that forces, motions, and other sorts of behavior can be decomposed into components in a variety of ways. The general problem of component realism is the problem of saying whether or not such decompositions ever correspond to a part-whole structure in the forces themselves, and if so, when? The chief objection to component realism is based on the thought that

---

case of micro-reductions was not that part and whole must have *every* property in common, but only that there must be some common property that allows us to understand what parthood means.

components of force, motion, and the like cannot be real features of a system because they are not *observable*. Marc Lange (1994) notes the unique character of this type of unobservable while discussing the problem of component realism as it relates to ancient astronomers:

> [C]omponent motions and forces differ fundamentally from, e.g., Salmon's (1985: 12) example of words in the Compact Edition of the *Oxford English Dictionary* that are typed so small as to be imperceptible without magnification. Since we have long observed real typed words of the same kind, only larger, there was – before unobservable words were posited – an established difference in practice between believing such a thing 'real' and believing that it accords with certain observations. The same cannot be said of components, fields, lines of force, the id, and some other entities posited by scientific theories. (Lange 1994: 122)

And similarly, in the examples of multi-scale analysis in chapter three, we don't directly observe any part of the systems following the trajectory associated with the slow-system. This unobservability leads the 'component anti-realist' to believe only that a decomposition into component forces or motions and the like yields accurate predications about observables, without truly describing any unobservable entities.

Since discussions of component realism typically focus on component *forces* (with some notable exceptions I will discuss below), I will begin by doing likewise. After reviewing some prominent positions on component realism, I'll return to the more specific cases arising from multi-scale analysis to see how the objections to component realism can be addressed.

The starting point for contemporary arguments about component realism is Nancy Cartwright's (1980, 1983) rejection of a realistic interpretation of physical force laws. Cartwright's argument comes as part of her famous rejection of the 'facticity' of the laws of physics. She argues that laws such as the universal law of gravitation and Coulomb's law of electrical attraction/repulsion are not 'factual' because in general they don't correctly describe how the things they apply to behave. That's because, contrary to what the law of gravitation says (for instance), two massive bodies will not attract each other with a force $F = Gm_1m_2/r^2$ in the quite common event that they are also charged. In such cases the occurrent force between two bodies will be quite different from that given by the law of gravitation *and* from that given by Coulomb's law, so neither law correctly 'states the facts'.

The obvious response to this suggestion would be to maintain the facticity of laws and insist that those occurrent forces are the resultants of real component forces 'added' together: that is, that component forces are *parts* of the resultant force. But though she admits that such laws are obviously useful for constructing explanations, Cartwright rejects this view:

> The vector addition story is, I admit, a nice one. But it is just a metaphor. *We* add forces (or the numbers that represent forces), when we do calculations. Nature does not "add" forces. For the "component" forces are not there, in any but a metaphorical sense, to be added: and the laws

which say they are there must also be given a metaphorical reading. (Cartwright 1983: 78)

I'll examine Cartwright's reasons for rejecting the parthood claim in a moment. But first, to make Cartwright's position clear, consider the famous 'oil-drop' experiment conducted by Robert Millikan to measure the charge on an electron. In simplified terms, Millikan sprayed electrically charged drops of oil into a chamber in which he could control the strength of an electric field. Millikan then measured the electromagnetic field strength required to suspend the drops of oil mid-air. From previous measurements of the mass of each drop and from knowledge of the force of gravity, he was able to calculate the electromagnetic force acting (in opposition to the force of gravity) on each drop. Knowing the field strength of the electromagnetic field, he was then able to calculate the electric charge on individual drops of oil. After varying the charge applied to individual drops, Millikan found that the charges were always multiples of -1.6×10^{-19} Coulomb, which he concluded was the charge on a single electron[60]. The description of an individual oil-drop is schematically illustrated in figure 4.1.

[60] The actual procedure and calculations involved were significantly more complex than this, though these were the basic features of the experiment.

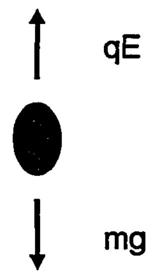

Figure 4.1 Oil-drop suspended in an electric field. E is the electric field strength, q is the charge on the drop, m is the mass of the drop, and g the acceleration due to gravity.

On the component realist interpretation, a drop of oil suspended in mid-air is subject to two equal and opposed forces. Since these forces *are* equal and opposed, the 'net' or resultant force acting on the drop is the null force. Cartwright's position is that while this is clearly a useful way of thinking about how electrical fields and gravitational fields interact, the only *true* ('factive') account is that a suspended drop of oil is subjected to *no* forces: after all, that's why it doesn't move. And this null force is not the sum of any actually present opposing forces: the component gravitational and electromagnetic forces are not there to be added in the first place.

Cartwright's position is that only 'resultant' forces are real, while component forces are merely useful theoretical postulates. Such a view could obviously be extended to the claims about component systems in multi-scale analysis: since it is the 'resultant' behavior that is observed, decompositions into 'fast' and 'slow' component systems are merely useful theoretical or 'metaphysical' devices that are not corresponding to anything 'there' in the world.

Note that there are actually two questions of component realism. We can ask about both the reality of *components* and their reality as *parts*. That is, one way of rejecting Cartwright's claim would be to defend the reality of *component* forces alone and argue that it is the *resultants* that are not real. This is the view defended by Lewis Creary (1981). Creary notes that the only realist position Cartwright considers (and rejects) is the one according to which component forces are *parts* of a given resultant force. Creary agrees with Cartwright in rejecting this suggestion, and admits that this leads to a difficulty in accepting the reality of both component forces and resultant forces, since "if one...took for granted the reality of overall resultant forces, then one would naturally be led to conclude that component forces are unreal, since one would otherwise have to regard them, most implausibly, as physically redundant real forces that 'shared' their effects with their (presumably real) resultants" (Creary 1981: 152). That is, if one rejects realism about *parthood*, accepting the reality of both component forces and resultant forces would lead to an unacceptable redundancy in the world. Creary argues that the preferable realist alternative to Cartwright would be to defend the reality of components while rejecting the parthood claim, and thus reject the reality of *resultants*.

Creary's argument is that realism about components rather than resultants offers a better explanation of the utility of various force laws since it provides a 'causal covering law' account of explanation, and it allows us to account for what would happen in counterfactual situations when not all of the actual forces operating are present (or when different forces are in effect). Clearly examples such as Millikan's oil-drop experiment make it difficult to see how we could seriously account for scientific activity without supposing that the reality of component forces – and not just their ability to 'save the phenomena', as the anti-realist slogan says – is an important assumption in scientific explanations. And if the purpose of Millikan's experiment was to measure the charge on an electron, and if that measurement is based on the assumption that the presence of electrons in a drop of oil (i.e., a charge on the drop) results in that drop being affected by components of electromagnetic force, then anti-realism about components of force seems to lead to anti-realism about charge as well. More everyday examples make the oddity of Cartwright's anti-realism clear as well: for instance, it seems perverse to suppose that a hovering helicopter experiences *no* forces, rather than holding that it experiences two exactly opposed forces, or to think that a parachutist gently drifting to Earth is being pulled down by a lesser force than that attracting a second parachutist (of the same mass) who has yet to open his chute.

Note that in her response to Creary's argument, Cartwright does concede that there are *some* legitimate reasons to believe in the reality of components. She grants that Creary's account of component forces gives "a plausible account of how a lot of causal explanation is structured" (Cartwright 1983: 63), but questions its general applicability. Her worry concerns the existence of general principles of composition that specify how different components combine to produce a collaborative effect: in the case of vector-valued forces, this principle of composition might be clear, but Cartwright doubts that this principle can be generalized, since "theories can seldom specify a procedure [for combining 'causes'] that works from one case to another" (Cartwright 1983: 63).

John Bigelow and Robert Pargetter (1990) defend a compromise position between Cartwright's and Creary's. They suggest that the reality of components or resultants is properly settled on a 'case by case' basis:

> Sometimes when we speak of component forces and their resultant it is the components that are real and the resultant which is not; sometimes it is the other way around. There is a principle solution, because which way to go is determined by the physical features of the situation in question" (Bigelow and Pargetter 1990: 108)

To some extent Bigelow and Pargetter's account is true. Consider their example of a case where components clearly seem real:

> Consider three protons, isolated from outside interference, one at the midpoint of the line between the other two. Changes to the ones on the ends will involve forces between each of them respectively and between each and the one in the middle. Yet the principle of action and reaction of forces, so fundamental to physical theory, thus has the middle proton the subject of two balancing component forces which jointly result in its non-change. (Bigelow and Pargetter 1990: 108)

This sort of example is clearly in the spirit of Creary's suggestion that it is unreasonable to reject component forces since positing them brings such widespread explanatory power to our scientific theories. For an example where component forces are clearly not real, Bigelow and Pargetter consider the behavior of a particle subjected to a force, yet confined to move in only one direction (presumably where the imposed force is not parallel to that direction):

> [In such cases] we standardly resolve that force finding two fictitious orthogonal resolutes, one in the direction in which the particle is free to move, which if taken as component forces would have when aggregated result in the same effect, that is, in the same change in state of the particle. (Bigelow and Pargetter 1990: 108-9)

While they may be right that in such cases the 'fictional' nature of components is clear, it seems unwarranted to expect that in all cases the physical nature of the situation will make it obvious whether it is component forces or resultant forces that are 'real'. I mentioned earlier that in the case of multi-scale analysis, scientists *talk* as if those components are real, and as if they reflect real physical features of the system of interest. Yet since the 'resultant' behavior is the observed behavior, it is difficult to see how we could choose between the components and their resultant just by taking their 'physical nature' into consideration. For a concrete example of a different sort, consider the following case described by Neil Sheldon (1985). As a challenge to both Creary's and Cartwright's views, Sheldon examines how radio waves can be represented in various ways. In radio broadcasting, an amplitude modulated (AM) wave consists of a carrier wave at a particular frequency, which can be represented as

(1) $$y = a\cos(2\pi\varphi t)$$

where φ is the frequency of the wave and a is the amplitude. Variations in the amplitude of such a wave represent the signal. For example, a signal of frequency f and amplitude b can be represented as

$$(2) \qquad y = [a + b\cos(2\pi ft)]\cos(2\pi\varphi t)$$

This represents a single wave at frequency φ, but with variable amplitude $[a + b\cos(2\pi ft)]$. But we can also write equation (2) as

$$(3) \qquad y = \tfrac{1}{2}b\cos(2\pi(\varphi - f)t) + a\cos(2\pi\varphi t) + \tfrac{1}{2}b\cos(2\pi(\varphi + f)t)$$

This represents a decomposition of the wave into three component waves of different frequencies ($\varphi - f$, φ, and $\varphi + f$) each have a constant amplitude ($\tfrac{1}{2}b$, a, and $\tfrac{1}{2}b$, respectively). Figure 4.2 illustrates these contrasting representations.

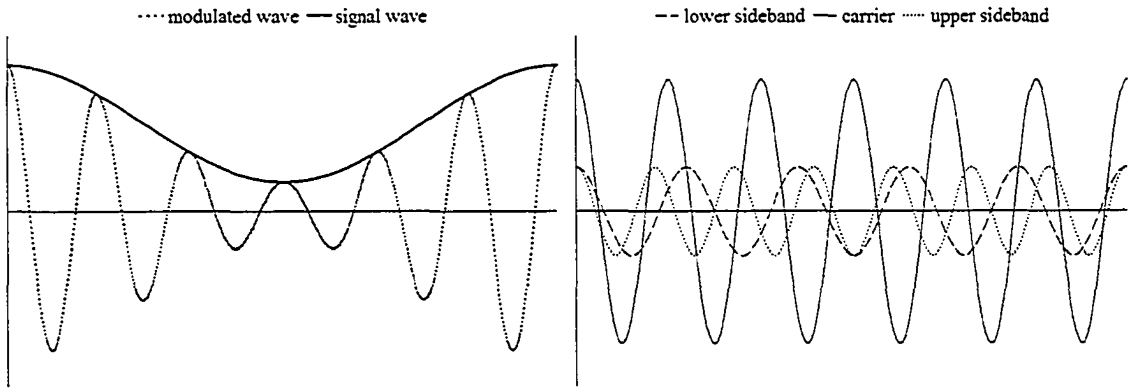

Figure 4.2 One Wave or Three?

Sheldon points out that (3) is a simple example of Fourier analysis of equation (2). According to Cartwright's view, equation (2) should represent 'the facts' about the wave, while equation (3) represents some artifice of mathematics; conversely, on Creary's view, (3) might represent the facts, while the 'resultant' wave represented in (2) is merely artificial; finally, on Bigelow and Pargetter's view, which equation is to be regarded realistically and which treated as merely artificial is to be determined by the physical nature of the situation.

However, Sheldon argues that there are good reasons to regard *both* representations (2) and (3) realistically. For instance, the wave described by (2) is the natural result of the way radio signals are produced: a carrier wave (given by equation (1)) is passed through an amplifier which selectively 'boosts' the amplitude of the carrier wave from moment to moment to conform to some signal. We apparently have a causal (even 'mechanical') account of how (2) is produced, giving it a strong claim to reality, even by Cartwright's strong demand in terms of observable behavior. There is no obvious analogous causal account of the origins of the component waves described in (3). However, these component waves are relevant to causal explanations of interference between AM waves at close frequencies. Remember that in (2) the broadcast wave is a wave at a single frequency with varying amplitude. Such waves can interfere with each other, disrupting signal transmission if their frequencies are too close together (a phenomenon known as 'sideband interference'). The explanation for this interference is given by equation (3): according to *that* representation, a signal consists of waves at a variety of frequencies centered on the 'carrier' frequency. These waves at slightly higher and lower frequencies are known as 'sidebands', and if two carrier signals are too close to one another, their sidebands will interfere with one another. Again, even by Cartwright's standards, this representation seems to have a legitimate claim to reality as well. This example puts Creary in as much of a bind as it does Cartwright, since there doesn't seem to be any reason to single out the components (the three wave representation) as real and dismiss the resultant as artifactual. And contrary to Bigelow and Pargetter's expectation, the physical nature of the situation does not appear to determine *one* representation as the 'real' one.

