Levels of Organisation in Scientific Practice

By

Matthew Baxendale

Submitted to

Central European University

Department of Philosophy

In partial fulfilment of the requirements for the degree of

Doctor of Philosophy

in Philosophy

Supervisor: Maria Kronfeldner

Budapest, Hungary

2018

I hereby declare that the dissertation contains no material accepted for the completion of any other degrees in any other institution and no materials previously written and/or published by anther person unless appropriate acknowledgement is made in the form of bibliographical reference.

Budapest, 13th April 2018. Matthew Baxendale

Abstract

Until recently, in philosophy, the concept of ‘levels of organisation’ was synonymous with the Layer Cake Model (LCM), (Oppenheim and Putnam, 1958). The LCM posited a global account of levels according to which every object of scientific inquiry could be placed into a single hierarchical structure; with the organisation of scientific inquiry reflecting the very same hierarchical structure. Whilst this is arguably a well-known organisational structure for scientific inquiry and the world it investigates, it has long since been widely rejected as playing any significant role in contemporary scientific practice or reflecting the organisation of our scientific inquiries (Brooks, 2016; Feibleman, 1954; Guttman, 1976; Potochnik and McGill, 2012; Waters, 2008). A central theme within discussions of levels of organisation has concerned whether the concept has outlived its utility both in scientific practice and the philosophical analysis of it (DiFrisco, 2016; Eronen, 2013, 2015), or whether the concept can still serve to illustrate an important aspect of scientific practice in a very different form to the LCM (Brooks, 2017; Craver, 2015; Kaiser, 2015).

In this dissertation I provide a new pluralistic analysis of levels of organisation with the aim of demonstrating the concept’s centrality to many forms of scientific practice and its role as a fruitful component in analyses of those practices. My account proposes that levels of organisation are collections of purpose-relative vertical and horizontal principles deployed as representative tools in the conceptualisation of a target system of inquiry. Further, I argue that, understood this way, the concept of levels of organisation serves as a lens through which to analyse key epistemic practices, specifically the manipulation, integration, and transfer of information of, and between, systems of inquiry. In the process of developing and defending my account, I examine the use of levels across a wide-range of cutting-edge research in the life sciences, including the Human Microbiome Project, cancer research, homeostatic processes of the human digestive system, and protein signalling networks. I argue that the plausibility of my account is evidenced in these diverse research fields and also demonstrate how levels of organisation can serve as a conceptual tool with which to illuminate key aspects of the dynamic and complex processes of knowledge production in the sciences.

Acknowledgements

I’d like to thank Maria Kronfeldner for all the time and effort put into my work over the last four years, not just the dissertation itself, but also my development in many aspects of academic life. I count myself exceptionally lucky to have crossed paths with Maria at CEU, and I am extremely grateful to have had the opportunity to develop my work under her watchful eye.

A portion of the ideas in this manuscript came together during my Ernst Mach Fellowship at the University of Salzburg. To that end I’d like to that Lena Zuchowski for stimulating discussions about complex systems, to Charlotte Werndl for supporting my application and acting as my academic supervisor during my time there, and to the OeAD for the award of the fellowship.

Thanks to Rachael & Robert Lee for providing a warm and comforting home from home anytime I went back to the UK, and to David & Jenne McDonald for doing the same whenever I spent time back in Ireland. Thanks to my parents, David & Alison Baxendale for their love, support, and encouragement.

I’ve met some wonderful people over the past four years to whom I owe a massive debt of gratitude for keeping me afloat. Thanks to Maria Temmes for being my writing-up companion over the last 6 months. To Frank Karioris for his friendship, particularly during my early years in Budapest. To Adina Covaci for beers on Castle Hill and everything since. Thanks to Garrett Mindt and Michele Luchetti for their encouragement, friendship, and support over the past three years. Finally, thanks to Anna Kocsis for taking care of the rest. Additional thanks to Anna and Michele for reading final drafts of the manuscript and providing thought-provoking comments and suggestions.

Some of the material in this dissertation has been published independently. The main ideas contained with Chapter Three are published as Baxendale, M. (2018). ‘Mapping the Continuum of Research Strategies’, *Synthese*, doi: 10.1007/s11229-018-1683-1. Thanks to the two anonymous reviewers for their detailed and constructive criticism, reflections on which have made it into this dissertation.

Table of Contents

List of Figures

General Introduction

> "levels of organization are a deep, non-arbitrary, and extremely important feature of the ontological architecture of our natural world, and almost certainly of any world that could produce, and be inhabited or understood by, intelligent beings" (Wimsatt, 2007, p. 203).

> "These considerations suggest that developing a further account of levels of organization is likely to be pointless, and that we can wholeheartedly embrace the deflationary approach" (Eronen, 2015, p. 56).

Levels of organisation is a concept – or at least a phrase – that will be fairly recognisable to philosophers and scientists alike. The concept invokes some commonly found ideas across philosophical inquiry: the world can be usefully structured into distinct 'levels' each of which contains entities of roughly the same kind such as the 'cellular level', the 'molecular level', and the 'psychological level'. Scientific inquiries roughly cluster around these levels and give rise to the disciplines that are so familiar to us: physics, chemistry, biology, psychology, sociology, and ecology, for example. Further, each level is *constituted* by objects at an immediately lower-level allowing us to investigate the relationship between these level-bound objects via strategies of decomposition or theory reduction. A reader might be surprised, then, to see the distinctly opposing views expressed in the opening two quotes above. Philosophers might be used to seeing such opposing views in debates about mental causation, moral realism, or evolutionary psychology, but since when has 'levels of organisation' been such a divisive topic in philosophy of science?

As it happens, not only is the status of 'levels of organisation' a burgeoning topic in philosophy of science today, but just beneath the surface of mainstream discourse in the discipline, discussions of levels of organisation have been simmering for at least the last one hundred years. We can trace a thread of discussion that runs right throughout the 20th century. This discussion has, like many philosophical discussions over time, a 'boom and bust' flavour to it.

Several distinctive accounts of levels were offered (Needham, 1932; Oppenheim & Putnam, 1958; Woodger, 1930), and then critiques of those views developed, questioning either particular aspects of an account (Feibleman, 1954; Novikoff, 1945), or the coherence of the entire concept and it's applicability to scientific inquiry (Guttman, 1976). The process is then repeated. New accounts are developed that seek to learn from the criticisms of their predecessors (Churchland & Sejnowski, 1992; Wimsatt, 1976, 1994) and new criticisms begin to roll in (Kim, 2002; Waters, 2008).

In contemporary philosophy of science, this pattern is continuing. This time, the boom has been ignited by philosophers within the new mechanistic approach to philosophy of science turning their attention to levels of organisation (Bechtel, 2008; Craver, 2007, 2015; Povich & Craver, 2018). As any philosopher engaged in philosophy of the life sciences knows, once the new mechanistic approach turns it's gaze towards a topic, that topic tends to spread like wildfire across publications, conferences, and workshops. In line with the narrative I've been suggesting, the pushback against the mechanistic account of levels has begun to emerge. As with the 20th century cycle, some of this critique is specific to the mechanistic account (Eronen, 2013, 2015) while other criticism is directed at the utility of levels-style thinking altogether, including trying to identify the limits of its application (Ladyman, Ross, Collier, & Spurrett, 2007; Potochnik & McGill, 2012; Rueger & McGivern, 2010).

This brings us to the two quotes that open this introduction and that roughly capture of the state of play in the discussion around the time I began research for this project. On the one hand, through a mixture of historical, pedagogic, and conceptual research, several philosophers have attempted to explore a robust characterisation of levels of organisation that accurately captures its use in scientific practice whilst avoiding suspect metaphysical commitments of philosophical predecessors (Brooks, 2016, 2017; DiFrisco, 2017; Love, 2011; Wimsatt, 1994). Let's call this group 'levels revisionists', given their attempt to revise the concept of levels such that it learns from the lessons of those accounts that have gone before. On the other hand, what has been labelled a 'deflationary' view of levels of organisation has been developed and defended, according to which 'levels of organisation' has outlived its utility as a unified concept. Both philosophers and scientists alike can achieve their aims with less controversial and more precise notions such as 'composition' and 'scale' (Eronen, 2015; Potochnik & McGill, 2012).

In contrast to the ‘levels revisionists’, the deflationary view of levels is developed by philosophers who, on the whole, take a sceptical view towards levels of organisation (Eronen, 2015; Potochnik, 2017; Thalos, 2013) Accordingly, let’s designate this group the ‘levels sceptics’.

I will argue that ‘levels of organisation’ should be understood as collections of vertical and horizontal organisational principles deployed as a means to represent system structure in the conceptualisation of a target system of inquiry as well as for the integration, manipulation, and transfer of information of, and between, systems of inquiry. The central task of the dissertation will be developing, explaining, and demonstrating this characterisation of levels of organisation. Here, I will briefly mention how such a view of levels relates to the potted history of the debate just outlined above. My project enters the fray seeking to incorporate important insights from both the levels revisionists and the levels sceptics, whilst offering a fresh perspective on the role of levels of organisation in scientific practice and philosophical analysis of those practices.

From the sceptics, I adopt a rejection of an ontological or metaphysical significance for the levels of organisation concept, at least in terms of where an analysis of the concept should begin. Going a step further, my focus is firmly fixed on epistemic *practices* rather than the codified *outputs* of such practices, such as theories, laws, and explanations. Instead of theorising about the ways in which the world may or may not be carved into discrete compositional levels; about the causal relations that may or may not cross cut such joints; or the way theories may or may not exhibit hierarchically structured relationships that latch on to the world around us, the research question that drives this dissertation is as follows: does levels of organisation play a fruitful role in scientific practice and if so can it be used to illuminate those practices from a ‘meta-scientific’ or ‘philosophical’ perspective? By ‘meta-scientific’ or ‘philosophical’ perspective, I simply mean standing outside of any one particular scientific project and analysing the use of a concept – or practice – across a range of projects. Often, as is the case here, the broad aim of taking this perspective is to identify generalisable features across those projects that may help illuminate the complex processes of knowledge production in the sciences.

Where I diverge from the sceptics is that my answer to this driving research question will be yes. There is indeed a point to further developing an account of levels of organisation at least in response to the research question I am concerned with here.

Answering positively puts me in contact with the levels revisionists and I accordingly will draw on the influence of recent explorations of the concept of levels of organisation and its significance in biological education and research (Brooks, 2016, 2017). Considered through a revisionist lens, the distinctive aspect of my own account is that which has just been mentioned: I place the significance of the concept firmly within the epistemic *practices* of scientists rather than the *outputs* of those practices. More specifically, my account concerns the role that levels concepts play in representing structural features of a system of inquiry, as well as the development and refinement of those representations, rather than the structuring of outputs of scientific inquiry, such as explanations, theories, and laws.

This dissertation is a project based in a wider-framework of practice-oriented philosophy of science. I'll say more about this framework in the main body of the dissertation (particularly 3.2). For now I'll note that adopting such a perspective provides two dimensions of evaluation that regulate the aims this project. In other words, it provides the criteria for successfully addressing the research question outlined above. The first dimension evaluates how well the characterisation captures *deployments* of the concept in scientific practice. That is, the s*cientific plausibility* of the analysis. The second dimension evaluates how fruitful the characterisation is for the *analysis* of those practices. That is, the *analytic utility* that the characterisation can offer. In this context, analytic utility refers to what aspects of scientific practice an account of levels of organisation can help to illuminate and ultimately how such an account can contribute to a better understanding of the complex processes of knowledge production found in scientific inquiry. In sum, then, my aim is to develop a scientifically plausible characterisation of levels of organisation that is fruitful for the analysis of scientific practices in which the concept is utilised.

There are a few assumptions embedded in these dimensions of evaluation that I will make explicit here. The first is that a prerequisite for meeting the aims of scientific plausibility and analytic utility is to develop a *pluralistic* conceptual analysis of the concept. That is, I take it as an assumption that there is variance to be found in the use of 'levels of organisation' in scientific practice and thus a characterisation that claims to be scientifically plausible will have to be able to account for these variances. Adopting the assumption of pluralism, or taking the 'pluralist stance' (Kellert, Longino, & Waters, 2006) is a wide-spread methodological practice in philosophy of science regarding many prominent scientific concepts.

The concepts of ‘information’ (Floridi, 2011), ‘species’ (Kitcher, 1993), ‘causation’ (Cartwright, 2004; Hall, 2004), ‘genes’ (Waters, 2004, 2006), and ‘complexity’ (Ladyman, Lambert, & Wiesner, 2013; Zuchowski, 2017) have all been approached through the lens of pluralism. Indeed several of the newer analyses of levels also take this assumption of pluralism, (Brooks, 2016; Love, 2011; cf. Wimsatt, 1994) and in this dissertation I follow suit.

A closely related methodological assumption is the ‘domain’ of cases I focus on in order to satisfy these variances and to claim that my account displays scientific plausibility. Very broadly, I will be focusing on deployments of the concept in the life sciences. I’ll be dealing with cases from microbiology, cancer biology, homeostatic processes in human organisms, gene-regulation networks, and protein cascades. The main point I want to make here is that there are two considerations that are relevant to my choice of cases. The first is that I take it as a relatively uncontroversial assumption that developing a scientifically plausible characterisation of a concept does not require showing that the concept is utilised throughout *all* forms of scientific practice. Like most scientific concepts, ‘levels of organisation’ surely has limits of applicability both in terms of being unsuited to broad domains of inquiry, like physics as some have argued (Ladyman et al., 2007; Thalos, 2013), and for certain epistemic practices and goals. The domain of inquiry in which we do see levels of organisation being explicitly referenced in both scientific and philosophical literature is the life sciences. This hardly constitutes an argument for levels of organisation being a key concept in the life sciences, but it gives us a good place to start investigating that issue. Furthermore, I assume that if I can develop a characterisation that is scientifically plausible relative to the range of cases I examine from the life sciences, this will constitute a significant result in itself and be a large step towards cementing the utility of the concept. Similarly, it will certainly be the case that the concept has its limits regarding certain kinds of epistemic aims and activities. Accordingly, I focus on a specific kind of epistemic activity: developing the structural features of a system conceptualisation.

When combined, these two assumptions provide methodological edges to the project. The assumption of pluralism dictates that any claim to scientific plausibility must satisfy a requirement of pluralism – the characterisation must demonstrate how it can capture a wide-range of variation in deployments of the concept. But this requirement of pluralism is curtailed by the second assumption, namely that the wide-range of variation need not extend to domains or practices beyond those in which its usefulness is most likely to be exhibited. The strategy I adopt to negotiate this balance is to draw on cases differentiated by system *type*. Focusing on system type means that the characterisation of levels I offer should be able to contribute to an analysis of those system types wherever they're found, be it within the life sciences or other domains of inquiry. The upshot is that I'll be able to satisfy the requirement of scientifically plausibility based on types of inquiry rather than individual inquires, which should make my analysis generalisable beyond the main case studies I employ here.

In what remains of this general introduction I will outline the structure of the dissertation in a little more detail and give a birds-eye perspective on the argumentative threads that run throughout the work. The most general thematic division of this dissertation is into two parts. The first – spanning Chapters One and Two – is deconstructive. I seek to pick apart the most prominent conception of levels of organisation in philosophy, the Layer Cake Model, and identify key lessons that must be learnt for developing a new characterisation of levels that can avoid the pitfalls of this antiquated but still pervasive view of levels. The second, larger part – Chapters Three, Four, and Five – is constructive. Over the course of these chapters I develop a new characterisation of levels of organisation and demonstrate both its scientific plausibility and the role it can play in philosophical analyses of scientific practice.

The Layer Cake Model (LCM) is an account of levels of organisation that is as pervasive in philosophical thinking as it is inapplicable to contemporary scientific practice. Accordingly, the main goal of the first two chapters is to unpack this perplexing juxtaposition: how has the LCM become so cemented in philosophical thinking about levels, what aspects of the view render it so out of step with contemporary scientific practice, and how can we avoid this outcome when developing a new account of levels? In the opening chapter my focus is on explicating the key aspects of the LCM and understanding why its specific combination of features has such a problematic outcome. This will consist in a careful exposition of the main source of the LCM, Oppenheim & Putnam's (1958) paper 'The unity of science as a working hypothesis' and result in the isolation of *four* key aspects: global scope, monism, aggregation, and distinctness of levels. Additionally, I'll offer three possible interpretations of the LCM: an ontological, epistemological, and methodological interpretation respectively. I'll identify the account's inapplicability to contemporary scientific practice as the result of an essential tension at the heart of the LCM that consists in the combination of several of these key aspects.

The work done in Chapter One will provide insight as to what the LCM consists in and why it is not particularly useful for analysing scientific practice, but this does not explain it's continued influence of levels-talk in philosophy today: why have philosophers clung so tightly to this particular conception of levels of organisation and how might we rectify this dependence? I'll argue that, in fact, philosophers haven't clung tightly to the LCM, rather it has become a conceptual free-rider in the evolution of other prominent philosophical discussions over the last one hundred years. Specifically, I'll focus on in its deeply entangled relationship to the concepts of reduction and emergence, with the aim of understanding how the LCM seems to persist as the default view of levels throughout debates concerning reduction & emergence. My strategy will consist in demonstrating the compatibility of contrasting views of reduction and emergence with corresponding interpretations of the LCM. The compatibility exhibited is not claimed to imply that proponents of, for example, ontological emergence are committed to the LCM. Rather, I use the compatibility to pick out what issues a view of inter-level relations (i.e. an account of reduction/emergence) must explicitly address in order to be *incompatible* with the LCM. I conclude that before a general characterisation of levels can be given I'll have to address these lessons and offer a different framework for thinking about inter-level relations.

In Chapter Three I take on that task. First, however, I take a step back to introduce in more detail the practice-oriented approach to philosophy of science that strongly influences my approach to issues of reduction and subsequently to developing my characterisation of levels of organisation. My own version of the practice-oriented approach consists in developing the notion of a research strategy as a framework for analysis. To jump ahead slightly, the constituent parts of my characterisation of levels of organisation will be situated as different aspects of a research strategy; interacting with other components of a strategy. In this chapter, I start with what will become one part of the characterisation: inter-level relations (or as I go on to label them 'vertical organisational principles'). The main claim I develop in the chapter is that focusing on reductive and non-reductive *practices* results in a rejection of the distinction between the two. Rather the concept is essentially comparative: a research strategy is only more, or less, reductive than another according to a frame of reference. This frame of reference is provided by what I label 'the continuum of research strategies' (Fig. 6). The continuum is a sliding scale of possible system conceptualisations based on: (1) assumptions regarding the decomposability of the system and (2) the foregrounding or backgrounding of certain kinds of system-wide dynamics. I test my proposed framework by applying it to a recent discussion in cancer research. I show that prominent 'reductive' and 'non-reductive' approaches to carcinogenesis can be re-conceived as boundary-markers for the field, delineating the area on the continuum within which progressive inquiries in the field are exploring. I conclude by showing how this new account of reduction explicitly addresses the lessons learnt at the end of the previous chapter and, in doing so, exhibits explicit incompatibility with the LCM.

Chapter Four contains the bulk of the development of my characterisation of levels. I develop and explain key features of the account. The outline of my account is that levels of organisation should be understood as a collections of purpose-relative organisational principles that are deployed to represent system structure. I divide organisational principles into two categories: vertical and horizontal. Vertical principles pertain to level differentiation and the two sorts of vertical principle I will examine in detail are *decomposition* and *control* – each of these will be given much more detailed and refined meanings. Discussions of levels mostly centre on these kinds of vertical (or 'inter-level') relations and for good reason: they provide the sense of 'higher' and 'lower' that is probably the most distinctive feature of organising individuals into *levels*. However, I will introduce another kind of organisational principle that belongs to a levels of organisation concept: horizontal principles. These principles pertain to *intra-level partitions* (divisions in the system at the same level rather than at different levels), *intra-level interactions,* and finally, considerations that relate to the process of carving a *boundary* for the system.

Also contained with my characterisation will be that a levels of organisation concept can be utilised to understand the epistemic practices of *integrating* different models, *manipulating* target systems of inquiry, and the *transference* of information between models. This aspect of my characterisation seeks to address the analytic utility of the levels concept by specifying exactly what kinds of epistemic practices such an account can help to illuminate. The rest of the chapter is concerned with putting this characterisation to the test by drawing on scientific examples. I'll be using these examples will to justify the account's scientific plausibility and analytic utility. I address cases studies from three 'paradigm points' on the continuum of research strategies that capture three specific system types.

I'll look at aggregative system representations from the Human Microbiome Project; component systems from studies of Glucose Homeostasis in humans, and integrated control systems that are used to represent MAPK cascades (a kind of protein signalling pathway).

In the final chapter I turn my attention back to the current philosophical literature on levels of organisation. My primary goal is to situate this newly developed characterisation of levels within the discussion of levels as it appears in contemporary philosophy of science and as a result to highlight its distinctive features and the new insights they bring to the table. My strategy is to offer a reinterpretation of the most prominent account of levels in the literature – mechanistic levels of organisation – such that the account can be understood *within* the context of the view I offer. Specifically, I'll argue that the mechanistic account offers an analysis of a subset of system types and accompanying strategies of investigating those systems. This re-interpretative analysis will enable me to underline the most distinctive features of my characterisation of levels; those that I hope offer the most interesting contribution to discussions of levels in contemporary philosophy of science. I'll also suggest that proponents of the mechanistic account could avail of my reinterpretation in order to respond to prominent criticisms of their view. To bring the project to a close, I'll zoom out and consider the general features of the LCM, the mechanistic account, and my own characterisation of levels side-by-side, in the hope of clearly illustrating the nuanced position offered by the general features of my own view of levels of organisation.

Overall, my hope is to show that 'levels of organisation' is a tool that is key part of the research strategies adopted by scientists towards the development of a target system of inquiry. Along the way, I'll develop a new framework for thinking about reduction and emergence that positions these concepts as comparative tools for understanding several key driving assumptions behind research strategies, rather than dichotomous labels that generate potential stand-offs. I also hope to show that the account of levels has complementary functions for the philosopher or scientist interested in analysing those strategies, including their dynamic development through integration, manipulation, and information transfer, and the exploration of new approaches to a problem via, amongst other things, the re-conceptualisation of the target system. I seek to achieve both of these goals without stretching the concept of levels beyond its justifiable limits of application or caricaturing the scientific case studies I invoke. I'll conclude by indicating some fruitful lines of research that this project has opened up.

1

What is the Layer Cake Model of the World?

1.1 Introduction

The Layer Cake Model is, without doubt, the most pervasive conception of levels of organisation found throughout 20th century philosophy and it continues to have a significant influence on the way levels of organisation are understood in philosophy today. In its simplest form it presents a neat picture in which the world, and our scientific inquiries into it, can be organised into a single model that is structured according to distinct levels of organisation, ranging from the most fundamental objects of inquiry ('lowest) to the most derivative ('highest'). It also happens to be a straightforwardly implausible account of levels from a scientific perspective. In order to develop a convincing account of levels of organisation that is scientifically plausible, I'll have to unpack this awkward juxtaposition between the LCM's prevalence and its scientific implausibility. That task will take place over the next two chapters. In this chapter I'll focus on carefully explicating the features of the LCM. The chief difficulty with tackling the LCM is pinning down what exactly it consists in. From the perspective of contemporary philosophy of science concepts that are central to the LCM such as, 'organisation', 'decomposition', 'distinct levels', 'part-whole' and so on, have a plethora of interpretations and are put to use in a variety of ways. The main question that drives this opening chapter is how exactly should we understand these – and related – ideas in the context of the LCM and how do those concepts interact to form this all too familiar picture of levels of organisation?

I begin with a detailed discussion of the LCM as found in Oppenheim & Putnam's work (1.2). Following Oppenheim & Putnam, I will then place the LCM in the context of the Unity of Science Project (1.3). This context is crucial both for understanding how the LCM was part of a package of ideas, the combination of which served as a working hypothesis for the unity of science project itself, and for understanding what the LCM *came to be*, long after this working hypothesis had become a degenerating research programme. This will allow me to pick out general features of the LCM which can be used as concrete indicators that the levels structure under consideration is indeed the LCM, or a structure that very closely resembles it (1.4). These features are: a global scope of application, a monism about levels of organisation, an aggregative system type, and distinctness of levels (Box 2).

Finally, I'll use these concrete indicators – or 'general features' – about the LCM to explain exactly what is so scientifically implausible about it. I'll argue that there is a deep tension at the heart of the LCM that lies at the intersection between three of its essential aspects: *stepwise decomposition*, *global scope*, and an inter-level relation of *aggregation.* Any one of these aspects alone, or even two together, might be sustainable, but together they yield a picture of levels that, when applied to contemporary scientific inquiry, will rapidly fall apart (1.5). I'll conclude that with this understanding in place, we will be in a position to tackle the second aspect of the LCM's awkward juxtaposition: explaining it's continued influence on the philosophical discussions in which 'levels of organisation' plays a significant role: reduction and emergence.

1.2 Oppenheim & Putman's Layer Cake Model

To find a clear presentation of the LCM, there's no better place to look than Paul Oppenheim & Hilary Putnam's (O&P) 1958 paper 'The Unity of Science as a Working Hypothesis'. O&P provide a conceptual analysis of levels of organisation, the result of which is now referred to as the Layer Cake Model, or sometimes 'layer-cake levels'. The main aim of the chapter is to lay out conditions that must be met – or more strongly, ways in which the world & our investigations of it must be – in order for the unity of science project to be successful (e.g., Neurath, 1937; Neurath et al., 1938). Thus, the LCM is an account of levels of organisation the function of which is to postulate a system within which the unity of science project may be attainable. A small disclaimer is required here. In what follows I'll be presenting the LCM pretty much as O&P envisaged it. By and large, I think their view corresponds to what philosophers have in mind when referring to the LCM. However, there are two idiosyncratic features of O&P's LCM that I'll eventually diverge from in order to provide an analysis of the role that the LCM has played in the history of philosophy of science.

The first is the 'correspondence thesis' (Brooks, 2016), which maps different variants of the LCM onto one another (1.4.4). The second is the account of *reduction* that O&P assume (1.3.1). Firstly, I'll introduce, what I consider to be an accurate interpretation of O&P's LCM before discussing these divergences in more detail where they arise.

Box 1 contains the six conditions of adequacy that must be fulfilled in order to construct a system that fulfils this role and, taken together, comprise the conditions that result in the LCM itself.

(1) There must be several levels
(2) The numbers of levels must be finite
(3) There must be a unique lowest level (i.e., a unique 'beginner' under the relation 'potential micro-reducer'); this means that the success at transforming all the *potential* micro-reductions must, *ipso facto*, mean a reduction to a single branch
(4) Any thing of any level except the lowest must possess a decomposition into things belonging to the next lower level. In this sense each level will be a 'common denominator' for the level immediately above it.
(5) Nothing on any level should have a part on any higher level
(6) The levels must be selected in a way which is 'natural' and justifiable from the point of view of present day science. In particular, the step from any one of our reductive levels to the next lowest level must correspond to what is, scientifically speaking, a crucial step in the trend towards overall physicalistic reduction.

Box 1 The Six Conditions for the Layer Cake Model.
Adapted from Oppenheim & Putnam (1958, p. 9)

The following instantiation of the model, i.e., a possible way of populating the different levels in the model that meets the six conditions, is proposed by O&P:

(L6) Social Groups
(L5) Multicellular living things
(L4) Cells
(L3) Molecules
(L2) Atoms
(L1) Elementary Particles

Hopefully, this token instantiation of the model will be familiar and so I will use it to explicate O&P's six conditions. It is important, however, to bear in mind that this is merely an instantiation of the LCM and not the account of levels of organisation itself. By an instantiation I mean that O&P put no restrictions on the kinds of entities that might populate the actual levels, nor do they put a restriction on the amount of levels in the model. So, one could produce different token instantiations of the model and still comply with the six conditions of the LCM. Perhaps one has good reasons, given the conditions of the LCM, to think that there should be a level between (L3) & (L4) called 'macromolecules' giving us seven levels in the model. Further, perhaps there are good reasons to add a new highest level above (L6) called 'Ecosystems', giving eight levels in the model altogether. All of this is perfectly possible under the conditions of the LCM and, importantly for later, criticisms of the LCM as an account of levels of organisation should not focus on whether any one token instantiation of the model gets it right from the view point of either contemporary science or philosophy.

1.3 Putting the LCM in Context

As I will argue in Chapter Two, one erroneous assumption in contemporary philosophy has been that the LCM stands or falls with the unity of science project. The unity of science project is no longer pursued, at least not in any recognisable form from the early twentieth century, but the LCM has not been jettisoned from philosophical reasoning about levels of organisation; the two seemingly come apart. However, in order to fully understand the conditions of the LCM it is necessary place them in their context as conditions for an attainable unity of science project, not only in a general sense, but specifically how O&P understood such a project to develop. Without this context it is difficult to get a grasp on where some of the conditions come from and how they depend upon each other.

1.3.1 The Unity of Science Project

The contemporary historical roots of the unity of science project can be found in the work of the Vienna Circle (*Wiener Kreis*) and early 20th Century Logical Empiricists (sometimes referred to as logical positivists). The Circle contained a group of philosophers and scientists who met regularly in Vienna roughly between 1924 and 1936, although the term is often used to capture the work of its members beyond these dates. Prominent members included Moritz Schlick, Rudolf Carnap, Otto Neurath, Hans Hahn, Olga Hahn, Philipp Frank, Viktor Kraft, Gustav Bergmann, and Herbert Feigl.[1] For present purposes the most salient aspect of the Vienna Circle was their work on a unity of science or unified science (*Einheitwissenschaft*) beginning with the organising of several International Congresses for the Unity of Science[2], and the subsequent publication of *The International Encyclopedia of Unified Science* (Neurath et al., 1938). The scholarship on the unity of science movement is vast, partly due to the nuanced differences between the positions of the members of the Vienna Circle, as well as those who followed them.[3] For the purposes of providing a succinct overview, the following paragraph is inspired by Ian Hacking's (1996) summary.

'Unity' can refer to a plethora of categories but is most often taken to refer to three central themes. Unity of scientific language can refer to the desire for a single language of scientific discourse (Carnap, 1934, p. 32), or at least the potential translatability of different scientific languages to aid co-operation (Neurath et al., 1938, p. 15). Unity of scientific reasoning or methodology, often manifested in the debates surrounding the demarcation of science from non or pseudo-science (Popper, 1959).

1 This abridged list of members is taken from Cartwright et. al., (1996, p. 77).
2 The first being held at the Sorbonne, Paris 1935 (Neurath et al., 1938, p. 26).
3 For a comprehensive overview of both the Circle and its key members see Friedrich Stadler (2001).

Finally, metaphysical unity of science concerns the relationship between phenomena in the world and the branches of science that study those phenomena (Hempel and Oppenheim, 1948; Nagel, 1979, p. 335ff). Hacking (1996, p. 40) subdivides metaphysical unity into three further categories:

A. ‘Interconnectedness’ - all kinds of phenomena must be related to one another
B. ‘Structural’ - there is a unique, fundamental structure to the truths in the world
C. ‘Taxonomic’ - there is one fundamental, ultimate way of system classifying everything: nature breaks down into natural kinds.

All three of these central themes implicitly require a sense of reduction. The first theme, which O&P label ‘the unity of language’ requires the reduction of one scientific discourse to another. The second theme they label as the ‘unity of method’ and pertains to the unification of standards of explanation or evidence. This theme does not straightforwardly require a sense of reduction and, perhaps as a result, they do distance themselves from this particular goal. The third theme roughly corresponds with what they label ‘the unity of laws’ and pertains to the reduction of laws, and by implication the entities to which those laws refer, to one another (Oppenheim and Putnam, 1958, pp. 3–5). Therefore, to complete the context in which O&P provide their account, it is crucial to understand the concept of reduction they have in mind.

Firstly, they make it clear that they are interested only in synchronic reduction – reduction of one thing to another at a particular time – rather than diachronic reduction: reduction or replacement of one thing by another over time. For example, they are interested in how the laws, theories, and explanations of, say, cellular biology (L4) might be reducible to the laws, theories, and explanations of atomic physics. They are not attempting to provide an account of historical development; for example, the replacement of the phlogiston theory of combustion with theories of oxidation. They do, of course, include the caveat that the reductions they posit have not yet been achieved; it is a working hypothesis.

Secondly, they straightforwardly adopt a particular account of reduction developed by Kemeny & Oppenheim (1956) according to which:

Given two theories T1 and T2, T2 is said to be reduced to T1 if and only if:

(1) The vocabulary of T2 contains terms not in the vocabulary of T1
(2) Any observational data explainable by T2 are explainable by T1
(3) T1 is at least as well systematized as T2.

(1) Covers a linguistic sense of reduction – the vocabulary of the reduced theory contains more terms than the reducing theory. Conditions (2) & (3) pertain to explanatory and predictive power, capturing the theme of ‘unity of laws’. Condition (2) concerns observational data – the reducing theory can explain all observational data that the reduced theory can. Condition (3) concerns theoretical virtues, broadly captured by the notion of ‘systematisation’. For Kemeny & Oppenheim, the idea of one theory being more, or at least as well, systematised than other is meant to reflect a balance between simplicity and ‘strength’, where strength is best understood as explanatory and/or predictive power. The idea is that reduction may not always produce simpler theories, but this is acceptable insofar as that increased complexity is counterbalanced by an increase in the explanatory or predictive power of the theory (Kemeny and Oppenheim, 1956, pp. 11–12). Whole branches of science can be reduced to another if we take a branch of science to consist of simply the sum of its theories; that is we can define T2 as the entire collection of theories in a branch of science B2. Thus, B2 reduces to B1 (another branch of science) when T2 reduces to some theory T1 in B1.

A final point on reduction, O&P are clear that with the reduction of a theory comes a corresponding reduction of entities in the world to which that theory refers. More specifically, we must assume that, “corresponding to each branch we have a specific universe of discourse UB1, and that we have a part-whole relation, Pt” (Oppenheim and Putnam, 1958, p. 6). The reductive relationship between two universal branches of discourse captures the relation between all the terms in that branch of science and directly corresponds with a part-whole relation of entities in the world. In sum, a reductive relation, should it be attainable, holds between individual theories that can be collected together into whole branches of science, the language of which actually maps directly onto entities in the world, such that any reduction of one theory to another will directly map onto a part-whole relation between the entities that the theories describe.

Before heading back to the six conditions of the LCM, let's pause to summarise O&P's understanding of the unity of science project and it's relationship to the LCM. Their model is supposed to provide a structure to the world in which a unity of science project is possible to attain. They understand the project as primarily concerned with the unification of scientific language and laws, where the latter pertains to scientific explanations and the world they latch on to (a form of metaphysical reduction). These goals are to be constrained by the account of reduction they employ. On that account a successful reduction of one theory to another requires that the reducing theory (1) has less terms than the reduced theory, (2) can explain any observational data that the reduced theory can, and (3) makes up for any increased complexity by providing more explanatory and/or predictive power than the reduced theory.

1.3.2 Understanding the Six Conditions

Let's return to the six conditions of the LCM (Box 1). Condition (1) states that 'there must be several levels'. O&P position this as a fairly intuitive necessary condition for any account of *levels* of organisation. Of course, there's nothing incoherent about the idea of a system with only one level of organisation, but this is tantamount to saying that there's no *vertical* structure in the system. For individuals in a system with just one level of organisation we might get groupings of organised structure: they may interact in specific sequence; be partitioned according to shared features and/or the methods used to discover them. However, want they won't be is positioned in a 'higher-lower' relation to one another. For that, we'll at least need two levels. O&P are concerned only with vertical structure; hence why, for them, condition (1) is necessary.[4]

4 To foreshadow a little here, my characterisation of levels will distinguish between the sorts of 'same-level' organisational features of a system just mentioned – i.e., *horizontal structure*, governed developed using what I will label 'horizontal organisation principles' – and the relations that give *vertical* structure, i.e., the sort of structure that gives a 'higher-lower' relations. There's no place for 'horizontal' structure in O&P's account so we won't be returning to this issue for a while, but there's lots more detail on this in Chapter Four particularly in 4.2.1. & 4.2.2.