Note that it's true that Sheldon's example is different from those discussed by Cartwright, Creary, Bigelow and Pargetter in that it is not specifically concerned with *forces*. Perhaps that difference would seem relevant to those authors, and so perhaps it would be unfair to assume that they would maintain the same positions with respect to wave components and resultants as they do with respect to force components and resultants. Still, Sheldon's example suggests that none of the preceding views can be maintained as *general* accounts of components and resultants.

Sheldon is content to present his example as a problem for both Cartwright and Creary, without drawing any further conclusions about the relationship between components and resultants. However, I think the reasonable conclusion to draw from this example is that Cartwright and Creary (and presumably Bigelow and Pargetter – though they don't discuss this suggestion explicitly) were wrong in thinking that components could not be *parts* of resultants. This was the one thing Cartwright and Creary agreed on, and that agreement led them to maintain opposing views about components and resultants. Since they agreed that components could not be parts, they each concluded that components and resultants could not both exist, on pain of over-determining their effects. Since it seems that both components and resultants *can* co-exist, Cartwright and Creary's rejection of the parthood claim must be wrong: components *can be* parts of resultants, after all.

With this suggestion in mind, let's turn to the second question of component realism – the question of whether or not components are *parts*. The answers to the *first* question – the question of whether components or resultants are *real* – all began with the assumption that if they were real, components could not be parts of resultants. Why not?

Bigelow and Pargetter's example of the three protons recalls Bertrand Russell's (1903) argument against viewing resultant forces as genuine 'sums' of their components:

> Let there be three particles A, B, C. We may say that B and C both cause accelerations in A, and we compound these accelerations by the parallelogram law. But this composition is not truly addition, for the components are not *parts* of the resultant. The resultant is a new term, as simple as their components, and not by any means their sum. Thus the effects attributed to B and C are never produced, but a third term different from either is produced. (Russell 1903: 477)

Ernest Nagel (1952) responds by pointing out that Russell's argument shows only that "by the component of a force...we do not mean anything like what we understand by a component or part of a length – the components of forces are not spatial parts of forces" (Nagel 1952: 22). Instead, Nagel argues that parthood between forces is a matter of the functional relationship existing between components and their resultant: "the 'sum' of a given set of elements is simply an element that is *uniquely determined* by some *function* (in the mathematical sense) of the given set" (Nagel 1952: 23). However, this response doesn't seem to get at the heart of the matter: Nagel appears to accept Russell's claim that 'resultant' forces are *simple* in the same way their components are. But *that* claim seems unjustified: why should we assume that forces are 'simple'? It's true that we *represent* quantities such as forces with vectors defined as elements in a vector space, and that the element representing a resultant force is distinct from the elements representing its components; but that doesn't mean that the things represented – the forces themselves – can't have structure.

Cartwright's reasons for rejecting the parthood claim are somewhat mysterious. For example, in considering John Stuart Mill's claim that "[i]f a body is propelled in two directions by two forces, one tending to drive it to the north, and the other to the east, it is caused to move in a given time exactly as far in *both* directions as the two forces would separately have carried it; and is left precisely where it would have arrived if it had been acted upon first by one of the two forces, and afterwards by the other." (Mill 1967: 243), Cartwright suggests that such parts are simply counter-intuitive:

> Mill's claim is unlikely. Events may have temporal parts, but not parts of the kind Mill describes. When a body has moved along a path due north-east, it has traveled neither due north nor due east. The first half of the motion can be a part of the total motion; but no pure north motion can be a part of a motion that always heads northeast...The lesson is even clearer if the example is changed a little: a body is pulled equally in opposite

directions. It does not budge, but in Mill's picture it has been caused to move both several feet to the left and several feet to the right. (Cartwright 1983: 60-1)

Cartwright's argument against the reality of component *forces* is then that if the *behavior* caused by a force doesn't have parts, and if to 'state the facts' is to describe the causes of real behavior, then any explanation of that behavior in terms of 'parts' of forces cannot 'state the facts'[61]. However, Cartwright doesn't explicitly say *why* events couldn't have the parts Mill describes. The only argument she gives is rather cryptic. She writes "The first half of the motion can be a part of the total motion; but no pure north motion can be a part of a motion that always heads northeast. (We learn this from Judith Jarvis Thomson's *Acts and Other Events*)" (Cartwright 1983: 60-1). However, as far as I can tell, Thomson never discusses this sort of example specifically (and Cartwright makes no specific page reference). The only point – at any rate the only point *I* can find – in Thomson that is perhaps relevant to this issue is an assumption Thomson makes *on page 63*, to the effect that for any 'cause' C, and any events x and y, "C causes y if and only if C causes all of y's parts". From this we might conclude that a pure north motion cannot be part of a motion that always heads northeast

[61] Note that Cartwright herself proposes that component laws *do* describe 'causal powers'. "the law of gravitation claims that two bodies have the *power* to produce a force of size Gmm'/r^2. But they don't always succeed in the *exercise* of it. What they actually produce depends on what other powers are at work, and on what compromise is finally achieved among them. This may be the way we do sometimes imagine the composition of causes. But if so, the laws we use talk not about what bodies do, but about what powers they possess" (1983: 61). Obviously, Cartwright's view's on laws of nature are complex, and I don't intend to do justice to them all here; I'm simply trying to present the arguments she gives *against* component realism.

because something might cause the northeast motion without causing the pure north component motion. If this *is* the lesson we are supposed to learn from Thomson's book, it seems that the more appropriate conclusion to draw is that Thomson's assumption about the parts of events is just wrong[62].

Perhaps Cartwright's reasons for rejecting the parthood claim are similar to those given by Paul Teller in his (1995) discussion of component realism. Teller's discussion comes in the course of a discussion of the proper interpretation of Feynman diagrams in quantum physics. Feynman diagrams are graphical representations of equations for quantum states that depict events of particle generation, transformation, and annihilation. Teller argues that, though enormously useful, such diagrams should not be taken to be literal depictions of quantum-level events. One of the reasons he gives is that typically the equation corresponding to a given diagram will be only part of a much larger equation – the superposition of a variety of quantum states. And while it is tempting to see each individual diagram as pertaining to 'part' of the entire process of particle creation, propagation, and annihilation, Teller warns that such a conclusion would involve an equivocation over the term 'part'.

> An ordinary chair has what I will call *mereological parts*, such as its back and each of its four legs. A mereological part of a whole has its own

[62] Later in her discussion, Cartwright does give an example involving components of electron energy states in carbon atoms that she takes to be definitive proof that the parthood claim cannot be maintained. I suspect that this example can actually be accommodated by the component realist, but I lack the technical expertise to show this definitively. In any case, one fairly specialized counter-example does not show that the parthood claim can *never* be maintained. See Spurrett 2001 for discussion of this example.

trajectory in space and time. And a mereological part, in general, itself has parts that are also parts of the original whole – a leg of a chair has parts, say its top and bottom halves, each of which is also a part of the chair. (Teller 1995: 139-140)

In contrast to 'mereological' parts, Teller characterizes component Feynman diagrams, along with the components of Fourier analysis and component motions as 'analytic' parts, and insists that while these may 'compose' a resultant of some sort, this is not a genuine form of parthood. For instance, given a Fourier decomposition of a wave into component sine and cosine wave functions, "it seems dubious to say that the analyzing pure wave forms are 'parts' of the analyzed wave in anything like the sense in which the leg of a chair is a part of the chair" (Teller 1995: 140). Hence only parts of the former sort should be taken to reveal real components of things in the world.

We might think that the distinction Teller has in mind is that between *spatial* parts – "parts having their own trajectory in space and time", as Teller puts it, such as the legs of a chair – and non-spatial parts (parts overlapping in space and time). If this were the case, then his suggestion that only 'spatial' parts can be 'mereological' parts would make exactly the assumption about parthood we rejected earlier: nothing in mereology *per se* requires that parts be spatial parts. However, despite his definition of mereological parts in terms of space-time trajectories, Teller's argument against analytic parts is not really based on intuitions about whether or not genuine parts are spatial. Instead, he argues that analytic parts fail to satisfy the axioms of mereology: hence, assuming that mereology *is* the proper theory of parthood, analytic parts are not genuine parts.

Teller focuses on the axiom of transitivity: the claim that parts of parts are parts, as he puts it. He claims that transitivity is either problematic for or outright violated by analytic parts, and hence that 'componenthood' and parthood are very different relations. Teller doesn't say in detail why analytic parthood is *not* transitive, but does say that if we understand analytic parts as parts relative to a given basis (the basis for the decomposition of the vector, for example), "the relativity to a basis seriously interferes with seeing how the condition [that is, transitivity] is to be satisfied" (Teller 1995: 140). The reason for this is presumably that since the components of a vector relative to a given basis will be scalar multiples of those basis vectors, the components of *those* vectors – the 'parts of parts' – will simply be the basis vectors themselves: component vectors relative to a basis have no *proper* parts (at least not relative to that same basis). That doesn't actually violate transitivity, but it does make parthood somewhat uninteresting: vectors might have 'parts', but those parts won't in turn have any *proper* parts. Relative to a basis, there is only one way for a vector to be decomposed.

An alternative account of analytic parthood Teller considers is to drop the relativization to a basis and regard *any* set of vectors that 'add up' (via vector addition) to a given vector to be the parts of that vector: "starting with a given whole constituted by a vector quantity, all the vectors in any sum that add to the original whole count as analytic parts of that whole" (Teller 1995: 140). Since this appears to be the most general claim of component realism, call this the 'naïve' approach to component realism: any set of vectors adding up to a given vector are parts of that vector.

The naïve approach secures transitivity, but at a cost. As Teller points out, "my standing stock still [would have] as analytic parts my simultaneously moving at 100 miles an hour to the north and 100 miles an hour to the south", and, in a different context (but equally relevant here), "traveling at 50 miles an hour due east would count as an analytic part...of traveling 100 miles an hour due north" (Teller 1995: 140). This is the same sort of example that led Cartwright to reject Mill's claim about components of motion.

Teller and Cartwright – if her reasons *are* similar to Teller's, as they appear to be – both simply assume that it is obvious why this outcome is objectionable, and hence why the naïve approach to analytic parthood cannot be maintained. But it is helpful to spell the possible objections out in detail. The first possible objection is that viewing components as parts in this way trivializes the parthood relation: *any* vector will be part of any other vector, since for any pair of vectors V_W and V_P, we can always find vectors V_1, V_2, ... such that V_W is the vector sum of V_P, V_1, V_2, etc. If vectors, wave forms, and the like have all of their components as parts, then pointing out that one vector is part of another doesn't reveal any special relationship between them, since every vector will be part of every vector in this sense. The naïve approach to component realism clearly leads to an overabundance of parts: if all possible components are genuine parts, there are just too many parts to be believed.

Another objection to the naïve view suggested by Teller's examples (and Cartwright's) is that if components *are* parts, then things will have some surprising parts: it is simply incredible to think that my standing stock-still has my traveling 100 miles an hour to the north as one of its parts. Since such 'parts' are so counter-intuitive, it is supposed to seem obvious that motions don't have component motions as parts in the same way that a chair has legs and a back as parts. Thus, together with Teller's earlier worry about transitivity, we have the following objections to 'analytic' parts:

a. componenthood is not transitive, whereas parthood must be transitive; or

b. componenthood is trivial, whereas parthood is not trivial; or, components are overabundant, whereas parts are not.

c. Examples of analytic parts are simply *strange*: it is counter-intuitive to think that the entities or properties referred to by components are parts.

With these objections in minds, the appropriate response to the critics of component realism seems clear: for the first two objections, we need to show that the components we're trying to be realists about obey the axioms of mereology, and that they are not trivial or overabundant. The third objection is more difficult: counter-intuitiveness alone doesn't seem like a very good reason for rejecting the reality of anything, unless we want to defend a purely 'common-sense' understanding of the world (and I don't). However, there does seem to be *some* burden of proof on the part of the realist to show that it is not inconceivable that components are real parts. I'll now address these concerns about component realism in general and see how they can be met in the specific cases involving the non-spatial parts I discussed in chapter three.

Natural Parts

If component realism is incredible because it leads to an overabundance of parts, then we should be able to make it more credible by limiting the variety of components that count as parts. Just because not *every* decomposition of force or a wave or a dynamical system corresponds to real parts doesn't mean that none can. What we need is a distinction between decompositions that are 'artificial' and in some way 'unreal', and those that are 'natural' and reveal genuine parts. Call this the 'naturalized' view of parthood: to be a genuine part is to be an element of a natural decomposition of something.

The idea that some decompositions can be distinguished as particularly natural has been suggested by Marc Lange in his discussion of component realism in classical astronomy (see Lange 1994). The question at the center of Lange's discussion concerns whether or not ancient astronomers such as Proclus advocated views best characterized as 'realist' or 'anti-realist' about the possible decompositions of planetary motions. Ancient astronomers knew that the motion of planets can be described in two distinct ways. On the one hand, a planet's position in the sky can be predicted accurately by a model according to which that planet moves along (roughly) an epicycle turning on a deferent. The 'deferent' is a circular path centered on the Earth, while the 'epicycle' is a smaller circular path centered on the path of the deferent (see the left diagram in figure 4.3).

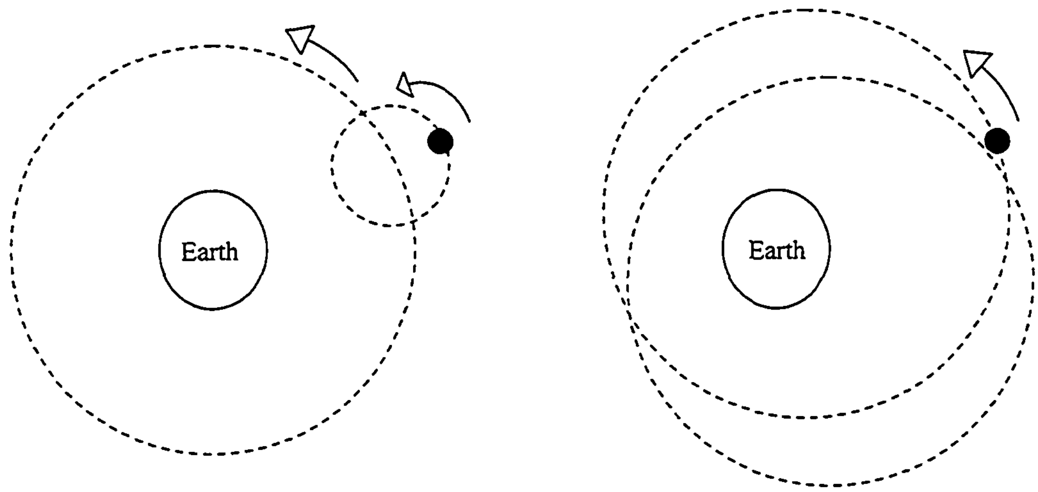

Figure 4.3 Orbital path of a planet represented as an epicycle and deferent (left) or as an 'eccentric' (right)

On the other hand, the same orbit can be represented as following an 'eccentric' path, moving uniformly around a circle (the 'eccentric') that contains earth but whose center does not contain with the center of the Earth. Additionally, the eccentric itself orbits the Earth, so that the over time the center of the eccentric traces out a circular orbit around the Earth. (See Lange 1994: 112, 125). These two decompositions of the motion of a planet will coincide in space, so there is no observable difference between the two.

On Lange's interpretation, Proclus suggests that *neither* decomposition can be regarded as 'real', but instead both represent merely convenient representations of each planet's motion. Lange's interest is in what it meant for Proclus to consider a particular decomposition 'real'. One realist view would be to hold that the planets are actually 'fixed' to rotating material spheres. If that *were* true, then obviously the decomposition correctly describing the motions of those spheres would be the 'real' one, whereas any phenomenally equivalent ones that described different combinations of spherical motions would not. But clearly the planets are not fixed to material spheres. Lange argues that this sort of realism was not the central issue for Proclus: instead, the real question concerned the reality of component *motions* of the planets. On his account Proclus was an anti-realist about these as well, rejecting the suggestion that either the decomposition into epicycles and deferents *or* the decomposition into eccentrics represent true components of the motion of the planets. This isn't because those component motions are not guided by material spheres, but instead because those components are not *natural*, in the sense of conforming to laws of nature:

> A planet's 'real' motions are the component motions natural to it, i.e., necessitated directly (rather than 'accidentally') by its nature. To say that a theory decomposes a planet's orbit into the planet's real component motions is to say that the theory consists of statements of natural law. Without using these laws, one can predict but cannot explain a planet's position in the sky at a given moment. (Lange 1994: 115)

So while many sorts of decomposition might 'save the phenomena' and adequately describe the net motion of a planet, only a decomposition that represents a planet's *natural* motion need be taken ontologically seriously. Natural components of motion are those described by natural law, and hence treating such decompositions realistically means that the same sorts of decompositions 'save' other sorts of phenomena as well:

> if one believes that the components of a given decomposition are real, then one must regard the fact that some other theory involves the same kinds of components (i.e., components governed by the same natural laws) as confirming that theory's capacity to save its phenomena. (Lange 1994: 113)

Lange contrasts Proclus' *rejection* of realism about the component motions of the planets with Newton's realism about the decomposition of the gravitational force exerted by a spherical body. Newton calculated this force by taking the sum of an infinite number of component forces, each corresponding to the gravitational force exerted by an infinitesimal portion of the sphere:

> Newton considered this sum to be the natural decomposition of the sphere's force-vector, in the same sense in which Proclus denied that epicycles and deferents constitute a planet's natural component motions. That is, Newton held that each component of this sum, representing the gravitational force between a point-particle outside the sphere and an infinitesimal element of the sphere, is given directly by a physical law: the two mass-point gravitational-force law $F = GMm/r^2$. (Lange 1994: 118)

By contrast Proclus denied that the decompositions into epicycles and deferents or into eccentrics revealed components that could be subsumed under any laws. Instead, such decompositions were "like seven different theories, each attributing to one planet several component motions that, although able to save the phenomena, bear no simple relation to each other or to the component motions of other planets" (Lange 1994: 119). On Newton's realist view, the decomposition of force was real precisely because the decomposition of the force exerted by a sphere produces components that could be characterized by a general law. Lange puts the same point in terms of natural *kinds*:

> [T]he epicyclic components given by decomposing all of the Ptolemaic planetary orbits do not constitute a natural kind, nor do the deferential components, whereas all of the component gravitational forces between pairs of point-masses constitute one natural kind, all of the component electric forces another. To distinguish a kind as 'natural' rather than 'artificial' is, at least in part, to say that in several distinguished respects, all cases of this kind bear some simple similarity to one another. (Lange 1994: 119)

This idea of the naturalness of certain decompositions seems appropriate for the case of component forces we've already discussed. In fact, along with rejecting the reality of resultants, Creary also draws a distinction between 'natural' components and 'mathematical' ones in a manner similar to that suggested by Lange:

> A distinction is made between *natural* component forces, which arise directly from the action of various real physical causes, and *mathematical* component forces, which arise merely from the artificial resolution of vectors, and thus lack physical existence. (Creary 1981: 151-2)

Creary doesn't say just what 'the action of real physical causes' means, but one plausible reading of it is that genuine forces are those described by physical laws: the real physical cause of the force 'pulling' the drop of oil in Millikan's chamber upwards is the result of the charge on the drop and the presence of the electric field in the chamber, as described by the laws of electro-magnetism.

Note that whatever the merits of the account of 'natural' decompositions as those whose components are described by laws of nature, it may seem question-begging to offer it as part of an argument against Cartwright's component anti-realism, since Cartwright's position was precisely that the laws of nature do *not* describe real components. We might think that the dialect is this: Cartwright argues that that the laws of nature do not 'state the facts' because they describe component forces and motions, and these things are not real; in response, the realist argues that component forces and motions *are* real, and that they are real precisely when they are described by laws of nature[63]. Thus the dispute between Cartwright and the realist seems simply to be a stand-off. However, I don't think this is quite the right way to view the argument. Cartwright's argument that laws don't 'state the facts' depends on her maintaining that component forces and motions cannot plausibly be *parts* of resultant forces and motions: she herself admitted that the vector-addition account that seemed to preserve the reality of both was 'a nice story'. The argument against that 'story'

[63] I should note that Lange's (1994) discussion is not specifically directed at Cartwright's views. In correspondence, however, he has expressed his interest in defending component realism of some sort against arguments such as Cartwright's, though not in a 'wholesale' manner.

was that the component forces and motions were not 'there' to be added in the first place, and *that* was because there was something unacceptable about treating such components as genuine parts. So Cartwright's argument against realism about laws is a reductio: we can't be realists about laws because if we were realists, we'd end up with unacceptable claims about parts of forces, motions, and the like. Since Cartwright didn't give much in the way of explicit argument against such parts, I've interpreted her objections as similar to those raised by Teller, and among *these* is the thought that if components *were* parts, there would be too many parts to be believed, and that parthood would in fact be trivial. But if it is only 'natural' decompositions that reveal genuine parts, then there are not too many parts to be believed, parthood is not trivial, and the nice 'vector-addition story' doesn't have to be just a story. Of course there were other complaints about treating components as genuine parts, but at least we can see that there are principled reasons to consider *some* but not all components genuine parts.

With this idea about the *general* issue of component realism in mind, let's consider the more specific case of realism about components in a perturbation expansion. Prima facie, defending the reality of these component systems shouldn't be that difficult: after all, the components involved are neither arbitrary nor overabundant, since we *can't* find an appropriate (that is, asymptotic) expansion using just *any* terms. This was illustrated in chapter three in sketching out the failure of the 'regular' perturbation expansion: the resulting expansion in that case was wildly divergent and hence couldn't be treated realistically. In general, finding the appropriate scales for a multi-scale expansion requires considerable care. While in general perturbation problems can be solved in a variety of ways, they cannot be solved using just *any* scale. What's more, the perturbation examples seem to conform to Lange's criterion of naturalness: the requirement that *real* component processes form natural kinds that conform to laws of nature. We typically expect this sort of conformity in perturbation expansions, since it is usual for the leading term in an expansion to represent the behavior of an 'ideal' system while the remaining terms represent perturbations of that ideal. That may not *always* be the case, but this is just a constraint on component realism: to the extent that decompositions *do* fall under laws of nature, their components can be considered real.

Note also that Cartwright's concession to Creary's argument – that component realism may be reasonable if a 'principle of composition' is available that tells us how 'component causes' combine to form composite causes – seems to be satisfied in general in the perturbation examples. Cartwright claimed that in Creary's (and her own) simple examples involving gravity and electro-magnetic charge, the idea of component realism *is* compelling, but argued that in general we don't have an account of how different 'causes' combine. However, in the perturbation cases, clearly we have *some* idea of how the different components combine. Admittedly, the combination of fast and slow components is not as simple and straight-forward as the vector addition characteristic of combinations of classical forces: to combine the component systems in the perturbation examples, we had to apply a 'matching' condition to particular limits of the components, and as I mentioned in chapter three, this is usually not a trivial procedure. However, while this sort of 'principle of composition' isn't so simple and isn't entirely general, once it has been accomplished in specific cases there doesn't seem to be any reason to doubt the reality of the parts on the grounds that we have no idea how they combine.

Aside from the worry about the overabundance of parts, the two other worries about component realism concerned the formal properties of components and the inherent strangeness of some components. In the *general* case of component realism, transitivity isn't obviously a problem. Teller's claim that 'analytic' parts violated transitivity was based on the assumption that analytic parts had to be defined relative to a basis; adopting the 'naïve' view that *any* decomposition into components revealed the parts of a vector-valued quantity resolved this worry, but at the cost of making parthood trivial. We've now added a condition on such decompositions: only decompositions into components that themselves form 'natural kinds' are to count as revealing genuine parts. Since transitivity wasn't the problem for the naïve view (overabundance and triviality were), it isn't obvious that it would present any difficulties for the 'naturalized' view.

In the specific case of non-spatial parts and perturbation theory, it's somewhat difficult to *prove* that transitivity is obeyed. Given a perturbation expansion for a given system, we'd need to consider expansions for the terms of that expansion and show that the original system could *also* be represented by a perturbation expansion involving the terms of *those* expansions. Intuitively, it might seem that this is straightforward, involving a simple substitution of terms:

given an expansion D = A + B, and an expansion A = E + F, we expect an expansion D = E + F + B. But the main point of the 'naturalized' view of parts I've suggested here is that simply being representable as a sum of terms isn't enough to support a claim of parthood (if it were, we would be open to the objection about the overabundance of parts). For the expansion D = E + F + B to count as revealing real components of D, we'd have to show that the other conditions were met: for instance, that the appropriate matching conditions between the limits of E, F, and B were satisfied. Perhaps this can be done, but I lack the mathematical expertise to show that it *can* be done.

With this unsettled question in mind, though, we might object to this insistence on transitivity: while transitivity *is* a standard axiom of the parthood relation, perhaps that's simply a mistake. The theory of mereology, after all, might be *wrong* in its account of parthood. Nicholas Rescher (1955) argues precisely this, citing three features of standard mereology that seem to violate intuitions about what count as parts and wholes. Notably, Rescher is explicit in his intent to provide an account of parthood suited to a 'scientific' conception of parthood (a goal he suggests also motivated Lesniewski's original formalization of mereology). Rescher presents a revised set of axioms for mereology that explicitly avoids a commitment to transitivity, on the grounds that many real world examples of parthood are *not* transitive. Admittedly, though Rescher frames his discussion as a development of a theory of parthood designed to accommodate a scientific world-view, his motivating examples of the failure of transitivity are not particularly compelling. For example, he claims that the organelles of a cell are not parts of the body containing that cell. A slightly more compelling, though even less 'scientific', example comes from Rosen and Dorr (2002): while it seems natural to say that Fred (a person) is part of the conga line, and that Fred's kidney is part of Fred, it seems somewhat unnatural to say that Fred's kidney is part of the conga line.

David-Hillel Ruben (1983) argues that there is a non-transitive sense of parthood, but suggests examples of the failure of transitivity show that a distinction must be drawn between the relation of being 'a part of' and being 'one of the parts of'. He considers the following example:

> Suppose that I have an alarm clock, one of whose parts is a light by which one can tell the time at night. Suppose that I use the clock to make a bomb that detonates when the clock's hands reach a specified position. The light makes no contribution to the working of the bomb. The light is one of the parts of the clock, and the clock is one of the parts of the bomb, and yet it is false that the light is one of the parts of the bomb. This seems to show that the relation of being a part of is non-transitive. (David-Hillel Ruben 1983: 231-2)

Ruben argues that while being 'a part of' is transitive, being 'one of the parts of' is not: being one of the parts of something has "quasi-functional implications", presumably meaning that to be one of the parts of something, a thing must play some role in the 'functioning' of that composite entity.

Perhaps there is linguistic evidence to support Ruben's contention about the use of the phrases 'a part of' and 'one of the parts of'; in any case, whatever the phraseology, I think Ruben's point is a useful one about the parthood relation. Being a genuine part of something is more than just being contained within that thing: being a part is a matter of being an element of a decomposition that is in some way 'natural'. Roughly, for Ruben's bomb, since the light in the clock plays no role in the functioning of the bomb, examining the decomposition of the bomb that includes the light doesn't tell you anything about which you can generalize to instances of other bombs. The decomposition of the *clock* that includes the light, on the other hand, *does* tell you something you can generalize to instances of other clocks[64].

So perhaps there is *some* room for debate over whether or not transitivity is an essential feature of the parthood relation[65]. And I'll repeat that the issue for

---

[64] I admit the example is somewhat strained, especially given the obvious fact that there are no *laws* about lights, clocks, and bombs. However, the idea of distinguishing some decompositions as in some way more natural than others does seem to be at the heart of the example Ruben discusses.

[65] Simons (1987) apparently disagrees. In introducing the basic axioms characterizing the part-relation, Simons comments "These principles are partly constitutive of the meaning of 'part', which means that anyone who seriously disagrees with them has failed to understand the word" (Simons 1987: 11). While I confess I'm not entirely sure I *do* understand the word, as I noted earlier Simons' entire book on the subject offers little help: he explicates the concept of parthood itself only by offering a list of a half-dozen trivial examples of part whole. In assessing Rescher's arguments against transitivity, Simons comments that "In all these cases [i.e., Rescher's examples], and in many others like them, there seems to be a way of understanding the term 'part' so that the objection is clearly wrong, namely that sense where the one object is spatio-temporally included in the other." (Simons 1987: 107). Simons goes on to suggest that Rescher's argument depends on a sense of parthood "narrower than this basic spatio-temporal sense" and hence that Rescher's examples do not touch on "more basic senses of 'part'" (Simons 1987: 108). While I would agree that Rescher's examples are not entirely compelling, I hope the arguments in this chapter have been enough to suggest that Simons is wrong in thinking that spatio-temporal inclusion is the 'basic' sense of parthood.

the non-spatial parts I've been discussing is not that they clearly *violate* transitivity, but instead that I cannot prove that they obey it.