Conditions (2) & (3) both concern the fact the that LCM is a working hypothesis for the unity of science project. Starting with condition (2): ‘the levels in the system must be finite’, this condition is important for the unity of science project because without it the project would never simply never be attainable.

Condition (3): that here must be a unique lowest level, indicates that the most important sense in which the levels must be finite is in terms of going towards the *lowest*, or most fundamental, level. Reducing all scientific terms to a unique language, or all explanations to a single unified account, requires that the process actually ends somewhere – thus the levels must be finite. The unity project was one of reduction, so, whilst perhaps there may need not be an upper limit on levels, there must be a lower limit. Consider that the bulk of a unifying project would be complete so long as a unified bottom-level had been established. In theory, new levels could be discovered upwards without threatening the success of the project, so long as these levels were made consistent with the unique lowest one. An infinite extension of levels downwards would make the unity project impossible, in principle. Presumably, for O&P at least, this requirement is also generated from adherence to condition (6): that there is a unique lowest level that is ‘natural and justifiable from the point of view of present day science’, this is makes perfect sense if you suppose that the unity project was capturing the state of affairs in ‘present day science’ at the time.

Condition (4): that ‘any thing of any level except the lowest must possess a decomposition into things belonging to the next lower level’, is crucial for understanding a major characteristic of the LCM – it is an aggregate, fully decomposable system. I will return to this in more detail in the next section, but for now the idea is that the each level in the system is formed is by a simple decomposition of individuals at one level into their aggregate parts; the aggregate parts then form a new, immediately lower level. As noted above, the reduction of theories or whole branches of science correspond directly to part-whole relations of entities in the world. If that notion of reduction is correct, then the way in which the levels will be populated should consist entirely in aggregative part-whole relations from bottom to top.

Condition (5): that 'nothing on any level should have any part on any higher level' is another result of the unity project understood within O&P's framework. It rules out cross-level or multiple-level entities and corresponding theories whose terms refer to them. Again, entities of these sorts would make a unity project almost impossible for several reasons. First, if entities could pop up at various levels at various times, it would seem much less likely that the reducing theories would contain less terms than the reduced theories. Second, if entities straddled different levels of organisation then there's no way that every term in every theory of a branch of science could map onto a fixed part-whole relation. Some, perhaps even most, still could do so, but if you had an entity popping up at, say, two levels of organisation at any given time, then it could be both a part of a whole at L3, putting it at L2 whilst also having parts at L2, putting it simultaneously at L3 – both a whole and a part at different levels of organisation. Finally, if any entity had a part on a higher level, then in principle that entity could not be reduced to a lower level thus blocking the unity project altogether.

I read condition (6) as an endorsement both of the unity project but also of a general attitude towards a scientific approach to philosophical reasoning. O&P spend a large portion of the paper proving reasons why their 'hypothesis' does, in fact, reflect science (at the time). Interestingly, that they do so via theoretical & methodological virtues one might use to evaluate bone fide scientific hypotheses and virtues that continue to be espoused today. They argue that their model is guaranteed to be *simpler* than then-alternatives such as vitalism or 'pyschism'. They are argue that the model is *fruitful* both because it can stimulate scientific inquiry and because even if it fails it will provide important insights about the relations between parts and wholes (1958, p. 12). This latter point in particular continues to be explored under the guise of reductionist heuristics – that is, research strategies (of which a hypothesis might be a component) that may be fruitful precisely because they fail in systemic ways (e.g., Bechtel and Richardson, 2010; Wimsatt, 2006; see also 3.2). Finally, they claim that the reductionist methodology is one inherent in scientific practice – something they call a 'Democritean tendency'.[5]

5 O&P (1958, p. 16) describe a Democritean tendency as, "to try, insofar as is possible, to explain apparently dissimilar phenomena in terms of qualitatively identical parts and their spatio-temporal relations." I take it that O&P are referring to the ancient Greek philosopher Democritus, known for his 'materialism' and 'atomism' that contrasted with Aristotelian teleology – i.e., explaining a phenomena by the interactions of its parts, rather than in reference to a 'final cause' or 'purpose'.

1.4. The General Features of the LCM

Hopefully, my analysis has confirmed the intimate relation between the construction of O&P's LCM and the goals of the unity of science project. In the next chapter I will try to tease the LCM apart from the goals of the unity of science project in order to demonstrate both that, and why, it remains so prevalent in contemporary philosophical reasoning. Before that, however, let's take a step back from the exact exegesis O&P's analysis of levels to get a sense of the overall character of the LCM itself. By 'character' I mean the broad-stroke features of the account that give structure to a system. In other words: if one adopts the LCM as an account of levels of organisation what conceptualisation of the target system will follow; what will our system look like and what options will we have for investigating that system? The purpose of this section is to make the answer to this question clear and thus to give an understanding of what impact adopting the LCM might have for thinking about constructing and investigating target systems of inquiry.

1.4.1 The Scope of The System

The first striking feature of the LCM is the scope of the system it is concerned with. The LCM is about the world as a whole, or at least all of the entities that come within the purview of our scientific investigation. In other words, the LCM is a global account of levels of organisation. This is perhaps to be expected, given its close ties to the unity of science project. That project is not concerned with the structural relationship between two particular target systems, or even between two entire branches of science. Rather, it is concerned with the structural relationship between all branches of science and, by implication, all entities to which those scientific investigations refer. Thus the LCM is an account of levels that construes the entire world as one system, partitioned into distinct levels.

Closely related to taking a global scope is that the LCM is a monistic account of levels of organisation. It is a small jump from attempting to stratify all scientific theories into one system that there should be only one way of successfully completing such a task. Of course, O&P are clear that the LCM is a working hypothesis, but presumably they believed that the hypothesis could be successful and if it were to be so then, given that it contains within it all branches of scientific inquiry, it would be the only account levels of organisation.

The claim that the account is global and monistic is not inconsistent with the idea that there may be several different instantiations of the model whilst it remains a *working hypothesis*. As noted in 1.2, there could be many numbers of levels populated with many different entities, but eventually the outcome of a completed LCM would mean no further levels of organisation concept could apply to the world.

It is important to note that these ideas – a global scope and a monistic account – are very closely related, but they are not the same claim. To see this consider that you could have a global account that is nevertheless pluralistic. You could contend that *everything* in the world can be put into system, but there are many different legitimate ways to develop the *structure* in that system; ways that are perhaps even inconsistent with one another. For example, you might think that these differences can be generated by *different* notions of part-whole relations.[6] As we'll see in the next subsection O&P understand a part-whole relation to be one of aggregation not of causal-role or functional decomposition. Thus, a major sense in which the LCM is a *monistic* account comes from having only *one* way of building levels in the model.

In response to the LCM, it is fairly common to see talk of levels as having a 'local scope' and being pluralistic about how those local levels concepts are deployed (Craver, 2015; Kaiser, 2015; Potochnik and McGill, 2012), but we can find at least one account of levels that might be understood as having a local scope and yet be a monistic account. In William Wimsatt's (1994) account of levels of organisation there is really only *one sense* in which he understands what levels of organisation are, i.e., it is a monistic account. However, not everything in the world can be organised into levels. Rather, as complexity increases in systems, the clear structure of discrete levels of organisation begins to break down, first into 'perspectives' and then into 'causal thickets'. Hopefully, this shows that the LCM's commitment to a *combination* of a global scope and a monistic account is not a trivial matter.

6 In fact, it is not *just* different part-whole relations that can will result in different system structures. Exploring different *vertical* ordering principles will form a core part of the pluralism of my characterisation of levels (3.3, 4.2, and 4.3 will give much more detail on this).

1.4.2 An Aggregative System that is Simply Decomposable

The LCM takes its system of interest to be the entire world, but what kind of system does it take the world to be? The answer is an aggregative system that is simply decomposable. It is not clear whether O&P would have used these exact terms – they didn't in the key 1958 paper at least – but their explication of the part-whole relation fits these contemporary labels. In contemporary work on kinds of system, a tripartite division is made according to the susceptibility of systems to decomposition. Aggregative systems are simply decomposable; component systems are near decomposable; and integrated systems are minimally decomposable (Bechtel and Richardson, 2010, pp. 25–27; Wimsatt, 2007, pp. 277–281). This tripartite division will be a major theme throughout this dissertation and I will return in detail it in Chapter Three, albeit providing my own modifications to the divisions in kinds of system. There, I'll consider some more technical notions of 'aggregative' that have been offered (Wimsatt, 2007, pp. 280–281). For now, I'll try to give as clear and basic understanding of what I mean by aggregative, it's relevance for thinking about system types, and, by extension, what I claim O&P had in mind when thinking about the type of system captured by the LCM.[7]

Aggregation is the collecting, assemblage, summation, or accumulation of objects, process, or units of any kind in fact, together. A 'heap' of sand is an *aggregation* of individual grains of sand. The final points score of a football team at the end of the season is the *aggregation* of the points they gained in each individual match. I could go on but the bare notion of aggregation is pretty much just what you pre-theoretically think it is: when you start with discrete individuals and then collect them together, the resulting collection is an aggregation of those discrete individuals. The relevance of the concept of aggregation to thinking about systems is really about about marking features that the system *doesn't* have rather than those it does. In this sense, aggregation means that there's *nothing more* to the development of system structure than this process of adding things together or breaking things apart; hence labelling the process by which such systems are built as 'simple

7 As you might anticipate, my criticism the LCM will involve the fact that it conceptualises the world as an aggregative system (1.5). However, in In Chapter Four, I'll argue against the idea that there's nothing particularly useful about representing systems as having an aggregative structure. I'll argue that conceptualising a system as aggregative can be a powerful tool for the organisation of large data-sets that serve as blueprints for developing inquiry, particularly large-scale research projects (see 4.3.1).

decomposition'. If we were to start with an unstructured system containing three objects, then to 'simply decompose' that system and thereby build aggregative system structure, would be to form a lower-level than contained all and only the spatio-temporal parts of the three objects.

Turning back to O&P, we can actually find two senses of 'decomposition' utilised in the paper (Oppenheim and Putman, 1958. p. 11). It is not clear if they want to advocate both but I will assume they do because either one (or both taken together) results in understanding the world as an aggregative system. They begin with a 'wide' sense of decomposition according to which x is a part of y when x is spatiotemporally contained in y. This is precisely the sense of aggregation I introduced above. Levels in the system are constructed simply by understanding what is spatiotemporally contained with the wholes at each level above. Levels are constructed under the assumption that every object at level n can be understood both as a part of a whole at level n+1 and a whole with parts at level n-1.

The iterative process of decomposing wholes into their lower-level spatiotemporal parts builds the 'higher-lower' structure of a system. O&P also refer to the part-whole relation in a 'narrower' sense, which requires developing a formal calculus of the relation where the notion of 'parthood' fulfils axioms of that calculus. Given the references to the work of Rescher (1955) and Rescher & Oppenheim (1955), I take O&P to referring to the field of mereology here. Mereology is a specialised philosophical project which takes the notion of parthood to be its basic object of study and attempts to develop a formal system with axioms that govern the part-whole relation (Casati and Varzi, 1999; Leonard and Goodman, 1940; Leśniewski, 1992; Simons, 1987).

O&P's reference to Rescher in particular affords a small departure here that can help sure up the understanding of O&P's target system as aggregative and simply decomposable. Rescher, among others (e.g., Ruben, 1983) explicitly argued that the mereological sense of 'part-whole' seems to fail when applied to certain scientific and everyday examples. Specifically, he argued that plenty of 'part-whole' talk seems to rule out the relation being transitive, despite transitivity being one of the axioms of classical mereology.[8]

8 ((Pxy & Pyz) → (Pxz)) Any part of any part of a thing is itself part of that thing (Casati and Varzi, 1999, p. 35). Rescher (1955, p. 10) argued that in biology, for example, the fact that a DNA molecule is part of the cell, doesn't make it also part of the organ in which that cell is embedded

The way in which the issue of transitivity was diffused in classical mereology is important when considering the LCM. One work-around for this problem was to note that these apparent counterexamples to the transitivity of part-whole relations concerned a different kind of relation from the one classical mereologists originally had in mind. Specifically, it concerned a (quasi) functional or causal-role relation between parts and wholes. In other words, when x is a part of y not because (or not only because) x is spatiotemporally contained within y, but because x plays a causal role in the functioning of y. Thus, mereologists argued that classical mereology contains the axiom of transitivity because it is concerned only with structural part-whole decomposition and not causal-role part-whole decomposition. Such a system could of course be developed but it would presumably not contain an axiom of transitivity (Ruben, 1983; Simons, 1987; Varzi, 2006).

We can use the presence or absence of transitivity in O&P's account as a key indicator of whether they had a structural or causal-role sense of decomposition in mind; even when referring to the 'wider' sense of decomposition. O&P are clear that, not only can the relation of reduction be thought of as transitive; it must be so if the unity of science project is to be plausible. They argue that the transitivity of reduction avoids a common misconception about the unity project, namely that all theories would have to be explicable in terms of theories making reference only to elementary particles (the lowest level in the system). Transitivity, however, makes the project cumulative. Reduction need take place only one level at a time (and never more than one at a time), because with each stepwise reduction, transitivity ensures that all levels get reduced to the lowest one without big jumps in levels being required (Oppenheim and Putnam, 1958, p. 7).[9] Thus, the transitivity of levels in the LCM gives a clear indication that the decomposition, be it in 'narrow' or 'wide' sense, is strictly aggregative and does not concern causal-roles or functional decomposition.

O&P have a little more to say about the status of the levels throughout the cumulative unity of science project. They note that, due to transitivity, all entities will be a part belonging to the lowest level, but "the highest level to which a thing belongs will be considered the 'proper' level of that thing" (Oppenheim and Putnam, 1958, p. 10). This is an important clarification.

9 I.e., if a theory from (L3) is reduced to one from (L2), and subsequently, that theory from (L2) is reduced to one from (L1) then, by transitivity, the theory from (L3) has been reduced to (L1), without the terms in (L1) needing to map directly onto those from (L3).

Not only will reduction take place on level at a time but, despite reduction, the levels remain stable because an entity's 'proper' level will still be the highest at which it is to be found in the model. Put differently, even though the LCM is an account of levels designed with the unity of science project in mind, the levels themselves do not disappear as the project develops. For now, the main point to note is that kind of system that the LCM is designed to structure and that is an aggregate system, subject to simple decomposition.

1.4.3 Distinctness of Levels & Step-wise decomposition

Closely related to the issues of aggregation and cumulation is the *distinctness of levels*. Levels in the model must be distinct. There are two components that I'm referring to in the conception of 'distinctness' here. The first is that entities cannot appear on more than one level: higher or lower; they are level-bound. The second is that decomposition into distinct levels must take place in a step-wise fashion, that is they must take place one at a time. To see both of these components in action we need to combine conditions (4), (5) and the recognition that the decomposition relation utilised by O&P is *aggregation*. Condition (5) states that: nothing on any level should have any part on any higher level. This is a slightly convoluted way of saying that if an object has parts, its parts on *one* lower-level in the model. Condition (4), specifies that step-wise aggregative decomposition is exhaustive, so its parts are on the *next lowest* level.

In sum, if objects have parts then they are at a lower level, and they are *all* at the *next* lowest level. Next we can note that an aggregation of parts (a whole) is not identical to any one of its parts taken in isolation.[10] If this was not the case, then there would be no way to maintain the levels in the structure, they would collapse. As just noted above, O&P wanted the levels to maintain as least some robustness in order to make the unity project cumulative

10 I'm using the 'wide' sense of decomposition here but if I'd followed the 'narrower' mereological sense of part-whole relation, this point could be made using the mereological relation of *proper part*. *X* is a proper part of *y* iff *x* is a part of *y and x* is not identical to *y* (Casati and Varzi, 1999, p. 36). Understood this way, distinctness of levels refers to the claim that all of an objects *proper parts* are at the next lowest level of organisation. If levels are differentiated as a relation between wholes and proper parts, and two levels are not identical, then it's clear that an object cannot be at a lower level (it can't be at the same level as its proper parts), and it cannot be at the higher level because, *ex hypothesi*, it is a proper part of an object on that level.

rather than ‘utopian’. This is presumably why they explicitly state that each level in the model is *necessary*. (Oppenheim and Putnam, 1958, p. 10).

Taking this altogether we can see that objects in the model can only appear on *one* level in the model. Take an object *X* at level Ln. Firstly we know that its parts are on the next lowest level (Ln-1), so it cannot belong there; secondly if we know that it is a *part* of something at a higher level (Ln+1), then it cannot belong there either. So it’s stuck precisely at the level between; every object in the model is level-bound.

1.4.4 Three Kinds of LCM & The Correspondence Thesis

Throughout the discussion so far I have moved between several *variants* of LCM interchangeably. The final point I want to make is to get clear on the differences between them. We can pick out three variants of LCM differentiated according to *what* they structure in accordance with the six conditions in Box 1. One variant is an *ontological* LCM according to which the constituents of the model are *entities* in the world. By ‘entities’ I refer to individuals and their properties collectively. Another is *epistemological* according to which the model structures theories and/or explanations; things I will label ‘epistemic items’. A final variant is *scientific inquiry* according to which our scientific endeavours are organised around the model, for example the ordering of scientific disciplines like biology, chemistry, and physics; I label these ‘epistemic practices’ (Fig. 1).

The LCM, as presented by O&P, can be understood as the unification of these three possible interpretations. Our scientific endeavours (epistemic practices) are organised around the formulation and interrogation of theories and/or explanations codified in a scientific language (epistemic items). These epistemic items come structured into distinct levels of organisation giving rise to an isomorphic structure for scientific practice (i.e., a level-bound structuring of scientific disciplines). Finally, the theories and explanations come structured in this way precisely because the objects or entities that they are about *also* exhibit a hierarchical structure; that is, because objects in the world cluster into discrete levels of organisation. Thus, each kind of LCM maps directly onto one another giving an overarching structure to the world and our investigation of it.

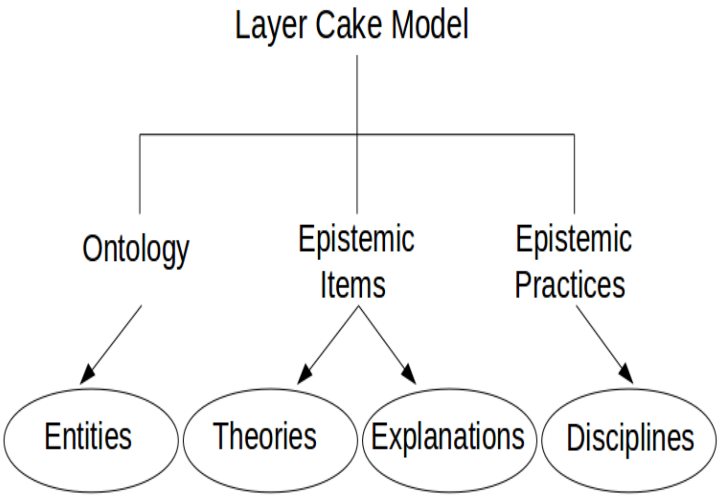

Fig. 1 Three variants of the LCM. Each variant is an interpretation of the LCM, differentiated according to what is being structured into levels in accordance with the six conditions of the LCM.

This feature of the LCM is something Daniel Brooks (2016, p. 52) calls “the correspondence thesis”, which he describes as, “a 1:1 relation between a particular level of organization and a specific scientific discipline that investigates this level.” We’ve seen the correspondence thesis in O&P’s explication above. Individual theories (epistemic items) collectively comprise a whole branch of science (an epistemic practice), the terms of which designate a ‘universe of discourse’. A given universe of discourse describes entities (ontological levels concept) who *qua* universe of discourse have at least one thing in common – they are at the same level of organisation in the model.

Brooks (2016, p. 46) argues that if you had to pick one or two defining features of LCM-style thinking it would be either stepwise compositional continuity[11] or the correspondence thesis. I agree that the correspondence thesis is an important part of the LCM, particularly embedded within it's proper context as a framework for the unity of science project. However, I will treat the correspondence thesis as an idiosyncratic part of O&P's explication rather than as an important aspect of the LCM. By this, I mean that we can consider the variants of the LCM in isolation from one another and still correctly identify them as LCM structures. The correspondence thesis is an important part of O&P's thinking, but it is not *necessary* to hold the correspondence thesis in order to be influenced by the LCM.

This decision is driven by my aim to understand why the LCM has maintained such philosophical currency whilst being plainly implausible from a scientific perspective. I'll argue in the next chapter that this is in large part because discussions of reduction & emergence – that is, discussion concerning the relationship between levels – have neglected issues of system structure and this has allowed the LCM to persist. However, in order to show this *and* to pick out important lessons for developing an alternative account of reduction/emergence, I'll need to consider the different variants of the LCM and their relationship to corresponding kinds of reduction and emergence. In sum, then, the correspondence thesis is undeniably a component in the LCM, but I'll be at backgrounding the issue to some extent from here on.

To make it straightforward to refer back to these key aspects of the LCM, I have summarised them in Box 2. A main result of this chapter has been to move from Box 1 to Box 2. This move consists in taking O&P's six conditions as a starting point, and transitioning to a full understanding of the LCM both in the context of the unity of science project and beyond that initial motivation towards a fully fledged analysis of levels of organisation with four key characteristic features.

11 Brooks means something slightly different by 'stepwise compositional continuity' than I mean by 'stepwise decomposition'. Included in his understanding of the phrase is that *all phenomena* are connected by step-wise part-whole relations (Brooks, 2016, p. 46, 2017, p. 146). This clearly makes sense in the context of O&P's model; after all, all phenomena are supposed to be contained within it because of its *global* scope. But I separate the ideas of step-wise decomposition and global scope out in order to isolate the essential tension between aspects of the LCM (see 1.5 below). All structuring concepts that build 'levels' require *iterative* processes of differentiation (i.e., step-wise decomposition) even ones that only contain a very small subset of individuals (i.e., with a very local scope of application). Failing to the separate the concepts of step-wise decomposition and global scope makes this point easy to miss.

(1) Global Scope: the system under consideration is the entire world
(2) Monism about the Levels Concept: there will only be one correct model of levels of organisation
(3) System Type: It is an aggregative system that is simply decomposable
(4) Distinct Levels: no entity 'properly' belongs to more than one level and all its parts are found at the next lowest level

Box 2 A summary of the key features of the LCM.

1.5 What's wrong with the LCM: An essential tension

Among contemporary philosophers who have directly engaged with the LCM, there seems to have been a resounding consensus that it is not a scientifically plausible account of levels of organisation (Brooks, 2017; Craver, 2007, pp. 172–176, 2015; DiFrisco, 2016; Eronen, 2015; Love, 2011; Potochnik and McGill, 2012; Rueger and McGivern, 2010; Waters, 2008). With that in mind, my task here is not so much to argue that the LCM is scientifically implausible, but to pick out the key criticisms that unite those rejections of the LCM. I should make an important clarification first. Not all the philosophers I just referenced explicitly reject the LCM, in name at least. Rather, they question the utility of 'levels of organisation' as a concept more broadly, or so it would seem. Brooks (2017), however, has developed a convincing argument – which he labels 'guilt-by-association' – to the effect that these criticisms are rooted only in rejections of key aspects of the LCM, and/or projects within which the LCM is deeply embedded, such as inter-theory reduction, rather than 'levels of organisation' per se.

Regardless of whether Brooks is right that rejections of aspect of the LCM *exhaust* the content of these critiques, it is undeniable that they reject the central tension of the LCM that I will develop here. They may reject more than this central tension, but for my purposes this will suffice to show both that and why that the LCM is a scientifically implausible account levels of organisation, in a way that I believe picks out a common thread of criticism on which all of the above would agree.

Oppenheim & Putnam's paper was published in 1958, but we can see accounts with subtle variations from the LCM around the same time. For example, Brooks (2016, pp. 184*ff*) carefully explicates the 'organicist' conception of levels of organisation (e.g., Needham, 1931, 1932), which pre-dates O&P's work and shows how their view contrasts with the LCM in a variety of ways, not least their pluralism about different inter-level relations. Additionally, Joseph Feibleman (1954) and Alex Novikoff (1945) developed *anti-reductionist* views of levels; explicitly designed to find a balancing point between 'atomistic' or 'mechanistic' views, and the 'holist' or 'organicist' views of Needham, for example. As I'll argue in the next chapter, anti-reductionism about the units of analysis within a system structure doesn't suffice to reject the LCM structure.

Nevertheless, these accounts contain interesting departures from the LCM. For example, Novikoff (1945, p. 209) argues for the existence of 'mesoforms', objects that exist only at the boundaries between two levels and so are, in some sense, present on more than one level of organisation (contra distinctness of levels). He argues that viruses and 'colonial organisms' occupy this sort of role. For example, a colony of bees sometimes functions in such a unified way that it's not clear whether it should be considered as a mere collection of individuals (a social group, L5) or as a unified multicellular organism (L4). Novikoff is suggesting that a colony of bees should be considered as residing at the edges of *both* levels and thus belonging 'properly' to neither one.

The question of how to conceptualise colonial organisms re-emergences in recent behavioural ecology under the guise of the 'supra-organism' hypothesis (Seeley, 1995). Feibleman (1954, p. 61) argues that "for an organisation at any given level, its mechanism lies at the level below and its purpose at the level above" (p. 61). To investigate an entity at one level, we must investigate how it functions in relation to its parts at a lower level *and* what it's function is, in relation to its place as a part on the level above. This foreshadows contemporary views on multi-perspectival or integrated modelling.

For example, mechanisms (in the contemporary sense) are argued to be *essentially* multi-level phenomena (Bechtel, 2008; Craver, 2015, more on this in 5.2 & 5.3).[12]

Feibleman & Novikoff's views already suggest that the LCM has problems with scientifically plausibility. When we get to Burton Guttman (1976), these problems are laid out more explicitly. We can use Guttman's essential complaint against levels of organisation to reveal the crux of the matter when it comes to the scientific plausibility of the LCM. Overall, Guttman argues that 'levels of organisation' is at best a useless concept in biology and, at worst, is actively misleading for biological pedagogy and inquiry.

Guttman considers two interpretations of 'levels of organisation', which focus on key aspects of the LCM: one *structural* and one *functional*. The structural claim is that "every system of level n, is made entirely and exclusively of systems at level n-1" (Guttman, 1976, p. 113). Guttman proceeds to give several counterexamples against this claim drawn from many levels of organisation. His examples include ecosystems, multicellular organisms, polymeric molecules, and tissues. A particularly vivid case of the examples I think Guttman has in mind is the extra cellular matrix (ECM). The ECM is defined as "a *non-cellular* component present within all tissues and organs" (Frantz et al., 2010, p. 4195, my emphasis) The ECM is composed of, amongst other things, two macro-molecules: Proteoglycans and Fibrous Proteins (elastin, fibronectin, and laminin). These macro-molecules are responsible for physical scaffolding of cells as well as the initiation of biochemical and biomechanical cues required for tissue morphogenesis, differentiation, and homeostasis. More to the point, Frantz et al. also state that, "each tissue has an ECM with a *unique composition* and topology […] Indeed, the physical, topological, and biochemical composition of the ECM is not only tissue-specific, but is also markedly heterogeneous" (2010 p. 4195, my emphasis).

12 There's a possible way that the LCM could accommodate examples of this sort. Recall that the levels given by O&P were one possible instantiation of the LCM, compatible with the addition of new levels if required. Perhaps O&P could simply have added a new level of 'supra-organisms'. On the other hand, a more damaging reading of Novikoff's objection would be that that no matter how many levels you add to the model, entities will be legitimately placed on different levels relative to different theories and explananda. Perhaps for explaining *certain aspects* of the bee colony's behaviour it will be important to consider the colony as a collection of individual organisms. But for others – like swarming, or homeostatic mechanisms of food, temperature, and water regulation, it will be more productive to think of the colony as a unified organism (Seeley and Tovey, 1994). So, colonial organisms have no 'proper' level at which they reside. This kind of 'level-shifting' is in tension with the distinctness of levels at a global scope.

The compositional heterogeneity of the ECM is damning for the LCM. Macro-molecules should not be appearing as the parts of *tissues* as they already appear as the parts of *cells*. Tissues should find their parts at the next lowest level of organisation, cells, and in turn cells should find their parts at the next lowest level of organisation, molecules. However, when considering the ECM, molecules appear as the parts of *both* tissues and cells. Further, tissues have *both* cells and molecules as their parts – parts that appear on *different* levels of the LCM. Molecules, then, appear to be freely moving up and down the LCM structure in stark contrast to the idea of distinct levels. If you can't maintain distinctness of levels, then this opens up the possibility of a plurality of ways in which a system might be structured in contrast to the monism aspect of the LCM.

Guttman's second point is a functional version of the first one, namely that: "Interactions between systems of level n are mediated through objects of level n-1 (or some other specific level less than n)" (Guttman, 1976, p. 112). Again, Guttman draws counter-examples from many levels of organisation. One such example is the complex web of interactions that 'mediate' multicellular organisms. According to the LCM, we should be investigating cellular interactions, and *in turn* molecular interactions.
But Guttman points out that cells interact *via* molecules such as hormones. Further, the immune system provides a site of complex interaction between cells, molecules, tissues, that is not structured into neatly defined aggregative levels in accordance with the LCM.

There's an objection to Guttman's points on the horizon here. There's at least an interpretation of the LCM available that might be able to accommodate free floating molecules, for example, *to some extent*. One could claim that what O&P meant by 'distinctness of levels' was that a whole at level n is an aggregation of parts that are all at *a* lower level, but not necessarily *the next* lowest level. In other words, we drop the 'stepwise' component of distinctness of levels in simple decomposition (that is, aggregative level-building). According to the LCM structure, both molecules and cells are at a lower level than tissues, so there's no problem here and the LCM structure remains intact. Far from resolving the issue, this potential response on behalf of O&P, will actually help us to draw out the deep incompatibility between the key aspects of the LCM and contemporary scientific inquiry.

To start with consider that removing the 'stepwise' component is equally, if not more, problematic for the LCM because it destroys the very *structure* that the LCM builds. The problem with removing step-wise decomposition is that it becomes harder to see how individuals are related in a 'higher-lower' sense, in the first place. In other words, what justifies the different *levels* of organisation themselves? The advantage of a step-wise decomposition is that the 'levels' develop in a straight-forward fashion. The levels are constructed and justified by the claim that everything at a given level is composed of nothing but parts at the next lowest level of organisation. The very structure of the system is thus constructed *using* this *step-wise* part-whole relation. But with the requirement being only that wholes are composed of parts at any lower level, how are the levels to be distinguished? If ecosystems have individual organisms and molecules as *parts*, in what sense are they on *different levels*?

In fact, this point is not localised to the LCM, it applies to the concept of stratified system structure generally. In order to make an initial (vertical) partition in a system we need to distinguish one level as lower than other. For the LCM this is done with aggregation, but regardless it needs to be done with *some differentiation* relation such that we have a whole and its parts. Now we have two levels. In order to *build more levels* in the structure we have to reiterate this process over and over again, reapplying our differentiation relation again every time we want to generate a new partition. This is perhaps an obvious point but it's easy to miss that any kind of 'levels-talk' beyond having a system with just one level *requires* stepwise decomposition. In this sense, step-wise is just another word for 'iterative'.

So, if we remove the requirement for step-wise decomposition in the LCM, we'll generate what I'll label a 'mereological collapse' – the collapse of a part-whole relation. We won't be able to differentiate the strata in the system from one another, and so we won't have levels. If we combine this requirement with two essential features of the LCM: aggregation and global scope, we reveal an essential tension at the heart of the LCM which renders it implausible as an account of levels for contemporary scientific practice. The tension lies at this very intersection between the need for step-wise level building, the relation of aggregation to build those levels, and a global scope for the model as a whole. If we want to represent structures as containing levels of organisation, it'll have to be done stepwise, otherwise there will be no discernable (and stable) structure in the model.

However, if stepwise decomposition is understood as *aggregation* of wholes into parts, then eventually (and pretty quickly) as you move towards a *global* scope you'll run into entities that violate stepwise composition and appear on several levels of organisation at once; of just the sort that Guttman is pointing out. In other words, looking at scientific practice immediately throws up cases in which stepwise aggregation will be violated when attempted at a global scope. Indeed, contemporary accounts of levels have heeded this warning of mereological collapse and can be understood as tweaking this problematic triumvirate of concepts. Specifically, each of the accounts *rejects* global scope and *replaces* aggregation with a different partitioning relation, in order to save some sense of structure through step-wise composition relations. It's worth reiterating that point for the sake of clarity here: my argument is *not* that there's something inherently problematic with step-wise decomposition, aggregation as a part-whole relation, or global scope, taken separately. My characterisation of levels will maintain step-wise decomposition, and make room for the part-whole relation of aggregation (within a pluralism about part-whole ordering relations). My argument here is that when you put them *together* as general features of the way you think about levels of organisation – as the LCM does – then you're going to run into scientific counterexamples pretty quickly and this explains why the LCM lacks scientific plausibility.

1.6 Conclusion

The main aim of this chapter was to pick out concrete recognisable features of the LCM through an analysis of it's canonical presentation; embedded within the larger project of the unity of science. The pursuit of this aim has lead to the development of Box 2 and the features contained within. An immediate result was that it allowed me to offer an argument as to why the LCM is scientifically implausible in terms of the tension between three of these key aspects; thus tackling at least one part of the perplexing juxtaposition contained within the LCM: it's pervasiveness and its scientific implausibility. In this next chapter I'll turn to explaining the other part through an analysis of why anti-reductionist positions concerning ontology and epistemology have had little impact on the LCM, allowing it to slip through the net into contemporary philosophical thinking about levels of organisation.

2

The Layer Cake Model, Reduction & Emergence: A Complicated Relationship.

2.1 Introduction

The story of the layer cake model is also the story of reductionism. As we've seen in the previous chapter, Oppenheim & Putnam introduced their account of the LCM alongside an accompanying account of reduction; both of which were to work together as the structure within which the Unity of Science project might be realised. This was not just because O&P happened to be interested in giving both an analysis of levels and accompanying account of reduction. Rather, they understood the deep connection between the concepts of 'levels' and 'reduction'. If we understand analyses of levels of organisation as accounts of the structural *features* of systems; features like those in Box 2, and reductionism as specific views about the inter-level relationships *within* that system structure, then the two concepts become inextricably linked. The main purpose of this chapter is to explore the connection between these concepts: the LCM as an analysis of levels, and both reductionist and anti-reductionist positions regarding inter-level relations.

As I argued in the previous chapter, the LCM is not one unified view of levels, rather it comes in three variants differentiated according to their unit of analysis. The ontological LCM focuses on ontological entities (individuals & their properties); the epistemological LCM focuses on epistemic items (theories, laws, and explanations); and the methodological LCM focuses on epistemic practices (inquires and disciplines). My first step in this chapter is to introduce three kinds of reductionism (2.2) that overlap with these units of analysis, allowing me to map possible connections between different variants of the LCM and kinds of reductionism (Fig. 2).

If this mapping represented *necessary* connections the road to dealing with the LCM would be clear. Reject each kind of reductionism that corresponds with a variant of LCM.

For example, a rejection of theory reduction would mean a rejection of the epistemological LCM. If we want to develop a new conception of levels that is decidedly anti-LCM, then we should stick to an anti-reductionism concerning inter-level relations. However, I'll demonstrate that this neat solution is entirely wrong and in doing so unpack the complexity of the relationship between levels and reductionism.

In section 2 I'll demonstrate that a specific kind of reductionism (Nagelian reduction) became closely associated with the epistemological variant of the LCM via a mutual association with the unity of science project (Fig. 3). I will then proceed to examine the contrasting positions developed in rejection of Nagelian reduction, pertaining to ontological items (2.4.1) and then epistemological items (2.4.2). The take away from all this will be that epistemological and ontological variants of the LCM remain compatible with both reductive and anti-reductive positions that overlap with their respective units of analysis.

In the final part of the chapter I will argue that the issue crucially turns on the overlapping units of analysis. I'll argue that despite surface-level differences in units of analysis, ontological emergence, and epistemological anti-reduction share a deeper connection that explains their inability expunge the LCM from philosophical analysis: a 'property-first' approach. This is an approach in which properties are, in fact, the primary units of analysis. Whether they're being explained and collected into theories and robust laws (epistemology); or whether they reflect the entities that bear them and the kind of causal relationships those entities can stand in (ontology).