What about the final objection to component realism, the suggestion that such parts simply seem strange? As I said earlier, perhaps it's wrong to think that intuitions should be our guide here: many aspects of physical theory – the well-known oddities of quantum mechanics, for instance – are counter-intuitive. But I think I can do *slightly* better than that. Cartwright and Teller's examples of a body's lack of movement being decomposable into equal and opposite motions certainly seem implausible, but these aren't the sorts of decompositions found in the multi-scale case. What makes component realism seem so implausible in that example is that there is such a qualitative difference between the observed behavior – namely, rest – and the component behaviors. In the multi-scale case, though, while there *are* clear differences, it seems easier to imagine the observed behavior as the result of two or more component behaviors, at times competing with and at times complementing each another. And as I mentioned in chapter one, the physicists and mathematicians *talk* about such components as if they are real. That doesn't mean that physicists and mathematicians genuinely *think* that such components are real (perhaps it's just talk), but it does suggest that they at least don't find talk of such parts *so* strange.

Conclusion

In this chapter, I've tried to suggest how non-spatial parts can be understood in terms of the usual concepts of mereology, and how some of the arguments against the possibility of such parts can be met. I'll now turn to a related 'non-standard' concept of parthood also relevant to Kim's argument against non-reductive physicalism, and also involving multi-scale analysis. Unlike in the cases discussed in chapter three, the next examples will involve distinct *spatial* scales; decompositions on these different scales do not yield distinct parts of a system in the same sense as decompositions along different temporal scales do, but they do suggest how some 'realized' macro properties might be related to their micro-level realizers. The suggestion is that realized properties are in fact *parts* of their realizers.

V

Parthood and Realization

Introduction

The discussion of non-spatial parts in chapters three and four raises two related questions. The first concerns the wider applicability of the idea of non-spatial parts. While the neglect of this sort of parthood represents a significant oversight in the usual understanding of physicalism, we might wonder whether non-spatial parthood is limited to the sorts of physical systems amenable to multi-scale analysis, or whether concept can be applied on an even broader scale. The second question concerns the issue of causal exclusion, as discussed in chapter two. In chapter three I argued that since both 'fast' and 'slow' component systems are realized by the same 'micro-based' properties, there is no question of identifying those components with their realizer as Kim suggests we can. Since the causal exclusion principle implies that such irreducible properties cannot be causally active, we are led either to accept the epiphenomenalism of non-spatial parts, or to conclude that something is wrong with the causal exclusion principle. Since the former option is clearly unattractive, we need to investigate the latter. Some philosophers, such as Block (2003), seem content to argue that micro-reductions of the sort Kim imagines are unavailable. The conclusion Block draws from this is that the causal exclusion principle must be flawed, since if it were not, none of what we ordinarily take for causation would truly count as such. However, we might be more curious about *why* the causal exclusion principle fails and whether we can understand when properties do and don't compete as causes.

The connection between these two questions is the recent suggestion made by several authors that the concern over causal competition can be alleviated in general if we understand the realization relation as one of parthood[66]. According to this view, realized properties are *parts* of their realizers. Since this would appear a non-standard sort of parthood relation, we might wonder whether the idea of non-spatial parts and the claim of 'realization as parthood' are complementary ideas that could prove mutually illuminating.

In this chapter I'll discuss this suggestion about realization and examine its connection with the idea of non-spatial parts; in particular, I'll consider arguments from other examples of multi-scale analysis involving distinct *spatial* scales that appear to support this view of realization. I'll argue that the idea of non-spatial parts helps to understand the suggestion about realization, but that this alone doesn't license any across the board claims about realization and non-spatial parts. In keeping with the approach to parthood suggested earlier, claims about realization and parthood must be assessed on a case by case basis. I'll also suggest how such claims could be maintained in an example from outside of physics.

[66] See Shoemaker 2001, Clapp 2000, Watkins 2002, and Rueger 2004.

Realization as Parthood

The idea that realized properties might be parts of their realizers is an attractive one. Since part and whole are not identical, this would avoid a commitment to reductionism, yet it is also plausible that part and whole do not compete for causal efficacy. Hence 'realization as parthood' could allow for realized properties to be causally efficacious without accepting reductionism. Causal competition between realized and realizing property does not occur because the realized property is not distinct from its realizer, and the completeness of physics is not threatened because the realized property carries with it no *new* 'causal powers': instead, its causal contribution is part of that of its physical realizer.

This is an appealing compromise between the apparent need to reject novel causal powers characterizing inexplicably 'emergent' properties, and the common intuition that there *is* something distinct about higher-level properties that is not captured in purely reductive accounts. Consider the traditional line of argument about reduction and emergence, and in particular about the relationship between properties of wholes and those of their parts. Such disputes invariably turn on the claim that properties of wholes can't be identified with or entirely explained by those of their parts because properties of wholes are in some sense 'more' than the sum of those of their parts. This sort of claim is often put in terms of a 'causal powers' account of properties[67]. On this account, properties are individuated by the causal powers they bestow upon their instances: if two things share some property, then there is supposed to be some characteristic set of

[67] See Armstrong 1978 and Shoemaker 1980 for the classic accounts of properties in terms of causal powers.

powers or abilities those things have in virtue of their having those properties. Conversely, two things differing in their causal powers must differ with respect to the instantiation of some property. David Armstrong summarizes the relationship as follows:

> (a) The active and passive powers of particulars are determined by their properties; (b) Every property bestows some active and/or passive power upon the particulars of which it is a property. (c) A property bestows the very same causal power upon any particular of which it is a property. (d) Each different property bestows a different power on the particulars of which it is a property. (Armstrong 1978: 43-44)

While of course there are details to be filled out about just what 'causal powers' *are*, this view of properties is now widely accepted, and Kim in particular defends such a view. Claims about emergence can then be put in terms of properties of wholes bestowing 'novel' causal powers upon their instances, powers not explicable in terms of those of the parts. The suggestion of realization as parthood is that 'realized' properties are actually in a sense *less* than the sum of those of their parts. Thus, realization as parthood promises to avoid the dilemma of Kim's supervenience argument by denying an implicit premise of that argument, which Michael Watkins refers to as "Kim's dictum":

> Kim's Dictum: To be real and irreducible is to make causal contributions that are genuinely novel (Watkins 2002: 112)

If realized properties do *not* involve novel causal powers, then there is no apparent conflict with the completeness of physics. At the same time, the sorts of identities between realized and realizer posited by Kim can be avoided, since the realized property is only *part* of its realizer.

Consider a standard case of realization between a determinable property such as 'red' and a determinate of that property such as 'crimson'. The worry about causal competition in this case would be the worry that though an instance of 'crimson' might realize an instance of 'red', it isn't clear what causal role 'red' could play given the coinstantiation of 'crimson'. The suggested solution would be that though 'red' and 'crimson' are not identical properties, they are not entirely distinct, either: 'red' is *part* of 'crimson'.

Immediately, this might seem like exactly the wrong conclusion: after all, surely the 'sum' of all crimson things is *part* of the sum of all red things, so the solution seems to get the relationship backwards. But this objection misses the point: while that might be true as a claim about *particulars*, the claim is not about sums or sets or any other groupings of individuals, but is instead a claim about particular instances of *properties*: in a given instance, the property 'red' is part of the property 'crimson'.

However intuitively appealing this sort of claim is in cases where 'determinable' properties are realized by specific 'determinates', the idea of realization as parthood seems somewhat less plausible in cases involving 'structural' realization. Assuming that in such cases the realizing properties are structural or 'micro-based' properties – such as Kim's example of "the property of having two hydrogen atoms and one oxygen atom in a such-and-such bonding relationship" (Kim 1998: 84) – realization as parthood would seem to require that such realizing properties have *other* parts as well, properties which are not among the properties included in description of that property. Consider the example of *jade* being realized by *jadeite*. *Jadeite* is supposed to be a micro-based property – something like 'the property of having parts $x_1 \ldots x_{11}$ such that x_1 is *sodium*, x_2 is *aluminum*, x_3 is *iron*, x_4 and x_5 are *silicon*, $x_{6,7,8,9,10,11}$ are each *oxygen*, where all of the *x*s stand in specific bonding relations to one another'. Now – keeping in mind the discussion of structural properties and Lewis' arguments against the 'pictorial view' of structural *universals* – it seems conceivable that individual instances of the properties *sodium*, *aluminum*, *iron*, etc. could in some sense be parts of that structural property, but it seems less plausible think that *jade* is also one of the parts: 'jade' doesn't appear anywhere in the structural specification of *jadeite*. Another common example of structural realization is that of the property of *hardness* in a diamond being realized by the structural configuration of the carbon atoms constituting that diamond (as, for instance, in McMullin 1978 and Gillett 2002). To defend realization as parthood we would need to be able to say that *hardness* is part of its realizer. But while it is the properties of the parts of the diamond that realize its *hardness*, no property of the parts of the diamond – no property of the carbon atoms constituting the diamond – can plausibly have *hardness* as one of its parts[68]. It *could* be part of that decomposition in a fairly obvious sense if *hardness* were one of the properties attributed to part of the diamond on that decomposition. But this isn't what we expect in cases of structural realization, and in particular in cases of structural explanation:

[68] Or, at least, if some property of a carbon atom *does* have *hardness* as one of its parts, it can't plausibly be the *same* hardness as is instantiated by the diamond.

sometimes an individual's instantiating a particular property can be explained by one of its parts instantiating that same property, but in general this is not what we expect from a genuine explanation. For example, we *could* attribute a different structural property to the diamond, such as 'being composed of parts a, b, and c such that a, b, and c are each *hard*', but while this might be a legitimate decomposition of the diamond, it wouldn't correspond to a structural property *realizing* the hardness of the diamond.

These are the sorts of cases where realization as parthood appears to need some non-standard form of parthood. I'll discuss how useful the idea of non-spatial parts could be in this regard after reviewing current accounts of realization as parthood.

The Subset View of Realization

Accounts of realization as parthood are typically presented in terms of *sets* of causal powers (see Shoemaker 2001, Clapp 2001, Watkins 2002). The suggestion that realization is a form of parthood is thought to follow from the observation that the causal powers of 'realizing' properties often properly *include* those associated with the properties they realize: realizers are thought to bestow all of the causal powers associated with the properties they realize, and usually more besides. For instance, Sydney Shoemaker writes:

> The instantiation of the determinate entails the instantiation of the determinable, and can quite naturally be said to *include* it. It seems natural to me to say that being scarlet is in part being red. Likewise, the

instantiation of a realizer property entails, and might naturally be said to include as a part, the instantiation of the functional property realized. (Shoemaker 2001: 80-1)

In cases of multiple realization, where the same higher-level property is realized by distinct lower-level properties, the multiply realized property is thought to be characterized by the set of causal powers common to all of its various realizers, so that "the different realizer properties will differ from one another in the total sets of conditional powers they confer but will be alike in conferring the conditional powers conferred by the realized property." (Shoemaker 2001: 79) So while the various realizer properties, such as 'crimson', 'scarlet', etc., confer distinct sets of causal powers, if they in fact realize some common property, such as 'red', then this should be evident in their overlapping on the causal powers characteristic of that property.

Leonard Clapp (2001) presents the most fully developed version of Shoemaker's suggestion. Clapp presents his theory as part of an account of *disjunctive* properties and their role in arguments concerning reductionism. Most philosophers have argued that disjunction is not applicable to property formation, so that while conjunctions of predicates, such as 'red and triangular' may refer to legitimate 'conjunctive' properties, disjunctions of predicates, such as 'red *or* triangular' do not. This rejection of disjunctive properties plays a key role in the argument from multiple realizability against reduction: the suggestion is that multiply realizable properties cannot be identified with the disjunctions of their various realizers, since disjunction is not a legitimate operation for generating properties. Clapp argues against this widely accepted view, suggesting that disjunctive properties are those properties characterized by the causal powers forming the intersection of the causal powers of the properties disjoined. Multiply realized properties are thus characterized by the causal powers forming a subset of those powers characterizing each possible realizer. For example, the property *red* is plausibly multiply realizable, in that there are many different shades of color (*scarlet*, *brick red*, etc.) each of which realizes the more generic color *red*. Clapp's suggestion is that the predicate formed by the disjunction of all of these shades – 'Scarlet or Brick or …' – names a distinct, disjunctive property. Presumably, in the case of red and its realizers, those realizing properties will have some common set of causal powers, and it is the property associated with *those* causal powers that can be identified as *red*. The same sort of approach is suggested for other cases of realization: for example, mental properties can be identified with the disjunctive physical properties associated with whatever common causal powers are found across all possible realizations of those mental properties, and likewise for biological properties, chemical properties, and so on.

Clapp tests his account by considering two objections to disjunctive properties raised by David Armstrong (Armstrong 1978, 1997). These are:

(i) genuine properties should confer similarity upon their instances, but if disjunctive properties were legitimate properties then they could be shared by things lacking any significant similarity.

(ii) genuine properties are individuated by the causal powers they bestow upon their instances but if disjunctive properties were legitimate properties, then they could be shared by things differing wildly in their associated causal powers.

Armstrong illustrates both objections with the predicate 'is a raven or a writing desk': if this named a genuine ('disjunctive') property, then its instances (ravens and writing desks) would lack significant similarity and would differ considerably in their causal powers. Armstrong concludes that disjunctive predicates cannot name genuine properties.

Clapp rightly points out that while these objections may hold true for Armstrong's favorite example ('is a raven or a writing desk'), the failure of this predicate to name a genuine disjunctive property does not show that *no* property could properly be characterized as 'disjunctive'. The key failing on Armstrong's part seems to be the intuition that a disjunctive property should (like a 'conjunctive' property) be something 'more' than any of its constituent properties. We reject the idea of disjunctive properties because we expect that being told that something has the property of being 'A __ B', in addition to the property of being A, should supplement the characterization of that thing in some way. In the case of conjunctive properties ('A & B') this is certainly true, whereas in the case of disjunctive properties, it seems to be false: being told that something has the property 'A v B' seems redundant if we already know that it has property 'A'. Clapp's suggestion is that rather than being something more than their disjuncts, disjunctive properties are in a sense something *less*: being told that something has a disjunctive property 'A v B' tells us that it has some features common to property A and property B. If these properties overlap in their associated causal powers, then their disjunction is the property characterized by those overlapping powers. This doesn't guarantee that any two properties will have a disjunction, but it allows that some might. For instance, unlike ravens and writing desks, different shades of red *are* similar in a prominent way: they are all red, after all, and thus should be associated with whatever causal powers are associated with 'being red'. And recognizing that these causal powers will in general be a subset of the causal powers associated with each disjunct, we can see that disjunctive properties are still individuated by their causal powers, despite the fact that their instances may vary in their 'total' causal powers, since those associated with disjunctive property make up only some of these.

Since Clapp's account is designed to defeat the suggestion that multiply realized properties cannot be identified with the disjunctions of their realizers, his account is actually a defense of a form of reductionism. However, Clapp contends that this view of realization should be amenable to a 'non-reductive' view of physicalism, even though – as the name suggests – this is typically associated with the *denial* of reducibility. The reason Clapp gives is that while on his account 'special science' properties or predicates are identifiable with (disjunctive) physical ones, any account of reduction in these terms will have to involve "an idealization of a rather extreme sort" (Clapp 2001: 135). Since a given mental property might have an infinite number of physical realizers, "[t]he reductionist's appeal to disjunctive predicates must then be understood as a claim concerning what is *in principle possible*; it is a claim to the effect that some sort of epistemologically ideal being could reduce mentalistic predicates by exhaustive infinite disjunctions of physicalistic predicates." (Clapp 2001: 135) Accordingly, Clapp defends his claim as one of *epistemic* non-reducibility: though mental (and other) predicates are in some sense 'in principle' reducible to physical ones, this is not a sense which is practicable for *us*. This account, he argues, should be satisfactory to all concerned (i.e., "psychologists, economists, and nonreduction-minded philosophers" (Clapp 2001: 136)) in providing for the autonomy of the special sciences.

However, there are two reasons to be dissatisfied with this result. The first is that to claim that the question of reductionism is a purely epistemic one based on what is 'practically' possible seems to miss the point of most discussions of reduction. Is there really *any* practical concern about the autonomy of economics, say, from particle physics? It seems exceedingly unlikely that anyone would seriously think that a complete account of economic behavior could *ever* be given in terms of particle physics, and not just because economics involves properties such as 'being a unit of currency' that are multiply realizable. Economic entities (people, money, commodities) involve enormous numbers of constituent particles, and the resulting complexity of the purely physical descriptions alone makes it incredible to think that there is any *practical* question of whether or not we could do economics from a microphysical perspective. Instead, I take the chief purpose of the argument from multiple realizability to be to refute claims about what is possible *in principle*, rather than in practice. For instance, the argument is that even if we *could* give a molecular level account of our psychological states in terms of the states of *our* brains, we wouldn't have captured the true nature of those psychological states, since (presumably) other creatures with other types of brains could share those same states. The thought can't just be that giving other such accounts would be prohibitively complex in practice, since presumably giving such an account of our *own* brains is prohibitively complex *in practice.* Instead the thought behind the argument from multiple realizability is that there is some difficulty *in principle*, due to the *endless* variety of possible realizations of a given psychological state.

The second difficulty for Clapp's acceptance of reductionism is that it seems to remove any sort of empirical content from the debate about reduction. Surely any question of reducibility should be at least in part a matter of empirical investigation: the reduction of 'heat' to 'mean kinetic energy', for example, ought to depend in some way on the relationship between our best theories of heat and our best theories of the kinetics of gases. But it's unclear that Clapp's account *as a theory of reduction* would involve any appeal to empirical theories. Provided we think that kinetics of gases *determines* their thermodynamic properties, appealing to infinitely disjunctive predicates will allow us ('in principle') to identify bodies of gas with a particular temperature as bodies with a particular molecular structure: but surely, if this *is* a sense in which one property or predicate is reducible to another, it isn't a very interesting one. Put another way, there *are* interesting questions about the relationship between our best theories of heat and our best theories of the kinetics of gases, and it seems reasonable to call these questions of reduction. For example, as we've seen in chapter three, we might ask whether one theory reduces to another in the limit of some parameter:

that *that* relationship holds between classical mechanics and the special theory of relativity, for instance, is an important, non-trivial fact.

Another worry we might raise about accounts of realization such as Clapp's concerns the association of properties with *sets* of causal powers, and of the realization relation with the subset relation. For instance, Shoemaker defines the realization relation as: "[P]roperty X realizes property Y just in case the conditional powers bestowed by Y are a subset of the conditional powers bestowed by X" (Shoemaker 2001: 78)[69]. Parthood between properties is to follow from this relationship between the sets of their associated causal powers. But what do *subsets* of causal powers have to do with *parts* of properties? Shoemaker writes:

> The instantiation of the determinate entails the instantiation of the determinable, and can quite naturally be said to *include* it. It seems natural to me to say that being scarlet is in part being red. Likewise, the instantiation of a realizer property entails, and might naturally be said to include as a part, the instantiation of the functional property realized. (Shoemaker 2001: 80-1)

Clapp's explanation of the relevance of his account of disjunctive properties for causal competition arguments is similar:

[69] Alternatively "the conditional powers conferred by a functional property will be a proper subset of the conditional powers bestowed by whatever physical property realizes it on a particular occasion." (Shoemaker 2001: 78)

> ...multiply realized mental properties, though real and causally efficacious, are better thought of as *parts* of their physical realizors...Just as there is no causal and/or explanatory competition between instances of a whole and its parts, so there is no causal and/or explanatory competition between mental properties and instances of their physical realizors. (Clapp 2001: 133)

However, while the relationship between a property and its realizer might be "better thought of" as that between part and whole, the important claim for the causal competition argument is that this relationship *be* one of part and whole. If realized and realizing properties are not related in this way, then the supervenience argument goes through, irrespective of any overlap between the causal powers of the properties concerned. After all, the fact that realized and realizing properties have at least some of the same causal powers is what gets the supervenience argument going in the first place: merely pointing out the subset relation between them only labels the problem.

I take this point to be at the heart of John Heil's (1999) objection to the idea of realization as parthood. Heil objects that RAP doesn't truly make sufficient 'room' for both realized and realizing properties: according to Heil, if a property M is realized by a property N_1 in an individual *a* "the supposed presence of *M* in *a* appears to be entirely absorbed by the presence of N_1" (Heil 1999: 194). Heil rejects the expected rejoinder, that the subset relation between the causal powers of the realized property and those of its realizer apparently *does* allow both properties to be present in the same individual, as committing a 'category mistake':

> It is a kind of category mistake...to assimilate the possessing of properties by objects to an abstract relation like set membership. The presence of *M* in an object is, unlike set membership, a concrete feature of the world, a feature the occurrence of which...makes a causal difference. (Heil 1999: 194)

One way of remedying this would be to strengthen the claim of realization as parthood and *identify* properties with sets of causal powers and then interpret the subset relation itself as one of parthood[70]. Such an approach would simply reject Heil's claim: if properties *were* sets and the subset relation *was* simply the parthood relation, then there would be no category mistake at all in pointing out that a realizing property does not 'swallow up' (to use Heil's phrase) its realizer, since the one is a subset – and hence a part – of the other. If properties were just sets of causal powers and the subsets of a set were its parts, then realized properties would be genuine parts of their realizers – provided, that is, that realized properties *are* characterized by subsets of the causal powers characterizing their realizers.

While neither Shoemaker nor Clapp *explicitly* endorse such claims, their discussions depend upon something like these assumptions. For instance, Clapp moves quite freely between the assumption that properties can be *individuated* by

[70] This latter suggestion, which I discuss below, is defended by David Lewis in his (1991); Bunt (1985) preemptively argues for a similar view.

sets of causal powers, and the assumption that properties can be *identified* with those sets:

> It will simplify matters if, instead of speaking of properties "bestowing" causal powers, properties are simply identified with sets of causal powers. Thus, I shall sometimes speak of a property being *constituted by* a set of causal powers. (Clapp 2001: 127)

It isn't clear from the remainder of Clapp's discussion how seriously he takes this simplifying assumption, but one way to develop an account of properties and their parts would be to take it at face value: properties *are* sets of causal powers. Shoemaker's discussion makes the claim about parthood and subsets all but explicit:

> The conditional powers conferred by the instance of the determinable or functional property are a proper subset, *and in that sense a part*, of the conditional powers conferred by the instance of the more determinate property that realizes it on a particular occasion." (Shoemaker 2001: 81, emphasis added)

However, both claims are contentious. First, we might object to the identification of properties with *sets* of anything: properties might be *individuated* by sets of causal powers, but this is distinct from the claim that properties *are* sets. In fact, Shoemaker himself insists that his causal powers account of properties is not meant to be a 'reductive' one involving identity between properties and sets, or as he says 'clusters', of causal powers. His main reason for this is a worry about circularity in defining properties as sets of causal powers and at the same time allowing causal powers to be conditional upon the co-instantiation of other properties. For instance, to use Shoemaker's favorite example, the property of being 'knife-shaped' supposedly confers the causal power of 'being able to cut butter if knife-sized and made of steel'. The fact that the properties of being 'knife-sized' and 'made of steel' themselves appear in the specification of the (conditional) causal power is supposed to prevent our being able to identify a property with its causal powers: "we must make use of the notion of a property in explaining the notion of a conditional power, so there is no question here of reducing properties to some more fundamental entity" (Shoemaker 1998: 64)

I think Shoemaker's worry here is unnecessary[71]. The use we must make of properties in describing a conditional power is that conditional powers depend in some way on the presence of other properties in an individual. If properties just *are* causal powers, then a conditional causal power will be one that depends in some way on the presence of other causal powers. Perhaps there is a need here to insist that clusters of causal powers be 'grounded' in some non-conditional powers, but there is no circularity simply in supposing that a given causal power is conditional upon the presence of other causal powers.

What about the second claim needed for the subset view of realization as parthood – the claim that the subset relation is a species of the parthood relation? The main proponent of this view of set-theory and parthood is David Lewis (Lewis 1991), though Lewis didn't develop the suggestion with any concern about

[71] Michael Watkins draws the same conclusion; see Watkins 2002: 119-20.

realization or causal powers in mind. Instead, Lewis' claim comes from his attempt to develop a theory of sets based only upon the singleton relation, which relates each entity to the set containing only that entity, and classical mereology. His 'First Thesis', that subsets are *parts* of a set (Lewis 1991: 4), is defended on three grounds:

i. There are certain linguistic parallels between terms for 'subset' and terms for 'parts'. For example, the German word for 'subset' is 'Teilmenge', which means 'part-set'.

ii. There are certain formal parallels between the subset relation and the parthood relation. For example, the subset relation, like the parthood relation but unlike the membership relation, is transitive.

iii. The assumption that the parts of a set are all and only its subsets allows us to build set-theory from (mere) mereology, and set-theory is exceedingly useful.

The appeal to common language in the first argument seems slightly anachronistic as a way of doing philosophy, but lest one think Lewis himself was not serious about his evidence, he adds that the explanation for this parallel is that "subclasses just are parts of classes, we know it, we speak accordingly" (Lewis 1991: 5). Talk about ontology without honest-toil! However, such linguistic evidence is hardly decisive. For instance, Alex Oliver (1994) points out that in French, the word for set is 'ensemble', which can be literally translated as 'whole', and the word for subset 'partie', which can be translated as 'part'. But the 'whole' referred to by 'ensemble' is the whole of the *members* of the set, not its subsets: "the 'ensemble des parties' is not the original set, but the power set" (Oliver 1994: 218). And while the formal analogy between the part-whole relation and the subset-set relation is often noted, this alone seems insufficient grounds for identifying the two, especially since the formal characteristics of both relations are quite minimal.

This leaves the argument that assuming the First Thesis is fruitful, and that fruitfulness makes endorsing it a worthwhile ontological investment, since it allows us to reconstruct set-theory from mereology. But, as Potter (1993) points out, it is unclear how 'fruitful' Lewis' account actually is, given its unavoidable appeal to what Lewis himself calls the 'mysterious' singletons needed to build set-theory from mereology. A singleton is a set with only one member: an entity's singleton is the set containing only that entity. The mystery lies in understanding this form of a 'generation relation' (as Lewis calls it): how is it that along with every entity comes a distinct entity that is its singleton (and that singleton's singleton, and so on)? *That* relation cannot be explicated in terms of parthood: singletons have only two parts, an improper part (themselves) and a proper one (the empty set, which is part of every set). So a singleton is not a 'whole' with its member as 'part'. Hence the mereological view of sets seems to end up having to adopt exactly the sort of unattractively mysterious principle it set out to eliminate from set-theory. Perhaps such mysterious principles are unavoidable[72], but it's unclear that the mereological view of sets makes any

[72] Lewis argues that there is no getting around the mystery, but that the fruitfulness of set theory forces us to accept singletons; of this he is adamant:

headway at reducing the mystery. Judged on fruitfulness, it's unclear that the mereological view of sets is worth defending.

Another objection to the set/subset formulation of realization as parthood concerns its ability to accommodate cases of structural realization. Gillett (2002, 2003) argues that the views suggested by Shoemaker and Clapp are ill-suited for cases of structural realization. He distinguishes between two views of realization: the standard 'flat' view of realization and the less prominent 'layered' or 'dimensioned' view he wishes to defend. Flat views are characterized by the requirement that realized properties and their realizers be instantiated by the same individual, while layered ones allow that realizing properties might be instantiated by one or more of the constituents of the individual instantiating the property realized. More specifically, the flat view says (Gillett 2002: 317):

I. A property instance X realizes a property instance Y *only if* X and Y are instantiated in the same individual

II. A property instance X realizes a property instance Y *only if* the causal powers individuative of the instance of Y match causal powers contributed by the instance of X (and where X may contribute powers not individuative of Y).