The major output of the chapter will be the development of four important lessons uncovered by exploring these relationships. In short these are (1) rather than relying on an account of reduction *or* emergence as a strategy to avoid the LCM, we need a new account of reduction *and* emergence that is explicitly incompatible with the LCM. (2) That account will need to provide a clear sense of what it means for a system to have an 'aggregate' structure, as well as how and why aggregation fails. (3) That distinctness of levels and monism are mutually reinforcing aspects of the LCM. And finally (4) that the account requires an explicit position on issues of scope that are decidedly incompatible with LCM-style thinking. I'll conclude that these lessons will serve as desiderata for the account of reduction that I'll develop in the next chapter.

2.2 The LCM and Kinds of Reductionism

There are a couple of different ways that the concept of 'reduction' has been carved up in the literature. Probably the broadest and most common division is made into: epistemological reduction; methodological reduction; and ontological reduction (Ayala, 1974; Brigandt & Love, 2017; Sarkar, 1992; Wimsatt, 2007). Here I'll try to spell out a clear understanding of each. I won't be able to give an exhaustive characterisation of each kind of reductionism here, such a task constitutes a full project in itself (e.g., van Riel, 2014). Instead, I'll require that there be at least a clear and meaningful sense of what the *contrasting* view to that kind of reductionism would be: a 'contrast-requirement'. Without this I won't be able to clearly show that the LCM remains compatible with reductionist and anti-reductionist views alike. Let's start with the kind of reductionism about which it is hardest to meet that requirement: ontological reductionism. It's hard to meet the requirement because ontological reduction comes in several flavours; varying in the strength of claim they make.

The first, and weakest, is a commitment to *compositional materialism* or physicalism.[1] This is merely the claim that wholes and parts are fundamentally the same kind of 'stuff'. It can also be considered a commitment to substance monism – there's one kind of substance in the world out of which all things are composed. The second, more substantive view, is a metaphysical flavour of ontological reductionism in that it concerns the *causal efficacy* of wholes and their parts. In this sense, ontological reductionism pertains to the claim that wholes have no causal efficacy 'over and above' their parts.[2] Finally, the strongest version of ontological reductionism is *eliminative materialism*, in which the reduced entities/ properties are argued to be superfluous or merely useful ways of talking that do not track any 'real' properties (P. M. Churchland, 1981; P. Churchland, 1986). What unites these flavours together is their units of analysis: ontological items.

1 I'll just be treating 'materialism' and 'physicalism' as the same view. Sometimes 'physicalism' pertains to the objects studied by *physics*. Other times it is operationalised solely as a contrast to realism about 'mental states' (in a similar way to eliminative materialism). For detailed discussions on physicalism, its relation to materialism, mental properties, and to physics see Dupré (1993, chpt. 4); Loewer (2001); Papineau (2001); Stoljar (2010).

2 I'll be using 'causal powers' in a completely neutral way throughout this chapter. It can be understood as specifying changes in values of variables; intrinsic dispositions; counterfactual dependencies; probabilistic dependencies; physical mark transmissions; and any other causal notions of choice. Nothing I have to say here turns on a particular view of causation.

As van Reil (2014, p. 19) puts it, we're dealing with ontological reduction when the predicates involved refer to "kinds, types, properties, events, substances or individuals". I'll be sticking to individual entities and their properties.[3] This is primarily because the account of reduction that is of central importance to the discussions in this chapter – Nagelian reduction – concerns properties and their related bearers (individuals), rather than relations between events, substances, or truths.

I should add that the only assessment I'm offering here is in terms of meeting the contrast-requirement. None of what follows is intended as judgement on the plausibility of these flavours of ontological reduction. With that in mind, I'll be focusing on the metaphysical flavour. Compositional materialism is too weak to establish a proper contrast. Without a sense of what it means to be fundamentally a different 'kind', there's no clear sense of what it would mean to deny that wholes and parts are of one and the same fundamental kind. On the other hand, eliminative materialism is too strong *outside* of usage in philosophy of mind. In those debates the elimination of mental states seems a pressing issue and both it, and its contrast, are viable options, i.e., eliminativism or realism about 'mental states'. But outside of this particular debate, eliminativism is not usually on the table when discussing whole and parts. Rather, what is on the table is the metaphysical flavour of ontological reductionism. Finally, the metaphysical flavour adds the specificity that was lacking from compositional materialism for it to meet the contrast-requirement: wholes are of a different kind from their parts insofar as they have novel causal powers over and above those of their parts.

Epistemological reduction is a little more straightforward. Again, the unifying factor amongst different brands of epistemological reduction is the unit of analysis, this time epistemic items such as theories, laws, and explanations. Below we'll see how Nagelian reduction combines all three of these epistemic items. More often, accounts of epistemological reduction have focused on one of these items, primarily explanations. The relationship between epistemic items can be deductive; logical relations that link predicates from higher-level theories, laws, and explanations to lower-level explanations. They can also consist in showing that the higher-level explanation is explicable using only lower-level explanations, and higher-level theories can be shown to be 'limiting cases' of more general, lower-level theories.

3 Sometimes individual entities are also referred to as 'concrete particulars' to contrast them with abstracta like concepts, numbers, and so on.

Additionally, a reductionism about explanation can pertain to causal specificity, arguing that lower-level explanations contain more causal specificity, in virtue of which they are *always* preferable to higher-level explanations.

The final kind of reduction in my tripartite division is methodological reduction. I mention methodological reduction here only in passing to complete the introduction of reduction and won't be returning to it in much detail in the rest of the chapter. It will, however, become the main focus for the development of a new framework of reduction the next chapter. Accordingly, much more detail will be provided there. Methodological reduction focuses on reductive epistemic practices. In short, the idea is to move away from examining the reductive relationships between epistemic outputs like theories, laws, and explanations, towards analyses of what scientists do in practice, in the field, in the lab, and in complex computer simulations. Some examples of methodological reduction couched in terms of epistemic practices include a focus on reductive practices as heuristic strategies (Bechtel & Richardson, 2010; Wimsatt, 2006); problem-agendas (Love, 2008); research traditions (Laudan, 1977); and research directives (Lausen, 2014).

There's one final distinction that will be important in what follows: the distinction between synchronic and diachronic reduction/emergence. For simplicity, I'll explain this distinction just in terms of theories. Synchronic reduction concerns the reduction of one theory to another ahistorically. It concerns whether, for example, the predictions of a higher-level theory are entirely deducible from those of a lower-level theory, or the higher-level theory describes only causal properties that can be accounted for in the lower-level theory; at any given time. Diachronic reduction can be understood in two ways. The first is *successional*; when, in the history of science, one theory replaces another, or two theories are unified to create a new theory. An active area of debate concerning theory succession is evolutionary theory. In the early 20th century, the Darwinian theory of evolution as natural selection was fused with a genetic theory of inheritance to create the 'modern synthesis' (Dobzhansky, 1937; Mayr, 1942). The modern synthesis (sometimes called 'neo-Darwinian evolution') *replaced* Darwinian evolutionary theory. Currently, there are active discussions concerning whether the modern synthesis itself needs to be replaced with a framework that can incorporate, for example, niche construction, developmental plasticity, and epigenetics; dubbed 'the extended evolutionary synthesis' (Huneman & Walsh, 2017; Laland & Brown, 2011; Pigliucci & Müller, 2010).

The second form of diachronic reduction concerns the evolution of dynamics in a system; that is, the relationship between higher and lower-level phenomena over time (Bedau, 2008; Humphreys, 1997; O'Connor & Wong, 2005; Silberstein, 2006). Accordingly I'll label this form 'dynamic emergence' to keep it separate from the successional sense of diachronic In what follows I'll be dealing only with *synchronic* forms of reduction and emergence. Again, this is due to a focus on Nagelian reduction. As I'll argue shortly, my claim is that the LCM became closely associated with Nagelian reduction, itself a synchronic account. I'll then be attempting to show that rejections of Nagelian reduction – anti-reductionist and emergentist positions – remain compatible with the LCM. For now, I'll limit my focus to those anti-reductionist positions that are *not* founded on the basis of rejecting the synchronic aspect of Nagelian reduction, but I will return to diachronic forms of reduction & emergence in the next chapter.

2.2.1 Three variants of the LCM meet three kinds of reductionism

From the introduction of different kinds of reductionism it should be apparent that each kind shares something in common with a variant of the LCM: a unit of analysis. To make this explicit I have expanded the structure in Fig. 1 and mapped onto it the three kinds of reduction that correspond with the units of analysis of the variants of the LCM, the result of which can be seen in Fig. 2. If we're considering the LCM as pertaining to objects in the world then the question of reduction should also concern objects in the world; i.e., ontological reductionism. Similarly, if the LCM is understood as providing a hierarchical structure for epistemic items like theories or explanations, then the related notion of reduction also concerns the potentially reductive relationship between those epistemic items; i.e., epistemological reduction.

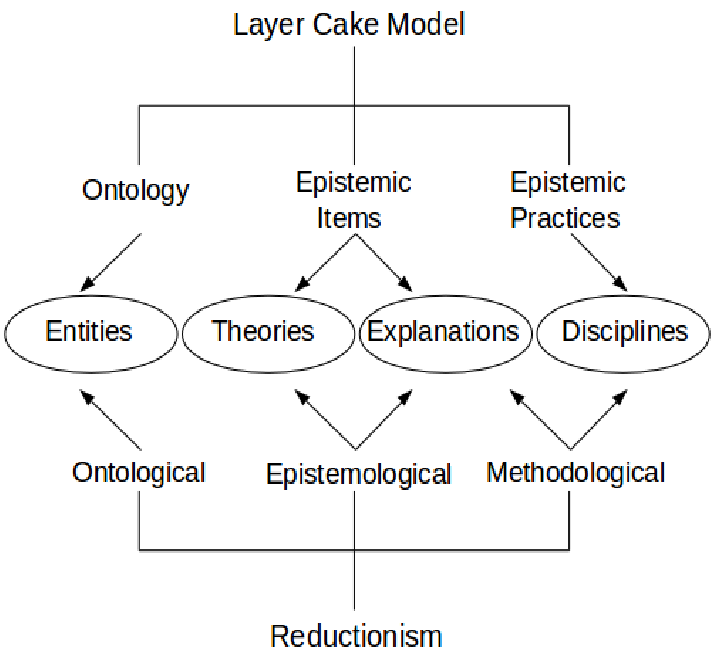

Fig. 2 Three variants of the LCM & Three Kinds of Reduction. The LCMs and Reductionisms are mapped onto one another via an overlap in focus e.g., on ontological items such as entities. The figure illustrates only possible compatibility between LCM variants and kinds of reductionism

The structure in Fig. 2 does not represent claims about how different variants of the LCM and different kinds of reductionism *must* be related. Rather, it outlines the possible connections; the compatibility between the variants of LCM and reductionism based on an overlap in their units of analysis. In 2.2.2 below I'll articulate the connection between the epistemological LCM and corresponding version of reductionism via Nagelain theory reduction. To get a general sense of how the LCM links with reductionism, let's take a look at the ontological versions. The ontological version of reductionism – interpreted in its metaphysical flavour outlined above – claims that any causal powers obtained by wholes are done so entirely in virtue of the causal powers of their parts. This yields a reductive inter-level relation for individuals and their properties. If I want to understand the role that, let's say, properties of cells can play in a system, I need to drop down a level of organisation and see what causal properties their molecular constituents have. Note that this is not a claim about *explaining* the causal properties of cells (although obviously it will be related to an explanation of those properties).

Ontological reductionism is merely a restriction on the ascription of causal powers to entities a different levels; whatever causal role I ascribe to cells cannot involve causal powers *not* found at the molecular level.

Ontological reductionism is neatly compatible with the ontological variant of the LCM. That the causal powers of entities at one level are to be fully constrained by the causal powers of entities at the level below makes sense in a structure that is *aggregative*. If cells are simply aggregated by molecular components then it stands to reason that the causal properties of cells are derived from those molecular components. Similarly, it is consistent with ontological reductionism that entities don't appear on more than one distinct level of organisation. Presumably, if entities could appear on more than one level, then the claim that the causal properties of entities at a given level are constrained by the causal properties of the *lower-level* constituents would be in jeopardy. Monism about levels is accommodated for similar reasons. Finally, whilst its not the case that ontological reduction *must* be set within a global scope, it certainly makes sense within such a scope: the causal properties of the occupants of one level to be found at the next lowest level. There's no obvious reason why, if this claim holds for one level, then it won't hold for them all except the lowest level of the system.

Methodological reductionism points to both explanations and strategies on Fig. 2. This is because work in certain forms of epistemological reduction has moved away from theories, laws, and explanations in the *Nagelian sense* I'll outline below, and towards explanatory *practices*. This work still often takes 'explanations' to be the unit of analysis but in a such a way that aligns much more with methodological reduction than epistemological reduction in the traditional sense. In the next chapter, I'll bring these two senses of reduction (and their anti-reductive equivalents) under the umbrella of 'practice-oriented analyses', in which epistemic practices can refer to either explanatory practices and/or research strategies. When both are practice-oriented in this sense, there is only a fuzzy boundary between 'epistemological' and 'methodological' reductionism. I have in mind examples such as mechanistic explanation (e.g., Craver, 2007; Craver & Darden, 2013; Machamer, Darden, & Craver, 2000) and Marie Kaiser's (2015) recent account of reductive explanation in biology. So, the 'explanations' bubble actually represents two different kinds of explanation *qua* unit of analysis: a more traditional form of explanation, as we'll see below, and a more practice-oriented form of explanation, which I'll discuss in much more detail in the next chapter.

2.2.2 Nagelian Theory Reduction

Fig. 2 can help make some headway on the relationship between the LCM and reductionism in the history of philosophy of science. Oppenheim & Putnam – perhaps unsurprisingly – suggested Kemeny & Oppenheim's (1956) account of reduction for the unity project but this is neither the account of epistemological reduction that became associated with the unity project, nor the prominent account around which discussions of epistemological reductionism were oriented in recent philosophical history (mid-late 20th century). That account was Ernst Nagel's (1979) version of epistemological reduction. My claim will be that the epistemological LCM and Nagelian reduction became closely intertwined via their mutual relationship to the unity of science project. So closely, in fact, that we see the LCM sometimes referred to as an account of reduction in its own right (Kincaid, 1990, p. 576; Steel, 2004, pp. 60–61). This is an important switch because Kemeny & Oppenheim's account was primary focused on diachronic theory replacement. They suggest that "an especially important case of [reduction] is the replacement of an accepted theory (or body of theories) by a new theory (or body of theories) which is in some sense superior to it. Reduction is an improvement in this sense" (Kemeny & Oppenheim, 1956, p. 7). Nagelian reduction, however, is most often understood as providing an account of synchronic ahistorical reduction relations between theories. Nagelian theory reduction was arguably the most central account of reduction in mid-late 20th century philosophy, as a consequence, critiques of reductionism were also focused on synchronic accounts rather than theory replacement.

I should be clear that I'm not claiming that Kemeny & Oppenheim's account literally *can't* make sense of synchronic reduction, nor that Nagelain reduction *can't* make sense of diachronic reduction. Indeed, this couldn't be right because as we saw in Chapter One (1.4.2), O&P explicitly argued that the levels in the LCM are *stable* even after reduction; they are not *replaced* by successional lower-level theories. Yet, O&P deploy Kemeny & Oppenheim's view of reduction. Presumably, they must have done so with a synchronic reading in mind, so the view must be amenable to both. Nevertheless, my claim here is about the perception of these accounts, and thus the roles they have played, throughout the history of recent philosophical discussions. I think it's fair to say that Kemeny & Oppenheim's view is more characteristically associated with diachronic reduction, while Nagel's account is more strongly associated with synchronic reduction.

More important for my purposes is that the LCM overlaps with Nagelian reduction at least partly in virtue of a synchronic view of reduction. To see if that's right, I'll start by laying out the basics of Nagelian reduction.[4]

Nagel is primarily concerned with what he labels 'heterogeneous reductions' meaning that there are terms that appear in the reduced theory that do not appear in the reducing theory (e.g., the term 'temperature' does not appear in the kinetic theory of gases, 1979, p. 342). Nagel provides two formal conditions that must be met for one theory to be reduced to another, these are (1) The Condition of Derivability and (2) The Condition of Connectibility. (1) states that the laws in the reduced science must be a logical consequence (must be derivable) of some of the theoretical assumptions contained in the reducing science.

The condition of connectibility is otherwise known as a requirement for 'bridge-laws' and is a direct result of the reductions being 'heterogeneous'. The thought is that because there will be no direct equivalent term in the reducing theory for us to map terms from the reduced theory to, we will have to develop intermediary laws that contain terms from both theories and thus 'bridge' the gap between them. Jerry Fodor's (1974, p. 100) reconstruction of bridge-laws helps illustrate this more clearly:

(1) $S1x \rightarrow S2x$

(1) is supposed to be read as 'All S1 situations bring about S2 situations', where S is a predicate of the special sciences and not a predicate of basic physics. The reduction then needs the following 'bridge-laws':

(2a) $S1x \leftrightarrow P1x$

(2b) $S2x \leftrightarrow P2x$

And finally a law in basic physics of the form

(3) $P1x \rightarrow P2x$

4 Some have argued for a distinction between the 'official' version of Nagelian theory reduction and the 'real' version. The 'official' version is the one referred to in most philosophical literature and captures how Nagelian reduction has *become* understood in the history of philosophy of science. The 'real' version is developed from careful exegesis of Nagel's work which, it is argued, differs in important ways to its less faithful counterpart (van Riel, 2014, pp. 157–162). To be clear then, I am working with 'official' version, rather than the 'real' version, precisely because I am interested in the role that the account played in the history of philosophy of science, even if that account diverged from Nagel's original intentions.