And so I have to say, gritting my teeth, that somehow, I know not how, we do understand what it means to speak of singletons. And somehow we know that ordinary things have singletons, and singletons have singletons, and fusions of singletons sometimes have singletons. We know even that singletons comprise the predominant part of Reality. (Lewis 1991: 59)

Gillett notes that Clapp's account of RAP is presented in a way that clearly suggests a flat view of realization: if the property *scarlet* realizes the property *red*, then it is most natural to think of *scarlet* and *red* as properties instantiated by the same individual. Gillett then argues that flat views in general cannot accommodate cases of structural realization. The argument he gives is simple. Consider the hardness of a diamond being structurally realized by the properties of and relations between its constituent carbon atoms. The diamond has certain causal powers characteristic of an instance of the property *hardness* (or a particular determinate of *hardness*). The diamond's instantiating this property is explicable in terms of its structure: it is composed of a tightly bound carbon lattice that is highly resistant to deformation. Clearly, we want to say that the diamond's structure in some way realizes its *hardness*. Yet the properties that realize *hardness* are the properties of and relations between those carbon atoms (their binding relations and so on). Contrary to the flat view, these realizing properties are not properties of *the diamond* at all.

Gillett acknowledges the familiar response of defenders of the flat view: the diamond itself instantiates the structural property of having such-and-such parts characterized by such-and-such properties and relations[73]. But Gillett argues that this defense of the flat view is not viable. He offers two arguments. The first is to invoke David Lewis' arguments against the viability of structural properties, while the second is his own argument, intended to be effective even if Lewis'

[73] Oddly, though Kim has articulated precisely this view in a number of publications, while Gillett acknowledges the possibility of defending the flat view with an appeal to structural properties, he attributes the suggestion to his APA commentator rather than to Kim.

objections can be overcome. Since I've already reviewed Lewis' objections to structural universals (see chapter two), I'll concentrate on Gillett's own argument, while I'll call the 'realization regress'.

The realization regress runs as follows. Consider a diamond with the property of being hard (H). This property is realized in the diamond by a variety of properties and relations holding between its constituent atoms, namely, whatever properties and relations make the constituent atoms tightly bound to one another and difficult to displace. These properties and relations realize H. Yet since they are not properties and relations *of the diamond*, but rather its constituent parts, the property H and its realizer are instantiated by different individuals, contra (I). Gillett's main claim is that invoking structural properties only postpones the problem. Suppose we argue that the diamond's hardness is realized by a structural property of *it*, called 'COMBO', where this is something like the property "has constituent atoms a1,…,an with such and such properties and standing in such and such bonding relations". COMBO realizes H, and COMBO is a property of the diamond, so it seems well suited to a flat view of realization. However, argues Gillett, COMBO itself must also be realized, eventually by 'fundamental' properties, such as spin, of 'fundamental' entities, such as quarks. COMBO is not a property of any particular quark, but is realized by properties of quarks. So while on the structural properties view, the immediate realizer of H (that is, COMBO) is instantiated by the same individual as H (the diamond), that realizer must in turn be realized, and this time by individuals other than the diamond. Again, (I) fails, and the flat view must be rejected.

In place of the flat view, Gillett offers a 'layered' or 'dimensioned' account of realization. It says:

> Property/relation instance(s) F1-Fn realize an instance of a property G, in an individual *s, if and only if s* has causal powers that are individuative of an instance of G in virtue of the powers contributed by F1-Fn to *s* or *s*'s constituent(s), but not vice versa.

The causal powers of the realized property in this case are not necessarily a subset of those of its realizer (or realizers). Instead, the causal powers characteristic of a realized property *P* are in had 'in virtue of' the causal powers characteristic of properties of the parts of the thing instantiating *P*. On Gillett's account, rather than being subsets of their realizers, realized properties are better characterized as *resultants* of those realizers.

However, while clearly there is some sense in which the causal powers associated with a diamond's *hardness* arise in virtue of those individuative of the constituent structure of the diamond, exactly how this relationship should be characterized is difficult to assess without a detailed understanding of the relevant structure. I'll now describe an example involving a property and its micro-structural realizer that is easier to evaluate; as we'll see, this example favors the suggestion of realization as parthood over Gillett's 'resultant' view.

Macro and micro-scale properties

In chapter three, I discussed examples of multi-scale analysis involving distinct 'fast' and 'slow' temporal scales. But multi-scale analysis also often appeals to distinct spatial scales, and we can use these to examine the relationship between macro and micro-level properties of the sort Gillett and others consider.

Let's begin by clarifying the distinction between macro and micro-level properties in physics. Philosophers often talk as if macro properties include only what we might call the 'gross' properties of an individual. Typical examples include properties such as an individual's *mass* or *hardness*, as the examples we've seen from Kim (i.e., Kim 1998: 83-85) and Gillett illustrate. What makes these properties *macro* properties is that they are instantiated by relatively 'big' *things*, such as tables and diamonds. Conversely, on the standard account, a property counts as a *micro* property just in case it is characteristically a property of a relatively small thing, such a molecule or atom.

However, this 'thing' oriented view misses out on a well established distinction between macro and micro properties in physics itself. The view from physics allows for macro-level properties to provide much more detailed specifications of a given property than the standard philosophical examples suggest. Macroscopic properties can be 'structural' properties, too, pertaining to the characteristics of *parts* of individual. For example, considering a macro property such as *temperature*, we can talk about the *distribution* of temperature throughout my body, and this pattern of distribution will be a macro property in two senses. The first is the ordinary sense of it being a property of a macro-level object; namely, me. The second sense, however, is that it is a property of an individual considered on a macroscopic *scale*. The scale at which we describe an individual is a measure of the level of detail we are concerned with: a macro-scale representation ignores small-scale variations in a property, while a micro-scale description takes those variations into account. Macro and micro descriptions are thus distinguished not so much by the relative 'size' of the things they describe as by the level of detail they take into account. Macro-level properties can still be structural. For example, we might describe the temperature of a heated bar with a function that gives a particular value for each spatial position along the bar. On the macro view, fine details are 'smoothed over' and the bar is represented as a homogenous body without any atomic structure. But this macro view still tells us about the *parts* of the bar and how they vary in their temperature.

Rueger (2004) considers an example of the relationship between macro and micro-level descriptions in a model of steady-state heat conduction in a one-dimensional rod. The relevant structural properties concern the distribution of conductivity and heat throughout the rod. In this case, descriptions on different spatial scales can be distinguished, based on whether or not they take the small-scale structure of the rod into account. On the micro description, the rod is assumed to be composed of atoms separated by empty space. Conductivity on this view is discontinuous, and varies rapidly with position. On the macro view, on the other hand, the rod is assumed to be homogeneous, and conductivity is continuous, without the rapid variations characteristic of the micro description. The general equation for steady-state thermal conductivity, assuming no heat sources, is:

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

We can read this equation as expressing a relationship between structural properties characterizing the distribution of temperature (specified by *T*) and conductivity (specified by *k*) in the rod. Note that in equation (1) neither *k* nor *T* is inherently micro-structural or macro-structural. Instead, equation (1) gives us a schema for representing causal relationships between various properties: the properties themselves are represented by particular solutions to (1), and it is the nature of a particular solution that determines whether it represents a macro or a micro structural property. A particular solution for *T*, for example, will designate a micro-structural property if it takes into account the micro-level structure of the body, and a macro-structural property if it 'smoothes over' these fine details. Formally, we can distinguish between these two levels of description in terms of the way we characterize conductivity, so that a function *T(x)* will represent 'macro' temperature if it is a solution to (1a):

$$(1a) \qquad (d/dx)[k_{Macro}(x)dT(x)/dx] = 0,$$

where k_{Macro} represents conductivity in a continuous 'macroscopic' way; and a function *T(x)* will represent 'micro' temperature if it is a solution to

$$(1b) \qquad (d/dx)[k_{micro}(x)dT(x)/dx] = 0,$$

where k_{micro} represents conductivity in a discontinuous 'microscopic' way. It's important to keep in mind that while k_{Macro} and k_{micro} both represent conductivity, they are distinct functions. Each is characterized by a distinct spatial scale, which we'll call *L* (the macro scale) and *l* (the micro scale).

To investigate the relationship between the macro and micro-level description of the rod, we can begin by considering whether or not the macro-structural property T_{Macro} reduces to the micro-structural property T_{micro}, in the sense of L-reduction discussed in chapter three (which, again, is the most appropriate sense of reduction for the case at hand). Testing this involves finding a representation where the macro description 'goes over' to – that is, L-reduces to – the micro description in the limit of some parameter. For convenience, we can choose our spatial scales L and l so that $L = 1$ (the units are arbitrary) and $l = \varepsilon L$, where ε is the 'small' parameter which we will take to the limit (it should be clear that ε is a small parameter since $\varepsilon = l/L$ and l itself – the micro-scale – is small compared with L; hence ε is a measure of the relative 'size' of the micro-scale).

To test the reducibility of the macro-scale description of the heated rod to the micro-scale description, we begin by trying to find a perturbation expansion for temperature that will serve as a solution to (1):

$$T(x) = T_0(x) + \varepsilon T_1(x) + \varepsilon^2 T_2(x) + \ldots \tag{2}$$

Our goal is to find a solution where T_0 represents the *macroscopic* solution to (1), while T itself represents the micro-scale description of the system. If we can do that, the micro description will be represented as a perturbation of the solution to the macroscopic description, T_0. T_0 gives us the 'core' solution, while the remaining terms in the series represent corrections to that solution. We can then investigate how (2) behaves in the limit of $\varepsilon \rightarrow 0$. For a successful reduction, we expect the solutions for the micro-scale description to go over to solutions to the macro-scale description as $\varepsilon \rightarrow 0$ (this represents the 'homogenization limit' of the micro-scale view). However, it turns out that that this 'regular' perturbation approach fails for equation (1): micro-level solutions to (1) do not uniformly converge on macro-level solutions as the perturbation parameter $e \rightarrow 0$. Since the perturbation expansion fails to be asymptotic, we can't regard the macro description simply as an approximation of the micro description: the micro description *doesn't* go over to the macro description as our 'approximation parameter' tends to zero. So we have an apparent failure of L-reduction: the macro-description is *not* a result of the micro-description in the limit of some parameter.

As in the case of fast and slow component processes, we seem to have a counter-example to Kim's assumption about properties in the physical sciences: we *can't* always identify these with micro-based properties. However, if we investigate how successful perturbation expansions can be found in cases like the heated rod, we arrive at an interesting conclusion about the relationship between these properties.

Equations such as (1) can be solved by explicitly introducing multiple spatial scales to distinguish between the macroscopic characteristics of the rod and the microscopic ones. To introduce these new scales, we assume that the variable x represents the microscopic scale and introduce a new variable $\xi = \varepsilon x$ representing the macroscopic scale. We then perform the perturbation expansion in terms of the two variables, treating them as distinct. The perturbation expansion of T is then:

(3) $$T(x, \xi) = T_0(x, \xi) + \varepsilon T_1(x, \xi) + \varepsilon^2 T_2(x, \xi) + \ldots$$

Now we use this expansion as our solution to the original equation (1) (the heat conduction equation)[74] to get:

(4) $$\partial/\partial x[k(x)\partial T_0(x, \xi)/\partial x] = 0$$

at the lowest order of expansion (that is, we insert the perturbation expansion and then solve for the terms multiplied by the lowest order of the expansion parameter ε). This gives us the 'core' approximation of T. To improve this approximation, we examine higher orders in the expansion. To keep the second order expansion asymptotic (that is, to prevent it from diverging unacceptably from the value it is supposed to be approximating), we need to impose a further constraint on T_0:

(5) $$\partial/\partial\xi[k_{eff}\partial T_0(\xi)/\partial\xi] = 0$$

where k_{eff} is the 'effective' conductivity:

(6) $$k_{eff} = \left[\frac{1}{l}\int_0^l \frac{dx}{k_{micro}(x)}\right]^{-1}$$

The effect of this constraint is to make k_{eff} a *macro* variable: it is indifferent to the small scale variations in conductivity (k_{eff} is the harmonic mean of the microscopic conductivity). So k_{eff} is k_{Macro} mentioned in (1a), and since T_0 satisfies (5), it can be regarded as the macroscopic description of temperature, as desired (note that (5) and (1a) do differ in that (5) involves partial derivatives while (1) does not, but this is simply a consequence of our introducing distinct scales into our solution).

[74] Or, rather, we use this expansion as our solution to the multi-scale equivalent of (1), where ordinary differentiation is replaced by partial differentiation to accommodate our multiple independent scales.

What this means is that to find a solution for the micro-level description of the relationship between temperature and conductivity in the rod, we've had to use a perturbation expansion, and to keep that expansion asymptotic (and thus to keep it a good approximation), we've had to posit two independent scales on which the properties we're interested in – temperature and conductivity – can vary, *and* we've had to constrain the first component of that expansion so that it represents the *macro* description of the relationship between temperature and conductivity in the rod. Rueger observes:

> The important point in this calculation is that the 'derivation' of [the macro-structural description] from the micro-structural description inevitably involves quantities at *both* length scales. You can't go in the limit from the micro-structural description…to the macro-structural description without formally introducing two independent spatial scales at which the quantities change. Once these scales have been introduced, the macro-structural description itself has to be invoked to ensure that the perturbation expansion of the micro-structural description remains asymptotic." (Rueger 2004: 9)

Now consider how this example relates to our interest in realization. We've just seen that in order to solve an equation such as (1) on the micro scale, we have to resort to methods that involve presuming that the solution is a perturbation of an idealized or macroscopic description of the same phenomenon. This suggests that the macro-structural property is in fact a *component* of its micro-structural realizer, just as the slow and fast sub-systems were components of the observed systems described in chapter three: in other words, macro-structural properties are *parts* of micro-structural properties.

This might seem like an odd conclusion: how can a macro property be part of a micro property? However, it is not that properties of a macro individual are parts of properties of some micro individual. Instead, a given individual (a rod, in the above case) has two properties: one is a micro-level property describing the fine details of its temperature distribution (for example), while the other is a macro-level property also describing its temperature description, but in a way that is indifferent to fine detail. These two properties are not identical, but they are not entirely distinct, either: one is part of the other. In this case, at least, 'realized' properties do not reduce to their realizers; instead, they are parts of them.

Realization and Non-Spatial Parts

What role could non-spatial parts play in understanding realization as parthood? In the case of macro and micro-scale descriptions, it isn't appropriate to regard the descriptions as applying to distinct entities: in this case, both are properties of the same entity, namely the rod. That's distinct from the case of slow and fast-scale decomposition, where the behavior of one observed process was treated as the result of the combined behavior of two distinct component processes.

However, in the macro/micro case the parthood claim isn't *so* different: in both cases, understanding the parthood relation fundamentally involves understanding the relation between different properties. In the same sense as in the slow/fast decompositions, the macro and micro-structural properties of the rod represent (some of) the same 'portion of reality', though in this case that portion is exclusively a portion of the properties characterizing 'reality'.

This connection between macro and micro-structural properties suggests how we could respond to Gillett's arguments against realization as parthood and the 'flat' view of realization it presumes. Rather than viewing the 'resultant' properties of particular structural configurations as the properties *realized* by that structure, we could treat Gillett's 'layered view' as an account of structural properties themselves: having a particular structure brings with it particular causal powers – the causal powers a thing has 'in virtue of' the causal powers of its structural constituents – and clearly these causal powers are what characterize that particular structural property[75]. However, rather than thinking that the causal powers characterizing realized properties are themselves the resultants of the causal powers of their structural realizers, the perturbation examples suggest that realized properties are components of *other* decompositions of those resultants. Realization, then, can be viewed simply as another variety of decomposition: just as resultant structural properties can be decomposed into the properties of their structure, they can also be decomposed into the properties they realize.

On this view the letter of the definition of the layered view is respected: an individual *s* has causal powers individuative of an instance of a (realized) property G in virtue of the powers contributed by the properties characterizing *s* and its

[75] I say 'characterize' and not 'individuate' since perhaps other structures could result in the same causal powers. The structural property itself is still to be understood as the property of having such and such parts with such and such properties and so on, and *having* such and such parts with such and such properties confers particular causal powers upon an individual.

parts. But the spirit is distinct: Gillett's view suggests that it is *realized* properties that are the resultant properties of the various properties of an individual's parts, whereas the suggestion here maintains the central claim from RAP that realized properties will (at least in general) only be *parts* of those resultants. That Gillett intends this account of realization is clear from a related discussion of Shapiro (2001) against the very idea of 'multiple' realizability. Shapiro presents the following dilemma for the defender of multiple realizability: either different realizer properties differ with respect to their "causally relevant properties" or they do not. If they *do* differ, then we have causally distinct properties, then there is no common 'realized' property, and we do not have a case of multiple realization after all. And if they do *not* differ – that is, if the different realizer properties do *not* differ with respect to their causally relevant properties – then these realizers are not truly distinct, and so we do not have a case of genuine *multiple* realization. Gillett argues that his layered view can make perfect sense of causally distinct realizers realizing a causally homogeneous properties: since the layered view is based on the structural features of constituents, and "although [the] properties/relations of constituents differ in the powers that they contribute, we nonetheless have distinct realizations of [a given property] because these different properties/relations both *result in* the causal powers individuative of *T*." (Gillett 2003, p. 601, emphasis added). In something like the sense Block suggested in chapter three's discussion of the supervenience argument (though for a different purpose), Gillett's suggestion seems to be that multiple realized properties can be 'differently decomposed'.

Consider Gillett's diamond example again. The diamond's property *hardness* is to be 'realized' by its having a structural property, and Gillett's regress arose once we started asking about the realizer of this structural property, and then *that* property's realizer, and so on. But once we distinguish between the property of having a particular structure and the causal powers had due to that structure, the problem becomes much less apparent. The diamond has a structural property in the sense that it has a decomposition into carbon atoms characterized by the description Gillett called 'COMBO': i.e., "has constituent atoms a1,…,an with such and such properties and standing in such and such relations". Simplify this to ease exposition, and take COMBO to be 'has constituents atoms a1 and a2 such that both a1 and a2 are *carbon* and are related by being *bonded*'. Having such a decomposition endows the diamond with certain causal powers: *these* characterize the 'resultant' property we can call '$COMBO_R$'. *That's* the property 'of the diamond', whereas COMBO is better thought of as a relation between the diamond and a particular group of individuals constituting it. And it's that *resultant* property that 'realizes' *hardness*: an individual diamond instantiating *hardness* has a particular structure (a particular decomposition – into a lattice of tightly bound carbon atoms); having that structure endows the diamond with a particular causal powers which individuate a property of the individual; the realized property is then part of *that* property. Since realizer and realized properties are properties of the same individual, this would still be a 'flat' view of realization. So the flat view can not only accommodate structural realization, but structural realization is consistent with realization as parthood after all.

As for the realization regress, we might wonder why the property COMBO needs to be 'realized' at all? COMBO characterizes a particular decomposition of the diamond, into carbon atoms with particular properties and so on. It may well be that *these* properties – properties such as *being carbon* and *being bonded* – are themselves realized by other properties (such as "having a nucleus with however many protons, etc.") and *these* properties may be realized in turn by other properties ('having a particular quark structure'). But at each stage the realized and realizing properties are borne by the same individual. Understood as decompositions, structural properties don't give rise to a regress.

So while the idea of realization as parthood doesn't involve non-spatial parts in exactly the same sense as discussed in chapters three and four, the idea of viewing the parts of properties as *components* is helpful for understanding realization as parthood.

Whether or not the idea of realization as parthood can be defended as a general view of realization remains to be seen. In particular, the same sorts of considerations about component realism discussed in chapter four should apply here, so that in order to make a genuine claim of realization as parthood in any particular case we would need to have a specific account of how the 'realized' and 'realizing' properties are connected, and we would need to have reason to think that that account involved a suitably 'natural' decomposition. This sort of relationship is satisfied in the case of the two structural descriptions of the heated rod, where we have a clear account of the relationship between the two properties, and where the multi-scale decomposition only worked once we constrained the first term in the perturbation expansion to conform to the *macro*-level law relating temperature and conductivity. In examples like Gillett's, justification for the claim of realization as parthood is more sketchy (but no more sketchy than his own account of the 'layered' view of realization); without a specific account of properties such as *hardness* (in the form of some laws governing hardness – or perhaps more general laws relating macro and micro-level rigidity or deformability), it's difficult to *show* that the claim of realization as parthood is justified in this case.

To illustrate how this idea of realization might be given a more formal treatment outside of physics, and how the ideas of non-spatial parts and realization as parthood might be drawn more closely together, consider the following suggestion about economic properties made by James Ramsey (1996). Ramsey argues for the importance of multi-scale decompositions for understanding macro-level properties in economics[76]. Though the mathematical details are somewhat different in Ramsey's discussion, they illustrate the same general principles as discussed here. The issue of interest for Ramsey is what is to count as a macro property in economics and whether or not such properties even exist. He rejects the suggestion that macro-economic properties merely represent averages over the true micro-level properties describing individual economic agents. According to *this* view, macro properties are simply *aggregates* or sums of micro-economic properties, and consequently they can be represented as the micro-level properties of some 'aggregate' or 'representative' individual.

[76] Note that I'm interpreting Ramsey's use of the term 'variable' to refer to properties.

Ramsey argues that macro-level properties are to be understood as genuinely 'system-wide' properties that cannot necessarily be represented in terms of properties of individual economic agents. Instead, he identifies *macro*-economic properties with particular components of the solutions to the equations governing the economics of a system. His central contention is that the appropriate definition of a 'macroeconomic' property is one characterized by the *slow*-scale component of a perturbation expansion: this slowly varying component represents the salient features of a system stripped of fast-scale 'micro' variations due to individual behavior. Discovering such components is an empirical matter: since they are not simply the sums or averages of properties characterizing individuals, there is no guarantee that there *are* such macro-level properties bearing stable relationships to one another.

This suggestion shares obvious affinities with the ideas suggested here. On the one hand, the behavior of an economic system is decomposed into distinct component systems operating on different temporal scales: those component systems are 'non-spatial' parts of the composite system. On the other hand, the macro economic properties are distinct properties in their own right, yet they are still 'realized' by the system that results from the combination of the slow-scale component and the fast-scale component: so in this case again macro properties are treated as parts of their realizers. In this case at least, the idea of non-spatial parts and realization as parthood appear to coincide.

Conclusion

I'll end by drawing some general conclusions about the ideas of parthood and realization discussed here and throughout the thesis and by suggesting some areas for further work.

The examples from perturbation theory give us one model for understanding realization, and the remaining task is to show that properties of interest conform to that model. This isn't necessarily easy or trivial: the question of which properties realize which other properties (or which *structural* properties realize which other properties) is no easier to answer nor any less 'empirical' than the question of which properties *reduce* to which others. Perhaps this model is not applicable to all interesting cases of realization. It is difficult, for example, to imagine how *mental* properties could be incorporated into it: understanding composition relations involving mental properties does not seem nearly as conceivable as does understanding the composition of physical (or even economic) properties. Perhaps some other account of realization is needed for mental properties: as I'm sure has been noted elsewhere, perhaps 'realization' is itself 'multiply realized'. But realization relations between physical properties, like relations of reduction between physical properties, are interesting enough in their own right to warrant attention even if they have no obvious bearing on worries about mental properties. Even if it was a worry about the nature of mental properties that generated the interest in realization in the first place, the nature of the realization relation between physical properties deserves our attention as much as that between mental and physical properties.

Another challenge for this view of realization involves the sorts of extrinsic properties I've neglected throughout this thesis. Robert Wilson (2001) argues that in many cases realization is 'wide' in the sense that the realizing property is instantiated in a system that properly includes the system or individual instantiating the realized property. For instance, a biological property of an organism such as *fitness* depends not only on the intrinsic features of that organism, but a variety of environmental factors as well. Thus 'realizing' a particular fitness – though that property is in fact a property *of the organism* – involves properties and relations that extend beyond the organism. This contradicts the standard view of realization (the sort of 'flat' view discussed earlier) according to which realized and realizing properties are properties of the same individual, though it contradicts it in the opposite way to that suggested by Gillett: realizing properties are properties of systems that are even broader than the individual bearing the realized property, rather than being properties of the parts of that individual, as Gillett claimed.

There seems to be some affinity between this 'wide' view of realization and the suggestion of realization as parthood: in Wilson's case, the realized property is a property of part of the system bearing the realizing property. However, developing that suggestion would require a more detailed study of how the properties themselves, rather than their 'bearers' are related. If a connection between the two views of realization could be made, then this could significantly widen the applicability of the idea of realization as parthood.

The general conclusion I want to draw from the work presented here is that like many other concepts in philosophy, concepts such as parthood and composition need to be *naturalized* in the sense of being brought out of the realm of pure thought and into closer contact with actual accounts of entities, processes, and properties found in science. Doing this can illuminate not only our understanding of other important concepts such as realization, but also our basic understanding of the variety of properties and structures there are in the world and the roles these play in our understanding of natural phenomena. Like Vemulapalli and Byerly (1999: 18), I wish to take to heart Abner Shimony's injunction to philosophers of science to pay attention to the "the great variety and subtlety, and often surprising nature, of the derivations of properties of composite systems from those of the components" (Shimony 1987: 401). I hope the preceding study has made a plausible case for some of that surprising variety not previously noticed by philosophers.

Bibliography

ARMSTRONG, D. M. (1978) Universals and Scientific Realism, Volume 2: A Theory of Universals, Cambridge, Cambridge University Press.

ARMSTRONG, D. M. (1997) *A World of States of Affairs,* Cambridge, Cambridge University Press.

BAKER, L. R. (1993) Metaphysics and Mental Causation. IN HEIL, J. & MELE, A. (Eds.) *Mental Causation.* Oxford, Clarendon.

BATTERMAN, R. W. (1995) Theories Between Theories. *Synthese,* 103, 171-201.

BIGELOW, J. & PARGETTER, R. (1989) A Theory of Structural Universals. *Australasian Journal of Philosophy,* 67, 1-11.

BLOCK, N. (1997) Anti-Reductionism Slaps Back. *Philosophical Perspectives,* 11, 107-132.

BLOCK, N. (2003) Do Causal Powers Drain Away? *Philosophy and Phenomenological Research,* 67, 133-150.

BOLTZMANN, L. (1974) On the Indispensability of Atomism in Natural Science. IN MCGUINNESS, B. (Ed.) *Ludwig Boltzmann: Theoretical Physics and Philosophical Problems.* Boston, Massachusetts, D. Reidel.

BONTLY, T. D. (2002) The Supervenience Argument Generalizes. *Philosophical Studies,* 109, 75-96.

BRACKBILL, J. U. & COHEN, B. I. (1985) *Multiple Time Scales*, Orlando, Academic Press.

BRITTAN JR., G. G. (1970) Explanation and Reduction. *The Journal of Philosophy,* 67, 446-457.

BUNT, H. (1985) *Mass Terms and Model-Theoretic Semantics,* Cambridge, Cambridge University Press.

BURGE, T. (1993) Mind-Body Causation and Explanatory Practice. IN HEIL, J. & MELE, A. (Eds.) *Mental Causation.* Oxford, Clarendon.

CARNAP, R. (1936) Testability and Meaning. *Philosophy of Science,* 3, 419-471.

CARTWRIGHT, N. (1980) Do the Laws of Physics State the Facts? *Pacific Philosophical Quarterly,* 61, 75-84.

CARTWRIGHT, N. (1983) *How the Laws of Physics Lie,* Oxford, Clarendon Press.

CHOMSKY, N. (1987) *Language and Problems of Knowledge,* Boston, Massachusetts, MIT Press.

CLAPP, L. (2001) Disjunctive Properties: Multiple Realizations. *The Journal of Philosophy,* 98, 111-136.

CRANE, T. & MELLOR, D. (1990) There is No Question of Physicalism. *Mind,* 99, 185-206.

CREARY, L. (1981) Causal Explanation and the Reality of Natural Component Forces. *Pacific Philosophical Quarterly,* 62, 148-157.

CROOK, S. & GILLETT, C. (2001) Why Physics Alone Cannot Define the 'Physical': Materialism, Metaphysics, and the Formulation of Physicalism. *Canadian Journal of Philosophy,* 31, 333-360.

DEHMELT, H. (1989) Triton,...Electron,...Cosmon,...:An Infinite Regression?

Proceedings of the National Academy of Science, 86, 8618-8619.

EDWARDS, D. A. (2000) An Alternative Example of the Method of Multiple Scales. *SIAM Review,* 42, 317-332.

FEIGL, H. (1958) The 'Mental' and the 'Physical'. IN FEIGL, H., SCRIVEN, M. & MAXWELL, G. (Eds.) *Minnesota Studies in the Philosophy of Science* Minneapolis, University of Minnesota Press.