The two bridge-laws connect the law-like statement from (1) to its counterpart in basic physics by mapping specific predicates from each law onto one another. Let's plug in a quick toy example to get clear here. First let 'S1' stand for 'green eyes' and 'S2' stand for 'brown hair' and the variable 'x' stand for 'humans'. Then (1) would be the claim that humans who have green eyes, will have brown hair. Next, let 'P1' stand for 'Gene_{GE}' and 'P2' stand for 'Gene_{BH}'. Now (2a) can be read as "for humans, 'green eyes' and 'Gene_{GE}' always co-occur'. Likewise, (2b) reads "in humans, 'brown hair and 'Gene_{BH}' always co-occur." Finally, we have the basic law "humans who have Gene_{GE} will also have Gene_{BH}".

According to Nagel's account, we have now *reduced* the law about phenotypic properties of green eyes and brown hair, to a (Mendelian) genetic law. Furthermore, we can see the stepwise nature of reduction by adding a new law-like statement into the mix. Suppose we added:

(4) DNA1x $\rightarrow$ DNA2x

DNA1x stands for a specific sequence of nucleotides[5], and DNA2x stands for a different sequence. If we iterated the process using (4) and (3) to construct new bridge-laws we could *reduce* the law stating the connection between instances of 'green-eyed humans' and 'brown-haired humans' with a law connecting two specific sequences of nucleotides. If we suppose that the laws we've been working with are components in theories – say, a phenotypic theory of eye colours, a Mendelian theory of genetics, and a molecular theory of genetics – then we have partaken in some *theory reduction* under the Nagelian account.

The Nagelian account also provides an account of explanatory reduction. The relationship between the reduction of theories and explanations can be understood in two ways. Firstly, Nagel's account simply doesn't recognise a distinction between the two, such that to explain a particular law from one branch of science *just is* to reduce it to a law from a more fundamental branch of science. In the toy example I've been working with, exhibiting the logical structure of the relationship between instances of Green-eyed humans and the carriers of Gene_{GE} is to explain – for example – the strong correlation between green-eyed humans and brown-eyed humans.

5 i.e., a combination of A-T C-G base pairs in some specific order, bound by a sugar phosphate. In other words a *molecular* understanding of 'Gene'. The examples of Gene_{GE} and Gene_{BH} are supposed to stand for a *Mendelian* understanding of 'Gene', which are defined in reference to an associated phenotypic trait such as 'green eyes' or 'brown hair'.

A second way to think about Nagelian theory reduction and explanatory reduction is to explicitly relate it to the deductive-nomological (or covering law) account of explanation (Hempel & Oppenheim, 1948). With the logical structured exhibited between phenotypic properties and genetic properties, we can do some formal re-arranging to produce a D-N explanation such as:

1. C1: This human has Gene_{GE}
2. L1: All humans who have Gene_{GE} have Gene_{BH}
3. L2: All humans who have Gene_{BH} have Brown Hair
4. Explanandum: Therefore, this human has Brown Hair.

The argument fits the criteria for an explanation under the D-N account.[6] I've used a statement from a lower-level theory (C1), a law of a lower-level theory (L1) and a bridge-law (L2) to explain a higher-level statement (4). So, this argument is an explanation that is also reductive.

2.3 The LCM, Unity, and Reduction

So how did the LCM become almost synonymous with the sort of reductionism just laid out in Nagel's account? My proposal is the specific links that both share to the unity of science project. For start recall Hacking's (1996) subtypes of metaphysical unity:

A) 'Interconnectedness' - all kinds of phenomena must be related to one another
B) 'Structural' - there is a unique, fundamental structure to the truths in the world
C) 'Taxonomic' - there is one fundamental, ultimate way of system classifying everything: nature breaks down into natural kinds

Nagel's account of epistemological reduction can be seen as complementary to these metaphysical theses; a way of putting those theses to work in the unity project. If we assume that 'truths' in the world are encapsulated in scientific theories, which themselves are codified in terms of laws, then Nagel's account gives us a way to build and investigate that structure.

6 Technically it doesn't meet all of Hempel & Oppenheim's (1948, pp. 137–138, 1948) conditions for a D-N explanation, mostly obviously it doesn't meet R4 – that the statements in the explanans must be *true*. But this is just a toy example to illustrate the relationship between theoretical and explanatory reduction and hopefully it suffices in that sense.

If we further assume that those theories and laws capture phenomena in the world, then we also get a way to relate all phenomena in the world to one another, at least in principle. Finally, an issue I will return to in 2.3.1 below, the relata of Nagel's laws are supposed to be *natural kinds*. This is largely because 'scientific laws' were thought to represent stable relations between natural kinds rather than relations between gerrymandered or artificial types. Thus, by providing an account that is based on laws (or law-like relations) between natural kinds, Nagel's account is amenable to the taxonomic unity mentioned by Hacking.

As noted in the last chapter, O&P motivation's for the developing the LCM was to provide a framework within which a unity of science project might unfold – a structure that is conducive to the aims of the project. On top of this point consider that the LCM is clearly *compatible* with Nagelian theory reduction. We can quite easily swap out Kemeny & Oppenheim's account for Nagel's account without changing O&P's main points about the LCM itself (i.e., Box 2). As noted above, Nagel's account was far more prominent in 20th century philosophy and so by being amenable to interpretation under the Nagelian account, the LCM prolonged its lifespan. Nagel's account can be seen as adding specifics about the relationship between levels, and how an overall reductionist meta-scientific project like the unity of science might take place. In the toy example I've been working with, I've essentially been operating with laws that might form parts of theories spanning O&P's proposed levels L5 (multicellular living things) to L3 (molecules).

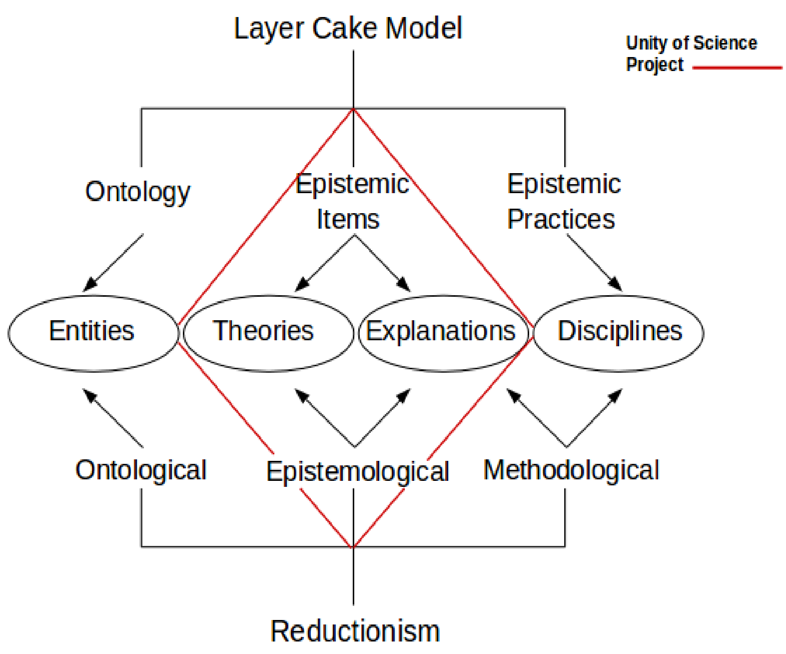

Fig. 3 The LCM & the Unity of Science. The diamond shape represents the unity of science project and encompasses a particular form of reduction & variant of the LCM.

This claim can be illustrated by seeing how the LCM pertaining to epistemic items and Nagelian theory reduction can be subsumed under the unity of science project (Fig. 3). When laying out arguably the most developed analysis of the LCM, Oppenheim and Putnam were explicit about their motivation being to provide a framework for the unity of science project. The unity of science project largely revolved around the co-ordination of scientific inquiry via exhibiting the structural (logical, if you like) relationship between epistemic items of scientific inquiry such as laws, theories, and explanations. The most developed and sophisticated account of how that structural relationship might be realised became Nagelian reduction – both in the sense of theory & explanatory reduction. On top of this overlap, the LCM and Nagelian reduction seem perfectly compatible as an account of levels on the one hand and an account of the reductive relations between those levels on the other. So by mutual association with the meta-scientific unity of science project – understood in a specific way – the LCM and reductionism seem to have become almost inseparable ideas.

Inseparable ideas perhaps, but are their fates tied together? That is, did the decline in enthusiasm for the unity of science project in 20th century philosophy precipitate a corresponding rejection of the LCM? Before that claim can be interrogated and, ultimately, rejected we'll need to see how the contrasting views – i.e., anti-reductionist positions – developed. That starts with seeing how Nagelian reduction has been criticised.

2.3.1 Rejecting Nagelian Reduction

A major set of critiques of the Nagelian account centred on problems with fulfilling the technical requirements of the proposed reductions; not least the requirements involved in developing 'bridge-laws'.[7] The main argument against the plausibility of bridge-laws involved the *multiple-realisability* of functional or 'higher-level' properties of the sort found in the 'special sciences' (Fodor, 1974; Hull, 1972; Putnam, 1975). Take, for example, the functional type 'money'. 'Money' could be made of (realised by) paper, plastic, metal, 1's & 0's, or even sheep – it doesn't really matter as long as it can fulfil the *functional role* of money in a transaction. This fact makes building bridge-laws very problematic.

The argument runs that: because these sorts of functional types are very prevalent in sciences such as biology, sociology, and economics, if the reductions of theories, laws, and explanations from those disciplines to 'more fundamental' sciences (supposed to be physics and chemistry) could possibly take place, then we'd need to build bridge-laws that connect those functional types to their physical *realisers* to be found in theories from lower-level sciences. But, due to multiple realisability, the best one could hope for is a bridge-law that connects one functional type – like 'money' – to a massively disjunctive collection of lower-level realisers: 'paper v plastic v metal v sheep v ...'

This objection doesn't apply only to functional properties, but also holds for properties that have 'one-many', 'many-one', or 'many-many' relations (Hull, 1972). We can find this sort of case in the toy example developed above. The property 'green eyes' is not usually thought of as a *functional* property with many different physical realisers.

7 Other problems included the scope of the Nagelian model being too restrictive due to its focus on theories and the lack of fully-fledged theories in sciences such as biology (Darden & Maull, 1977; Hull, 1974). Additionally, Nagel's account was criticised precisely because it *neglected* theory replacement, particularly the corrective aspect of historical cases of theory succession. This point was a driving force being Schaffner's (1967, 1993, 2006) 'General Reduction Replacement Model' of reduction which posited that reductions produce corrected analogues of the original theory and (as he argued in later work) patchy, fragmentary reductive explanations.

Rather, having the property of ‘green eyes’ can co-occur with a whole host of different combinations of ‘Genes’. In the example above, the bridge-law connected ‘green eyes’ to ‘Gene_{GE}’. The argument points out that we can see ‘green eyes’ instantiated at the same time as a whole bunch of different genes such that any possible bridge-law could only connect ‘green eyes’ to a massively disjunctive collection of ‘Gene_{GE1} v Gene_{GE2} v Gene_{GE3} v Gene_{GEn} ...’. What’s wrong with massively disjunctive bridge-laws, you might wonder? Massively disjunctive ‘types’ are not natural kinds traditionally understood, and putative laws that contain such types are not, themselves actually laws, traditionally understood. As Fodor (1974, p. 110) summarised, “either some of the generalizations to which the laws of special sciences reduce are not themselves lawlike, or some laws are not formulable in terms of natural kinds.” Either option was unpalatable for the proponent of Nagelian theory reduction and, by fiat (supposedly), the unity of science project more broadly.

A prominent proposal for circumventing the technical problems with fulfilling Nagel’s conditions for reduction was to swap out the relation of reduction with the relation of supervenience. Supervenience was a notion imported from metaethics.[8] In its original use, it operated as a coherence constraint between normative judgements (or facts) and descriptive (non-normative) judgements or facts. The idea being that if two situations were descriptively identical then it follows from supervenience that normative judgements about them must also be identical. Supervenience is asymmetrical such that normative judgements ‘depend’ on the non-normative descriptive facts, but not vice-versa. The asymmetrical dependence relation was put to use in accommodating multiply-realisable properties; the idea that two different types of property can be distinct yet one type remains wholly dependent on (or ‘fixed by’) the other. The property, x, is wholly dependent on the other, y, in the sense that there can be no difference in x without a difference in y. Returning to the toy example, we can say that the property ‘green eyes’ supervenes on the property $\text{Gene}_{GE,}$ when there can be no change in the property of having green eyes without a change in the property of having Gene_{GE}, yet the properties ‘green eyes’ and ‘Gene_{GE}’ are not identical (i.e., reducible) properties. Note how this picture is perfectly consistent with the multiple-realisability of the property ‘green eyes’.

8 Jaegwon Kim (2005, pp. 54–55), who has presented arguably the most detailed work on supervenience, attributes the origins of supervenience to G. E. Moore (1922) in outline and R. M. Hare (1952) in substance. The concept was exported from ethics into philosophy of mind in Donald Davidson’s (1970) famous paper ‘Mental Events’ and into Aesthetics by Frank Sibley (1959).

That ‘green eyes’ supervenes on Gene_{GE} means only that if some token Gene is present, say Gene_{GE1}, then ‘green eyes’ will be instantiated too, but if Gene_{GE1} is *not present* this in no way rules out the property ‘green eyes’ being instantiated anyway by some other ‘Gene’. Thus, supervenience fixes dependence relations among properties whilst still affording the multiple-realisability of those properties.

Finally, supervenience fixes a dependence relation between types but does not rule out identity between *tokens*. On a case-by-case basis, one particular person’s property of having green eyes may well be reductively explained by – for example – their having a particular Gene_{GE} token, say Gene_{GE2}. Supervenience merely allows that another person’s having green eyes may require an explanation to a different Gene_{GE} token than Gene_{GE2} due to the multiple-realisability of the property of having green eyes. As a broad non-reductive stance, this resultant view is sometimes referred to as ‘non-reductive physicalism’; ‘supervenience physicalism’ and/or ‘token-token physicalism’.

One path taken by philosophers convinced by these kinds of critiques has been to explore and refine successor accounts of reduction. Sometimes these views are labeled ‘new wave reductionism’ whose proponents include Bickle (2006); Churchland (1986); and Rosenberg (2006). A central component of these positions is to develop accounts of reduction that do not even require the bridge-laws that proved so problematic for the Nagelian account. Another path taken has been to explore the viability of various forms of *anti-reductionist* positions that attempt to utilise the supervenience relation to yield a coherent picture about the relationship between higher and lower level phenomena. These positions fall under the label of *emergence* or *anti-reductionism*. It is to these positions that I now turn with the aim of demonstrating how, whilst this path might lead away from reductionism, it doesn’t straightforwardly head towards new conceptions of levels of organisation.

2.4 The LCM, Emergence & Anti-reductionism

The concept of ‘emergence’ has seen somewhat of a resurgence in recent philosophical debates (Bedau & Humphreys, 2008; Clayton & Davies, 2006; Gillett, 2016; Humphreys, 2016; McGivern, 2015). Prior to this, emergent views were more closely associated with what is referred to as ‘British Emergentism’ (Alexander, 1920; Broad, 1925; Mill, 1882 Book III; Morgan, 1923).

British emergentism is supposed to a much stronger brand of emergentism than is discussed today.[9] I can't deviate from my aims in order to produce a full taxonomy of concepts of emergentism. However, thanks to the contrast-requirement, I can outline kinds of emergence that clash directly with the kinds of reductionism discussed so far. Accordingly, I'll discuss ontological emergence and epistemological anti-reductionism. Again, I'll be leaving issues of methodological approaches until the next chapter.

Two clarificatory points before proceeding. Firstly, I'll be discussing ontological emergence as a contrast to ontological reduction, but I'll be focusing on broadly *anti-reductionist* positions on epistemic items rather than 'epistemological emergence'. This is because epistemological emergence has come to mean something more than *just* a rejection of reductionist views of epistemic items like theories and explanations. For example, Mark Bedau's (2008) explication account of 'weak emergence' refers to a package of commitments; a balance of ontological reductionism and epistemological anti-reductionism; which *together* comprise an account of 'epistemological emergence'. Discussions of weak emergence are oriented around this balance and how a coherent picture can emerge through a reconciliation of this combination of commitments. What's more, weak emergentism usually pertains to a *dynamic* version of epistemological emergence rather than a *synchronic* version.

Given that my aim is to show that a mere rejection of the sort of synchronic reduction about epistemic items associated with the LCM (i.e., Nagelian reduction) fails to yield a view of inter-level relations that is incompatible with the LCM, I'll limit my discussion to epistemic anti-reduction, rather than 'weak' or 'epistemological' *emergence*. That is, a commitment to the idea that sometimes epistemic items, like theories, laws, and explanations, *cannot* be reduced. Furthermore, that this is no bad thing; sometimes the most apt explanation, for example, will be non-reductive. From here on out, then, when I refer to 'emergence' I refer only to the sense of ontological emergence that I'll spell out directly below; a metaphysical version of ontological emergence.

9 Although there are clear links between British emergentism and contemporary views of emergence. For example, Mill's notion of 'heteropathic' laws and effects pertains to laws which violate the composition of causes, that is, when the effect is *more* than would have been produced by the aggregation of causes taken individually. Stated this way, Mill's hetropathic laws seem to align very closely with violations of the superposition principle, which is itself a key indicator of a dynamic and non-linear system; the sort of system seen as a prime candidate as exhibiting emergent behaviour. See McLaughlin (2008) for a comprehensive overview of British emergentism

Secondly, I should reiterate the strength and purpose of my claim here. I am not suggesting that merely rejecting a reductionism about epistemic items, for example, commits one to a view of levels of organisation like the LCM nor that any proponent of such a view would advocate the LCM. I strongly suspect that they would not. What I am attempting to demonstrate is that the rejection of epistemological reductionism *alone* is not enough to expunge LCM-style features of the system as a whole. Further, to do this in such a way that we can discern some important lessons about inter-level relations that make sure that incompatibility with the LCM is explicit.

2.4.1 Ontological Emergence

The issue of what ontological emergence amounts to – including whether it is even a coherent position – is an active area of debate (Barnes, 2012; Humphreys, 2016; Kim, 2006; Silberstein & McGeever, 1999; Taylor, 2015). At the very least, there's a clear sense in which if it amounts to anything, it must start with the rejection of ontological reductionism. In 2.2 I argued that compositional materialism doesn't provide a clear enough contrast to emergence and we can now see why. The move from inter-level relations of identity found in Nagelian-style bridge-laws, to inter-level relations of supervenience doesn't yield a contrast with compositional materialism. That two properties are related by supervenience is perfectly compatible with them being of the same 'fundamental kind'. A whole being supervenient on its parts is likewise perfectly compatible with the view that the whole is exhaustively constituted by (is 'nothing over and above') its parts. Because of this, if ontological reductionism was understood merely as a commitment to compositional materialism, ontological emergence built on a conception of supervenience relations, would turn out to be just a variant of reductionism and not a view of 'emergence' at all.