FIELD, H. (1992) Physicalism. IN EARMAN, J. (Ed.) *Inference, Explanation, and Other Frustrations: Essays in the Philosophy of Science.* Los Angeles, California, University of California Press.

FODOR, J. A. (1974) Special Sciences, or the Disunity of Science as a Working Hypothesis. *Synthese,* 28, 97-115.

GILLETT, C. (2002) The Dimensions of Realization: A Critique of the Standard View. *Analysis,* 62, 316-323.

GILLETT, C. (2003) The Metaphysics of Realization, Multiple Realizability, and the Special Sciences. *The Journal of Philosophy,* 100, 591-603.

GILLETT, C. & LOEWER, B. (Eds.) (2001) *Physicalism and Its Discontents,* Cambridge, Cambridge University Press.

GILLETT, C. & WITMER, D. G. (2001) A 'Physical' Need: Physicalism and the *Via Negativa. Analysis,* 61, 302-309.

GIRILL, T. R. (1976) Criteria for the Part-Whole Relation in Micro-Reductions. *Philosophia,* 6, 69-79.

GIRILL, T. R. (1976) Evaluating Micro-Explanations. *Erkenntnis,* 10, 387-405.

HEIL, J. (1999) Multiple Realizability. *American Philosophical Quarterly,* 36, 189-208.

HEMPEL, C. G. (1966) *The Philosophy of Natural Science,* Englewood Cliffs, Prentice Hall.

HEMPEL, C. G. (1980) Comments on Goodman's *Ways of Worldmaking. Synthese,* 45, 193-199.

HENDRY, R. F. (1999) Molecular Models and the Question of Physicalism. *Hyle,* 5, 117-134.

HESSE, M. (1953) Models in Physics. The British Journal for the Philosophy of Science, 4, 198-214.

HINCH, E. J. (1990) *Perturbation Methods,* Cambridge, Cambridge University Press.

HOLMES, M. (1999) The Method of Multiple Scales. *Proceedings of Symposia in Applied Mathematics,* 56, 23-46.

HORNUNG, U. (1997) *Homogenization and Porous Media,* New York, Springer.

HUMPHREYS, P. (2000) Extending Ourselves. IN CARRIER, M., MASSEY, G. J. & RUETSCHE, L. (Eds.) *Science at Century's End: Philosophical Questions on the Progress and Limits of Science.* Pittsburgh, University of Pittsburgh Press.

HUSSERL, E. (1970) *Logical Investigations, Volume 2,* New York, Humanities Press.

JONES, C. K. R. T & KHIBNIK, A. I. (2001) *Multiple-Time-Scale Dynamical Systems*, New York, Springer.

KEMENY, J. G. & OPPENHEIM, P. (1956) On Reduction. *Philosophical Studies,* 7, 6-19.

KIM, J. (1971) Materialism and the Criteria of the Mental. *Synthese,* 22, 323-345.

KIM, J. (1989) The Myth of Nonreductive Physicalism. Proceedings and Addresses of the American Philosophical Association, 63, 31-47.

KIM, J. (1997) Does the Problem of Mental Causation Generalize? *Proceedings of the Aristotelian Society,* 97, 281-297.

KIM, J. (1998) *Mind in a Physical World,* Boston, MIT Press.

KIM, J. (1999) Supervenient Properties and Micro-based Properties: a Reply to Noordhof. *Proceedings of the Aristotelian Society,* 99, 115-118.

KIM, J. (2003) Blocking Causal Drainage and Other Maintenance Chores with Mental Causation. *Philosophy and Phenomenological Research,* 67, 151-176.

LANGE, M. (1994) Scientific Realism and Components: The Case of Classical Astronomy. *The Monist,* 77, 111-127.

LANGE, M. (1996) Laws of Nature, Cosmic Coincidences and Scientific Realism. *Australasian Journal of Philosophy,* 74, 614-637.

LEONARD, H. S. & GOODMAN, N. (1940) The Calculus of Individuals and Its Uses. *The Journal of Symbolic Logic,* 5, 45-55.

LESNIEWSKI, S. (1992 [1916]) Foundations of the General Theory of Sets. I. IN SURMA, S. J., SRZEDNICKI, J., BARNETT, D. I. & RICKEY, F. V. (Eds.) *Collected Works / Stanislaw Lesniewski.* Dordrecht, Kluwer.

LEWIS, D. (1970) Generalized Semantics. *Synthese,* 22, 18-67.

LEWIS, D. (1986a) *Philosophical Papers,* New York, Oxford University Press.

LEWIS, D. (1986b) Against Structural Universals. *The Australasian Journal of Philosophy,* 64, 25-46.

LEWIS, D. (1991) *Parts of Classes,* Cambridge, Massachusetts, Basil Blackwell.

LEWIS, D. (1994) Reduction of Mind. IN GUTTENPLAN, S. (Ed.) *A Companion to Philosophy of Mind.* Basil Blackwell.

LOEWER, B. (2001) From Physics to Physicalism. IN GILLETT, C. & LOEWER, B. (Eds.) *Physicalism and Its Discontents.* Cambridge, Cambridge University Press.

MARKOSIAN, N. (2000) What Are Physical Objects? *Philosophy and Phenomenological Research,* 61, 375-396.

MCLAUGHLIN, B. (1992) The Rise and Fall of British Emergentism. IN BECKERMANN, A., FLOHR, H. & KIM, J. (Eds.) *Emergence or Reduction?* Cambridge, Cambridge University Press.

MCMULLIN, E. (1978) Structural Explanation. *American Philosophical Quarterly,* 15, 139-147.

MEIRAV, A. (2000) A Mereological Criterion for Physicality. *The Southern Journal of Philosophy,* 38, 619-631.

MELNYK, A. (1994) Being a Physicalist: How and (More Importantly) Why. *Philosophical Studies,* 74, 221-241.

MELNYK, A. (1997) How to Keep the 'Physical' in Physicalism. *The Journal of Philosophy,* 94, 622-637.

MILL, J. S. (1967) *A System of Logic,* London, Longmans, Green and Co Ltd.

NAGEL, E. (1961) *The Structure of Science,* New York, Harcourt, Brace.

NAYFEH, A. (1973) *Perturbation Methods,* New York, Wiley.

NEEDHAM, P. (1996) Macroscopic Objects: An Exercise in Duhemian Ontology. *Philosophy of Science,* 63, 205-224.

NEEDHAM, P. (1999) Macroscopic Processes. *Philosophy of Science,* 66, 310-

331.

NEEDHAM, P. (2000) What is Water? *Analysis,* 60, 13-21.

NEURATH, O. (1973 [1931]) Empirical Sociology. IN NEURATH, M. & COHEN, R. S. (Eds.) *Otto Neurath: Empiricism and Sociology.* Dordrecht, D. Reidel.

NEURATH, O. (1983 [1931]) Physicalism. IN COHEN, R. S. & NEURATH, M. (Eds.) *Otto Neurath: Philosophical Papers 1913-1946.* Dordrecht, Kluwer.

NEWTON, I. (1973) *Mathematical Principles of Natural Philosophy,* Berkeley, University of California Press.

NICKLES, T. (1973) Two Concepts of Intertheoretic Reduction. *The Journal of Philosophy,* 70, 181-201.

NOORDHOF, P. (1999) Micro-based Properties and the Supervenience Argument: A Response to Kim. *Proceedings of the Aristotelian Society,* 99, 109-114.

OLIVER, A. (1994) Are Subsets Parts of Classes? *Analysis,* 54, 215-223.

OPPENHEIM, P. & PUTNAM, H. (1958) Unity of Science as a Working Hypothesis. IN FEIGL, H., SCRIVEN, M. & MAXWELL, G. (Eds.) *Minnesota Studies in the Philosophy of Science.* Minneapolis, University of Minnesota Press.

PAPINEAU, D. (1993) *Philosophical Naturalism,* Oxford, Blackwell.

PAPINEAU, D. (2001) The Rise of Physicalism. IN GILLETT, C. & LOEWER, B. (Eds.) *Physicalism and Its Discontents.* Cambridge, Cambridge University Press.

PETTIT, P. (1993) A Definition of Physicalism. *Analysis,* 53, 213-223.

PETTIT, P. (1995) Causality at Higher Levels. IN SPERBER, D., PREMACK, D. & PREMACK, A. J. (Eds.) *Causal Cognition: A Multi-disciplinary Debate.* Oxford, Clarendon Press.

PLANTINGA, A. (1996) Respondeo ad van Fraassen. IN KVANVIG, J. L. (Ed.) Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge. Lanham, Rowan and Littlefield.

POLAND, J. (1994) Physicalism: the Philosophical Foundations, Oxford, Clarendon.

POTTER, M. (1993) Critical Notice: *Parts of Classes*, by David Lewis. *The Philosophical Quarterly,* 43, 362-366.

PUTNAM, H. (1975) Minds and Machines. *Mind, Language, and Reality.* Cambridge, Cambridge University Press.

PUTNAM, H. (1975) The Nature of Mental States. *Mind, Language, Reality.* Cambridge, Cambridge University Press.

RAMSEY, J. (1996) On the Existence of Macro Variables and of Macro Relationships. *Journal of Economic Behavior and Organization,* 30, 275-300.

RESCHER, N. (1955) Axioms for the Part Relation. *Philosophical Studies,* 6, 8-11.

RESCHER, N. (2000) *Process Philosophy: A Survey of Basic Issues,* Pittsburgh, University of Pittsburgh Press.

ROSEN, G. & DORR, C. (2002) Composition as a Fiction. IN GALE, R. M. (Ed.) *Blackwell Guide to Metaphysics.* Oxford, Blackwell.

RUBEN, D. H. (1983) Social Parts and Wholes. *Mind,* 92, 219-238.

RUEGER, A. (2000) Physical Emergence, Diachronic and Synchronic. *Synthese,* 124, 297-322.

RUEGER, A. (2004) Reduction, Autonomy, and Causal Exclusion Among Physical Properties. *Synthese,* 139, 1-21.

RUSSELL, B. (1903) *The Principles of Mathematics,* Cambridge, Cambridge University Press.

SABATES, M. (2001) Varieties of Exclusion. *Theoria,* 16, 13-42.

SCHLESINGER, G. (1961) The Prejudice of Micro-Reduction. *The British Journal for the Philosophy of Science,* 12, 215-224.

SHAPIRO, L. (2000) Multiple Realizations. *The Journal of Philosophy,* 97, 635-654.

SHELDON, N. A. (1985) One Wave or Three? A Problem for Realism. *British Journal for the Philosophy of Science,* 36, 431-436.

SHIMONY, A. (1987) The Methodology of Synthesis: Parts and Wholes in Low-Energy Physics. IN KARGON, R. & ACHINSTEIN, P. (Eds.) *Kelvin's Baltimore Lectures and Modern Theoretical Physics: Historical and Philosophical Perspectives.* Cambridge, MIT Press.

SHOEMAKER, S. (1980) Causality and Properties. IN VAN INWAGEN, P. (Ed.) *Time and Cause.* Dordrecht, Reidel.

SHOEMAKER, S. (1998) Causal and Metaphysical Necessity. *Pacific Philosophical Quarterly,* 79, 59-77.

SHOEMAKER, S. (2001) Realization and Mental Causation. IN GILLETT, C. & LOEWER, B. (Eds.) *Physicalism and Its Discontents.* Cambridge, Cambridge University Press.

SIMONS, P. (1987) *Parts: A Study in Ontology,* Oxford, Clarendon Press.

SKLAR, L. (2000) *Theory and Truth: Philosophical Critique Within Foundational Science*, Oxford, Oxford University Press.

SMART, J. J. C. (1959) Sensations and Brain Processes. *Philosophical Review*, 68, 141-156.

SMART, J. J. C. (1981) Physicalism and Emergence. *Neuroscience,* 6, 109-113.

SMITH, D. R. (1985) *Singular-Perturbation Theory: An Introduction with Applications*, Cambridge, Cambridge University Press.

SMITH, P. (1992) Modest Reductions and the Unity of Science. IN CHARLES, D. & LENNON, K. (Eds.) *Reduction, Explanation, and Realism.* Oxford, Clarendon Press.

SPURRETT, D. (2001) Cartwright on Laws and Composition. *International Studies in the Philosophy of Science,* 15, 253-268.

STURGEON, S. (1998) Physicalism and Overdetermination. *Mind,* 107, 411-432.

TELLER, P. (1995) *An Interpretive Introduction to Quantum Field Theory,* Princeton, New Jersey, Princeton University Press.

THOMSON, J. J. (1977) *Acts And Other Events,* Ithaca, New York, Cornell University Press.

VAN DYKE, M. (1975) *Perturbation Methods in Fluid Mechanics,* Stanford, California, The Parabolic Press.

VAN FRAASSEN, B. C. (1996) Science, Materialism, and False Consciousness. IN KVANVIG, J. L. (Ed.) *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge.* Lanham, Rowan and Littlefield.

VAN GULICK, R. (1992) Three Bad Arguments for Intentional Property Epiphenomenalism. *Erkenntnis,* 36, 311-332.

VAN INWAGEN, P. (1987) When Are Objects Parts? IN TOMBERLIN, J. E. (Ed.) *Philosophical Perspectives, volume 1: Metaphysics.* Ataschadero, California, Ridgeview.

VAN INWAGEN, P. (1994) Composition as Identity. IN TOMBERLIN, J. E. (Ed.) *Philosophical Perspectives, volume 8: Logic and Language.* Atascadero, California, Ridgeview.

VEMULAPALLI, G. K. & BYERLY, H. (1999) Remnants of Reductionism. *Foundations of Chemistry,* 1, 17-41.

WATKINS, M. (2002) *Rediscovering Colors: A Study in Pollyanna Realism,* Dordrecht, Kluwer.

WEINBERG, S. (1987) Towards the Final Laws of Physics. IN *Elementary Particles and the Laws of Physics : The 1986 Dirac Memorial Lectures,* Cambridge, Cambridge University Press.

WILSON, R. A. (2001) Two Views of Realization. *Philosophical Studies,* 104, 1-31.

WIMSATT, W. (1997) Aggregativity:Reductive Heuristics for Finding Emergence. *Philosophy of Science,* 64, S372-S384.