However, inter-level supervenience is not so obviously compatible the metaphysical version of reductionism. Recall that metaphysical reductionism is the view that wholes have no causal efficacy over and above the causal powers of their parts. Put in terms of supervenience, this can be understood as a commitment to causal fundamentalism: that the causal powers of higher-level properties supervene on and are thus *determined by* lower-level properties. So, metaphysical emergence can be understood as a rejection of causal fundamentalism. Higher-level properties supervene on their lower-level constituents, but they have distinct causal powers.

As Wilson puts it, “token higher-level feature S has, on a given occasion, at least one token power not identical with any token power of the token lower-level feature P on which S synchronically depends, on that occasion” (Wilson, 2016, p. 356).[10]

This helps to makes sense of how higher-level properties can be supervenient on lower-level properties yet still be ontologically emergent; they have novel causal powers not conferred on them *in virtue* of their supervenience bases (the properties upon which they supervene). In this sense, ontological emergence can be understood as the combination of the following commitments:

(a) All higher-level properties supervenience on lower-level properties. But,

(b) Some, higher-level properties have causal powers that are not determined *solely* in virtue of their supervenience bases.

Does a commitment to (a) and (b) have an impact on the LCM? To see, we need to find an area of philosophy in which this position has been actively debated. The most obvious candidate is a debate that occurs in a particular strain of analytic philosophy of mind. In that debate the candidate ontologically emergent properties in question are mental properties, which for present purposes can be understood simply as propositional attitudes or intentional properties (beliefs, desires, and so on). Accordantly, I’ll run through the debate as it has manifested in philosophy of mind.

Touching on philosophical discussions of the relationship between ‘the mental’ and ‘the physical’ can be a tricky business as it is an area of philosophy rich in entrenched assumptions and esoteric concepts. However, dipping into this debate is a consequence of wanting to show that the *canonical* development of the debate about metaphysical emergence does not impact the plausibility of the LCM. I could have adopted a different example such as the genotype/phenotype toy example I used above. But this would be to gerrymander the development of the discussion. Instead, we’ll need to have a look at the discussion *in situ*. However, the point I’ll arrive at as a result of examining this debate is generalisable beyond the discussion concerning the status of *mental* properties.

10 Wilson calls this the ‘new power condition’ and goes on to define a conception of metaphysical emergence based on the condition. Wilson labels this view ‘strong metaphysical emergence’ and develops a modified version, ‘weak metaphysical emergence’ which she defends as her own position. The labels ‘strong’ and ‘weak’ are being used in a different sense to Bedau’s sense of ‘strong emergence (metaphysical) and weak emergence (epistemological).

It will concern the consequences of taking a property-first approach more generally and how such an approach is not conducive to investigating issues of system structure, not least structure developed in terms of levels of organisation (see 3.4 below)

The necessary requirement for illustrating the credentials of mental properties as emergent is to demonstrate that they can have downward effects. As Kim (2006, p. 548) suggests:

> "Emergentism cannot live without downward causation but it cannot live with it either. Downward causation is the raison d'être of emergence, but it may well turn out to be what in the end undermines it."

Why would downward causation be the crux of matter for emergent properties? The reason has to do with the specific combination of holding (a) and (b). To see this we'll need a quick run through of 'the causal exclusion problem'. The causal exclusion problem comes in many forms (Bennett, 2007, pp. 324–328). I'll be drawing on the most prominent version of the argument, presented by Jaegwon Kim (1998, 2005). The argument comes in two stages, the first denies the possibility of mental to mental causation. The second stage proceeds to deny the possibility of mental to physical causation – i.e., *downward causation* – leaving mental properties with an, at best, epiphenomenal status. The argument requires a few principles to get off the ground. The first is that the mental supervenes on the physical. This is a straightforward consequence of holding that (a) all higher-level properties supervene on lower-level properties.

The second principle is the 'principle of causal exclusion'. The causal exclusion principle posits that for any given event there can be no more than one distinct cause that is wholly responsible for the occurrence of that event, apart from in cases of 'genuine' over-determination. Rare instances of over-determination are presumed to be plausible enough, death by firing squad being an often-cited example, but the exclusion principle constitutes a constraint that over-determination is not *systematic*. The principle is supposed to capture the idea that the effects of mental causes could never be genuine cases of over-determination, or at least the conclusion that they are, would be wholly unsatisfactory (Bennett, 2007, p. 325). The final principle, 'the causal closure principle', states that if a physical event has a cause, then it has a physical cause (Kim, 2005, p. 15).

Again, this principle is supposed to be plausible, particularly in the conditional form presented here, as it makes no claims about the causal relationship between non-physical events, and allows for the possibility of physical events that have no cause.

The argument then runs as follows (Fig. 4): we start with the supposition that one mental property, M_1 causes the instantiation of another mental property M_2. Because of (a), M_2 must have a distinct physical supervenience base, P_2 upon which the instantiation of M_2 depends. As Kim notes, “Given that [P_2] is present on this occasion, [M_2] would be there no matter what happened before; as [M_2]’s supervenience base, the instantiation of [P_2] in and of itself necessitates [M_2]’s occurrence at t” (Kim, 2005, pp. 39–40). At this point it looks as though the instantiation of M_2 is guaranteed by two distinct events: its cause M_1 and its distinct supervenience base P_2. But, by (b) only one of these events can be responsible for the instantiation of M_2. The claim here is that it must be P_2 that is responsible because regardless of the occurrence of M_1, the very occurrence of P_2 necessitates the instantiation of M_2, thus the role of M_1 seems superfluous. This completes the first stage of the argument: mental properties do not cause the instantiation of other mental properties.

As for the second stage, consider that there may yet be a role for M_1 to play as long as something caused P_2. Let’s suppose that the cause is in fact M_1. If M_1 is the cause of P_2 then a causal role is preserved for mental properties – they cannot stand in causal relationships to other mental properties directly, but only indirectly through causing the supervenience bases of mental properties. In other words, in order to *be* causally efficacious, mental properties must cause the supervenience base of another mental properties. But this won’t work either because of (c). If we are supposing that P_2 does have a cause, which *ex hypothesi* we are, then, because of (c), P_2 must have a physical cause, P_1 and once again M_1 is excluded from playing a causal role in this story.

To summarise, we start by characterising emergent properties as both (a) supervening on lower-level properties and (b) *not* inheriting their causal powers from those lower-level properties. If (b) is correct then they can cause other properties to be instantiated *de novo* (not *qua* their supervenience bases). If (a) is correct, however, then that putative effect property will also have a supervenience base in virtue of which its instantiation is necessitated. So, if our strongly emergent property is to be involved in the instantiation of the effect property then it has to cause its *supervenience base*. This is why metaphysical version of ontological emergence requires downward causation.

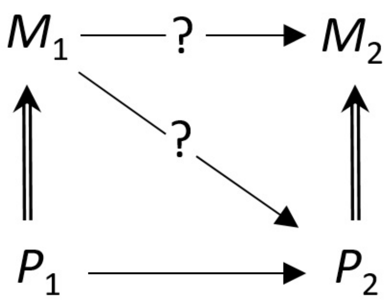

Fig. 4 The Causal Exclusion Argument. The arrows interrupted by question marks between M_1 and M_2 represent the first state of the argument and likewise between M_1 and P_2 for the second stage.

There's a lot of literature that discusses the fine-details of this argument (Kim, 2005, gives a substantial overview of this literature) and it is definitely not my intention to wade into that debate.[11] Rather, I can use the argument to show that adopting emergence has little effect on the LCM as the structure within which emergent properties might be situated. To start with, let's suppose we dropped the principle of causal closure – this is, in effect, what adopting emergence means. Seeing as downward causation is argued to be the crux of the matter for emergent properties let's also jump to the second stage of the argument.

Without the principle of causal closure, we can agree that M_1 has two potential causes, and that one of them has to go (as per causal exclusion). However, now we're free to argue that it is P_1 that has to go and not M_1. M_1 is the cause of P_2; we have an instance of genuine downward causation and an ontologically emergent property. Would this affect the claim that M_1 and P_2 reside on two distinct levels of organisation and, in virtue of this, could not reside on any other level? I don't see how it would. M_1 doesn't 'jump' levels by having a causal effect on a lower-level property, just as lower-level properties don't jump a level in cases of upward causation. Neither global scope nor monism are affected by the claim of downward causation alone.

11 The latest reboot of this debate can be found in the literature on the interventionist theory of causation. The main protagonists being Michael Baumgartner (2009, 2013), who offers causal exclusion arguments against interventionist higher-level causation, and James Woodward (2015) who defends it. Plenty of others have contributed to the debate on either side (Gebharter, 2015; Raatikainen, 2010; Yang, 2013).

The key to all of this is the *structural role* that supervenience is playing as an inter-level relation. Discussions surrounding emergence concern the potential of causal effects to happen upwards and downwards, so there has to be an 'up' or 'down' direction in the first place. Emergence, at least understood synchronically, doesn't remove the main inter-level relation being *dependence* or *necessitation* and, because of this, the LCM remains intact. We can already see this commitment in the first step of the causal exclusion argument. There, the supervenience relation between P_2 and M_2 excludes the potential causal relationship between M_1 and M_2. This makes it explicitly clear that *structure* is being understood as *non-causal metaphysical* dependence, not causal dependence. If the structure of the LCM was built via inter-level *causal* relations or if emergence denied that there is metaphysical dependence between levels, then the situation might be different. But the LCM is not built by inter-level causal relationships and emergence does not deny inter-level metaphysical dependence relations. Activating for the metaphysical flavour of ontological emergence is merely to argue that causal interactions can cross-cut these *non-causal* dependency relations; a view that can be accommodated within an LCM structure.

You might be wondering what happened to aggregation, the fourth aspect of the LCM, could this be a way out for strong emergence? I don't think so. Firstly, supervenience holds between *properties*, the entities that instantiate those properties remain fully aggregative even under emergence; no problem there. Secondly, supervenience relations are hardly *non-aggregative*. Rather, the conditions of aggregativity seem inapplicable to a structure maintained by supervenience relations. As noted above, the main claim of the relevant kind of emergence is that properties at different levels can have causal effects that 'cross-cut' the structure held together by the glue of supervenience. But this does not affect the very structure itself, which can be interpreted under an LCM framework.[12]

12 I am aware of arguments to the effect that examples from quantum mechanics both (a) support ontological emergence and (b) concurrently reject the LCM, or at least LCM-style structures (Humphreys, 2016; Primas, 1991; Silberstein & McGeever, 1999; Thalos, 2013). Two points on this. Firstly, of course those arguments are not conclusive and the debate remains an open one (Dickson, 1998; Healey, 1989). Secondly, regardless of the outcome of that debate it is clear that the proponents are working with a *dynamic* conception of emergence, rather than the synchronic conceptions being considered here (Humphreys, 2016; Silberstein, 2006). I agree that a dynamic conception of emergence is required to remove compatibility with the LCM, although I'll argue that it's not sufficient to do so alone, it needs to be adopted alongside a rejection of the distinction between reduction and emergence (see chapter 3).

When I introduced both the ontological interpretation of the LCM and ontological emergence, I simply stipulated that these labels included *both* entities and their properties (1.4.4). However, when discussing the exclusion problem Kim (1998, p. 77) is careful to distinguish between entities and their properties. He argues that some responses to the exclusion problem are based on a misreading of the argument. Specifically, that the argument does not apply to inter-level relations between individuals but only to 1st and 2nd order relations between properties instantiated by individuals.[13] This distinction makes clear that the exclusion argument concerns the causal status of properties belonging to individuals at the same level of organisation. For example, the reduction of intentional properties (mental properties) to neurological processes (physical properties) pertains to individuals at just *one* level of organisation – (L5) Multicellular living things – so the relationship between mental and physical properties is from a 2nd order property to a 1st order property but it is not an *inter-level* relation.

Regardless how Kim's distinction impacts the strength of the exclusion problem, it merely reinforces my point here. I'm interested how arguments surrounding emergence impact the plausibility of the LCM. We can understand Kim's response to be that the issue of the potential causal efficacy of properties that are related to other properties by supervenience is an entirely separate issue to the structural relationship between individuals that may instantiate those properties. So, for example, whilst it might be the case that we ought to assign causal roles to neurological properties rather than intentional properties, the multicellular living organism that instantiates *both* these properties is the causal actor *either way*.

13 This discussion comes up explicitly in relation to 'the expansion argument' (Baker, 2003; Burge, 2003). It's proponents point out that if the exclusion argument holds for the relationship between mental and physical properties, then it will apply to *all* properties other than those belonging to individuals at the most fundamental level. The argument will iteratively apply to any two properties related by supervenience until we try to apply it to properties that do not have a supervenience base (fundamental properties); only then can we stop causal powers of any given property being usurped by (excluded by) their supervenience bases. This would be a seriously worrying state of affairs, it is argued, because if only fundamental properties have causal efficacy, then strictly speaking there are no *true* causal explanations that do not cite such properties. Scientific explanation would be severely hamstrung as a result. Marras (2000) and Bontly (2002) provide responses to Kim to the effect that most sciences utilise functional properties in their explanations and these explanations would still fall victim to the exclusion problem even if we recognise the 'levels-orders' distinction.

So, the causal status of the properties instantiated by multicellular living things, for example, does not speak to the structural relationship between the multicellular living things and, say, the organs and tissues of which it is composed – i.e., it does not have an effect on the LCM.

So it seems that we can read the exclusion problem two ways, neither of which will speak to the issues of aggregate compositional relations between individuals arranged into distinct levels at a global scope. If we focus on the relationship between properties in the causal exclusion argument as *inter-level* it looks like the outcome of that discussion will not impact the structure of the system within which those properties interact. This is because the structural glue that holds the system together are non-causal metaphysical dependence relations, such as supervenience. Problems such as the causal exclusion problem may well speak to those kinds of metaphysical relations but not to the kinds of system features found in the LCM.

Alternatively, we could follow Kim and accept a clean distinction between orders of properties and level of individuals that instantiate those properties. But this just further underlines that the debate regarding the status of emergent properties – 1^{st} and 2^{nd} order properties related by supervenience – is simply disconnected to discussions concerning the structural relationship between individuals. In which case the fate of ontological emergence simply won't impact the plausibility of the LCM.

2.4.2 Epistemological Anti-Reduction

In his 1990 paper 'Why the anti-reductionist consensus won't survive' Ken Waters bucks a trend and argues that the arguments underpinning this consensus are unsustainable given historical and contemporary developments in science (specifically in molecular biology, in his specific case). He associates this consensus with canonical anti-reductionist arguments from Hull (1972), Hooker (1981) Kitcher (1984), Rosenberg (1985), Darden & Maull (1977) and others. Here I'll briefly lay out two types of these arguments. Once again, my purpose is not to weigh in the debate, but to examine the effects that the arguments would have on the resulting anti-reductionist position's compatibility with the LCM. They won't have much, as it turns out. This is a conclusion that Water's (2008) himself comes to in a later paper, going as far as label the position "layer-cake anti-reductionism". Waters thinks that most of the problems with the anti-reductionist consensus stem from a predominant *theory bias* in philosophy of science – i.e., an overemphasis on the role of theories in scientific inquiry by philosophers seeking to reconstruct and analyse those inquiries (see also Godfrey-Smith, 2008; Kaiser, 2012). I agree, but I'll point out a few considerations that go beyond Waters's argument that I think can be taken forward as lessons for a new account of inter-level relations.

The first set of anti-reductionist arguments fall under the label of 'unconnectability arguments'. They are arguments rooted in the multiple-realisability concerns pointed out in 2.3.1. As we saw there, multiple-realisability, as well as one-many, many-one, and many-many realisation relations, block the *identity* relations required in order to develop Nagelian bridge-laws. An unconnectability argument can be stated exactly as outlined there: Nagelian bridge-laws are not possible to construct because higher-level types are multiply-realisable. As Marie Kaiser points out, another objection of this kind made by philosophers of the life sciences is that the life sciences simply don't produce the kind of laws that would be required for a Nagelian reduction because "they typically have exceptions, are restricted in scope, and it can be argued that they are historically contingent" (Kaiser, 2012, p. 258). If the higher-level laws can't be produced in the first place, then Nagelian reduction is out, before we even get to bridge-laws.

The set second of arguments concern explanatory incompleteness. The idea here is that, not only is reduction impossible because of technicalities with Nagelian bridge-laws, but lower-level explanations are not always sufficient to explain their higher-level counterparts. So, even if you could reduce the higher-level *theory* to the lower-level one, you won't necessarily have *explained* the higher-level process. The most famous example of this kind of arguments, is Philip Kitcher's 'the gory details' argument; named after the following passage:

> "The distribution of genes to gametes is to be explained, *not by rehearsing the gory details* of the reshuffling of the molecules, but through the observation that chromosomes are aligned in pairs just prior to the meiotic division, and that one chromosome from each matched pair is transmitted to each gamete" (Kitcher, 1984, p. 370, my emphasis).

Kitcher, is essentially rejecting the claim that 'lower' is always better when it comes to explanations because whilst lower-levels will provide more specific details, higher-level explanations will provide more *generality*, which will sometimes be *required* to tackle the explanatory problem at hand. In this sense, higher-level explanations can sometimes be ‘autonomous’ from lower-level counterparts. Mark Bedau (2008, pp. 181–182) illustrates this with a nice everyday example, which I’ll summarise here. Strike action causes a traffic jam that makes everyone late for work. There are two levels at which we could explain this. We could put together all the causal histories of each individual car from the moment their drivers got in them that morning (the lower-level) or we could note that the strike action caused the traffic density above a *critical threshold,* causing a traffic jam (the higher-level). There’s an obvious sense in which the lower-level information is pointless detail here. But there’s a deeper point. The higher-level information will be applicable to more situations involving that road in the future. If I was investigating traffic flow in rush hour along that stretch of road, or was investigating punctuality of office workers who commute on that route, knowing the individual causal histories of the cars that day will not allow me to make the same sort of generalisations as knowing the critical density threshold and the factors that cause it to be exceeded.[14] This now resembles Kitcher’s point above. There are some explananda that *require* the higher-level information because the specificity involved in the lower-level information would mean it fails to apply to enough situations. This sort of argument continues to play an important role in debates about explanation and explanatory power (Jackson & Pettit, 1992; Sober, 1999; Strevens, 2009; cf. Potochnik, 2010).

With these two sets of arguments spelled out, we can see the minimal impact of anti-reductionism about epistemic items on the LCM. The sense in which they’re too weak is that they don’t tackle the structural features of the system in enough detail to make them explicitly incompatible with the LCM. We can see this more clearly as Kitcher’s argument proceeds. He suggests that a weak version of anti-reductionism consists in the claim that there are “autonomous levels of biological explanation. Anti-reductionism construes the current division of biology not simply as a temporary feature of our science stemming from our cognitive imperfections but as the reflection of levels of organization in nature” (Kitcher, 1984, p. 371). In this quote we get a feel for a few aspects of the LCM, particularly the *distinctness of levels*. In fact, that levels are ‘distinct’ is in some ways much more amenable to the anti-reductionist view than the reductionist one.

14 This point could be expressed using counterfactuals. The claim being that the higher-level counterfactual remains invariant (to a relatively higher degree) in cases when the lower-level counterfactual breaks down. Counterfactual dependency is deeply connected to explanation, at least for broadly ‘difference-making’ views of explanation.

Recall from the previous chapter that O&P had some delicate work to do in order to show that each level could be reduced to the next, but that each level in the model was still necessary because, “it would be utopian to suppose that one might reduce all the major theories or a whole branch concerned with one of our six levels to a theory concerned with a lower level” (Oppenheim & Putnam, 1958, p. 10). Anti-reductionism does not have to deal with this balancing act. Levels in the model are necessary because they *can’t* be reduced to one another. Similarly, claims about higher-level explanations being more general presuppose a fixed clustering of processes on these levels; otherwise it’s not clear how could you make the argument that higher-level explanations sometimes *require* the sort of generality that only obtains at the higher-level. This also speaks to a monism about levels, as does the appeal by Kitcher to ‘levels of organisation in nature’.

Of course, just using the word ‘levels’ doesn’t automatically mean that one has the LCM picture in mind. But the LCM view comes into focus more as Kitcher moves on to develop a stronger version of anti-reductionism:

> “Even if reductionists retreat to the modest claim that, while there are autonomous levels of explanation, descriptions of cells and their constituents are always explained in terms of descriptions about genes, descriptions of tissue geometry are always explained in terms of descriptions of cells, and so forth anti-reductionists can resist the picture of a unidirectional flow of explanation” (Kitcher, 1984, p. 371).

This sounds a more explicitly like a LCM structure: tissues, to cells, to genes ‘and so forth’. Kitcher is developing an epistemic analogue of the ontological emergence considered above. Rather than the unidirectional flow of causal relations, it’s the unidirectional flow of explanations that’s at issue here. A similarly analogous issue is that of aggregation and simple decomposition. The driving force behind the unconnectibility argument is still the multiple-realisability of higher-level properties. So, as we saw for ontological emergence, there’s nothing to stop the system being considered as simply decomposable; supervenience relations give structure to the properties and aggregativity gives structure to the objects that bear those properties.

This time, properties are embedded in theories and/or explanations rather than considered ontologically, but the result is the same.[15] At best, aggregativity is left up in the air; unaffected by rejections of reduction between epistemic items. Similarly, global scope is, at best, ambiguous here. The cases I've been working with pertain to the life sciences, so they tend to be focused on organisms, tissues, cells, and molecules. But there's nothing explicit in these arguments that rules out a global scope. It seems as though generality and specificity will be inversely proportionate all the way up and all the way down the levels structure; thus the explanatory incompleteness arguments should be applicable at any two levels in a global model. These observations about global scope and aggregativity are not decisive, but shortly I'll combine them with considerations drawn from ontological emergence, to pick out a broader lesson about these two aspects of the LCM.

2.4.3 Lessons Learnt

In the introduction to this chapter I stated that the main aim was to explore the complicated relationship between levels of organisation and accounts of reduction and emergence. I'll end the chapter by picking out the most important lessons that can be taken away from this discussion. I'll then take addressing these lessons as desiderata when developing a new framework for reduction and emergence that can complement a new conception of levels of organisation.

> Lesson (1): Rather than relying on an account of reduction *or* emergence as a strategy to avoid the LCM, we need a new account of reduction a*nd* emergence that is explicitly incompatible with the LCM.

15 Potochnik (2010) also questions the relationship between supervenience and epistemic levels but in a different way. She argues that in cases of 'complex realisation', all sorts of properties end up in the supervenience base for an instantiated property, ones that would sit on the next lowest level, perhaps lower levels than that, as well as same and even higher level properties, from the perspective of the LCM. A case of complex realisation would be camouflage in which the instantiation of camouflage (a property instantiated by an organism) requires properties of the organism, of the environment, and of predators (Potochnik & McGill, 2012, p. 128). I think that Potochnik's insight confirms the one I've been pursuing throughout this chapter so far: that metaphysical dependence relations between properties – such as supervenience – seem to come apart from the structural features of the individuals that instantiate them, including when those individuals are placed into a LCM-style levels structures, be that under an ontological or epistemic interpretation of the system itself.

This lesson is drawn from the main discussion that ran throughout the chapter; namely that both reductive and anti-reductive views, concerning both ontological items and epistemic items, remain compatible with the LCM. Mainstream debates in philosophy concerning reduction have focused on opposing Nagelian reduction. I argued that Nagelian reduction and the LCM became very closely intertwined via their association with the unity of science project. This probably explains the prevailing assumption that criticisms of reduction incorporate criticisms of the LCM. But we've seen that, not only is it incorrect to think that the LCM stands or falls with reductionism writ large, it's also mistaken to think that specific variants of the LCM stand and fall with specific kinds of reduction. The compatibility of the LCM with both reductive and anti-reductive views of explanation clearly demonstrates this. I suggest that the biggest lesson to take away from this result is that a strategy of rejecting the LCM by defending *either* reduction *or* emergence should be avoided. In its place, we need an account of inter-level relations – an account that encapsulates reduction *and* emergence – that demonstrates clear incompatibility with the LCM. Building an account of levels that is complemented by such a view of reduction and emergence should ensure that it avoids any problematic aspects of the LCM. The next three lessons concern how to avoid those aspects specifically.

> Lesson (2): The account of reduction/emergence must provide a clear understanding of the role that aggregation plays in inter-level relations and a clear sense of how and when aggregation fails.

A problem that surfaced in both ontological emergence and epistemological anti-reduction was the issue of aggregativity. For the metaphysical flavour of ontological emergence, *de novo* causal powers of instantiated properties didn't seem to force us to reject the claim that the bearers of those properties occupy fully aggregative levels. For epistemological anti-reduction, it wasn't even clear how to assess the claim that explanations that pertain to those properties are aggregative or not. Concerning anti-reductionism about explanations specifically, several philosophers have highlighted the focus on *theories* as a unit of analysis as a problematic assumption in the reduction/anti-reduction debate. Further still, Kaiser (2012, p. 259) highlights the combination of theory bias with the expectation that reduction should be a *deductive* relation amongst units of analysis, be them theories or explanations.

Seeing the ontological and epistemological issues side-by-side allows me to make an addition to Kaiser's point: the overlapping offender in both these cases seems to be *properties*, or more broadly a 'property-first approach'. This approach was manifested in both ontological emergence where properties stood as the potential causal relata and in epistemological anti-reduction where they were embedded within explanations or theories. When focused on properties in this way, the structure of the system is unaffected regardless of taking a reductionist or anti-reductionist view point.

I'm not suggesting that there's anything essential about 'properties' as a central unit of analysis that necessitates these problems. Rather, I'm proposing that the way in which the property-first approach has unfolded in philosophical discussions of emergence and anti-reduction has put issues of *system structure* in the background, in favour of metaphysical dependency, syntactic relations between statements that embed properties, and relationships between specificity and generality. The LCM positing a structure that is aggregative brings the issue of system structure back into the foreground and, I suggest, the lesson to be learnt here is that an account of inter-level relations needs to engage directly with issues of system structure as a primary issue. This can be put more concretely: an account of inter-level relations that can serve as a central part of a scientific plausible replacement for the LCM needs to give a clear understanding both of what aggregative systems are in relation to other kinds of systems, as well as how and why aggregative systems break down.

A consequence of pursuing this goal is that I will background issues of causal interaction between properties – be them in relation to supervenience bases or otherwise – in favour of focusing on a more metaphysically neutral sense of 'system dynamics' understood as interactions between components of systems and the way they affect the behaviour of the system as a whole and components at different levels in the system (see 3.3.2).

Lesson (3): Monism & Distinctness of levels are mutually reinforcing aspects of the LCM

As stated in the General Introduction, an assumption embedded within the dimension of evaluation governing 'scientific plausibility' is that an account of levels with any claim to scientific plausibility must be *pluralistic*. Clearly, then, I am not looking for ways to *reject* the monistic aspect of the LCM. Rather, its rejection is a starting assumption of my project.

Instead, what I'm looking for here is a way that the account can clearly *demonstrate* this pluralism and explain why the assumption of pluralism is indeed required for any claim to scientific plausibility. So, how could this be done? Presumably in several ways, but if we're focusing our attention on the four key aspects of the LCM, the most obvious candidate seems to be to deny the first part of *distinctness of levels:* deny that no entity 'properly' belongs to move than one level.

The salient question then becomes do the anti-reductive positions considered in this chapter help us to reject distinctness of levels? I don't think so. We saw that for non-reductive explanation, in order to address the question as to whether generality or specificity is always better the *structural* features of the system need to be fixed. As for ontological emergence, this fixity of the levels structure is similarly important. As we saw, even to question the potential 'upwards' and 'downwards' influence of causal effects, we need a fixed structure in place. We saw how that structure is provided by non-causal metaphysical dependence relations, meaning that neither the acceptance nor rejection of downward causation would affect this structure and neither would it affect the positioning of individuals in the system. Throughout the development of the discussion of emergence, we saw how considering the causal relationships between properties can come neatly apart from considering the structural relationship between the bearers of those properties – the individuals in the LCM structure. Figuring out which subsets of the properties of a multicellular living thing instantiates, for example, are actually causally efficacious and which other subsets are not, doesn't affect the compositional claims made by the LCM and therefore doesn't affect the fixed position of multicellular living things in the system: they remain level-bound.

To be clear, that we don't find a way to reject distinctness of levels from considering these anti-reductionist positions is hardly an indictment of them. Distinctness of levels is not even such a damning feature of the LCM until it's combined with *global scope*. We'll see how these anti-reductionist positions can't help us with issues of scope either shortly. But for now, the lesson I want to draw attention to from these considerations is that a monism about levels and a commitment to distinctness of levels are mutually reinforcing. Exhibiting the fluidity of the structural relationships between individuals across system conceptualisations (i.e., rejecting distinctness of levels) will provide the demonstration of pluralism required by scientific plausibility.

Lesson (4): The account must provide an explicit position on the scope of inter-level relations.

Global scope is a tricky aspect of the LCM. You get the feeling that many discussions of reduction and emergence are not primarily-oriented around issues of scope, not because they are *in favour* of a global perspective, rather because they think it's hardly worth mentioning.[16] In other words, issues of scope seem to be very much implicit in such discussions. However, failing to make a position on scope explicit leaves open a backdoor to global scope and in doing so leaves accounts of inter-level relations accidentally amenable to LCM-style thinking.

Given that my concern is explicitly with the LCM, I can't leave room for ambiguity about scope. Rather, the account of inter-level relations that I offer needs to provide an explicit position on issues of scope that avoids compatibility with the LCM. But what kind of explicit position does this need to be? Of course, it needs to reject the idea that 'levels of organisation' should have a global scope of application – like the LCM. However, there's also a danger of *overextending* this rejection. One of the fundamental issues with the LCM is that is makes structural claims about the world and scientific inquiry, *in principle*.

That is one of the reasons why it turns out to be so scientifically implausible; it simply doesn't match contemporary scientific *practice*. If the account of inter-level relations developed to reject the LCM also makes *in principle* claims, then it too will run the risk of quickly deteriorating in plausibility as scientific inquiry moves on. Put simply then, what we should want from an account of inter-level relations is the ability to explain why thinking about such relations in a global way is highly problematic but without recourse to 'in principle' arguments that put strong limits on how models might develop in the future.

16 I should note that within detailed discussions of supervenience a distinction is sometimes made between *global* and *local* supervenience, where the former applies to sets of properties in an entire (possible) world and the latter applies only between properties of an individual (McLaughlin, 1995). This is, then, a discussion of scope issues, but it's not one that affects the points I've been making here. Global supervenience quite obviously invites compatibility with the scope of the LCM. As for local superveniece, we saw in 2.4.1 how Kim harnesses a sense of local supervenience in his response to the exclusion argument by making the distinction between *levels* of individuals and *orders* of the properties they instantiate. But this only lead us further away from the idea that *de novo* causal powers of higher-level properties can impact the *structural* relationship between individuals in the system that instantiate those properties.

2.5 Conclusion

At the outset of this chapter I noted that the relationship between the LCM and reductionism was complex. I now hope to have unpacked some of that complexity and made the relationships between the LCM, reduction, and emergence sharper. The first step was to correlate kinds of reductionism with variants of the LCM via their overlapping units of analysis (Fig. 2). This allowed me to connect the epistemological LCM with Nagelian theory reduction and argue that the history of philosophy of science has driven via their mutual association with the unity of science project (2.3).

The next section (2.4) sought to interrogate the relationship between the LCM with views that were developed with the rejection of Nagelian reduction at their core: ontological emergence and epistemological anti-reductionism. I argued that both of these views remained compatible with variants of the LCM that focused on the same unit of analysis. I've attempted to stress in several places that nothing I've said here constitutes an argument for or against any of the reductionist or anti-reductionist positions considered in the chapter. Rather, the strategy was to show that, *if* we want a view about inter-level relations that is explicitly incompatible with the LCM, merely adopting an anti-reductionist stance will not suffice. Some philosophers have used these and similar insights to argue that the 'levels of organisation' concept has its limits both within science itself and as a conceptual tool in philosophical analysis (Love, 2011; Potochnik & McGill, 2012; Waters, 2008). I agree that the concept of levels of organisation has its limits in both these senses but I don't think that these limits can be identified on the basis of existing views of reduction and emergence because, as I've tried to show here, these views don't have much impact on the LCM. Instead, I'll offer a new framework for thinking about reduction and emergence.

To this end, the main outcome of this chapter was the development of four lessons, which I'll utilise as desiderata for a new account of reduction and emergence and in doing so provide an account of inter-level relations that is free from any features of the LCM. In turn, this framework will play a crucial role in the development of the account of levels I'll offer as a replacement to the LCM.

3

Beyond Reduction & Emergence: Constructing a Continuum of Strategies.

3.1 Introduction

The main aim of this chapter is to develop a new account of reduction and emergence that is incompatible with the LCM and can thus serve to provide a variety of inter-level relations that, in turn, will form a core component in my characterisation of levels of organisation. That account of reduction/emergence will consist in the development of a new framework: the continuum of research strategies. The continuum of research strategies represents a space within which different strategies can be compared according to how the system of inquiry is *conceptualised*. The main result of this will be that reduction and emergence are no longer definite categories, rather they become essentially *comparative* concepts.

I'll begin by exploring methodological approaches to reduction and emergence, understood through a 'practice-oriented approach'. This will lead to a replacement of 'disciplines' with 'research strategies' as a unit of analysis (3.2). I'll offer my own interpretation of research strategy consisting in four elements: a well-defined research question, a level of abstraction for carving a system boundary, methodological assumptions, and theoretical assumptions. (3.3) Theoretical assumptions can be used to understand the *conceptualisation of a target system* as more, or less, reductive relative to the application of *decomposition* heuristics and prevalence of key system *dynamics*. Different system conceptualisations can then be placed on a continuum of possible strategies and their more reductive or less reductive character exhibited in context.

In order to demonstrate how this framework can be utilised, I will consider a concrete case study: new developments in the field of cancer research (3.4). I'll plot putatively rival approaches on the continuum and show how they can be understood as offering boundary markers for the field as a whole: one marking the most reductive approach and the other marking the least reductive approach currently developed in mainstream inquiry.

To reinforce this point I'll show how recent approaches in the field can be understood as exploring the territory between these boundary markers, integrating important insights from both.

The main result of this work will be a new framework for understanding a plurality of inter-level relations; that is, a new way of thinking about reduction and emergence. I'll conclude by discussing how this new conception can accommodate the key lessons picked out in the previous chapter and in doing so make it apt for analysing inter-level relations within a scientifically plausible, pluralistic account of levels of organisation (3.5).

3.2 Methodological Levels & Reduction

Let's turn our attention to the neglected right-hand side of Fig 2. The overlapping unit of analysis there is *disciplines*. It's straightforward to understand what the LCM focused on disciplines looks like; it's probably the most familiar conception of the organisation of scientific inquiry. The LCM as a model for organising epistemic practices delineates scientific disciplines into distinct levels according to their *scope* of inquiry and their *objects* of study. Physics is at the bottom of the organisational system. It's scope of inquiry covers the smallest known objects like quarks and neutrinos, all the way up to the largest, solar systems, galaxies, and the universe writ large. It seeks the fundamental bits and pieces out of which everything else is composed and the laws that govern their interactions. If epistemological reductionism is correct, then *ipso facto* these are the laws that govern all interactions in the universe. As we move up a level the scope is narrowed. First to chemical reactions between elements & molecules; then to cells, multicellular organisms, populations thereof, and ecosystems. This hierarchical organisation of science yields the scientific disciplines with which we're so familiar – physics, chemistry, biology, psychology, sociology, ecology, and so on.

A move from ontological or epistemological reduction/emergence to a methodological form alone will not suffice to generate the incompatibility with the LCM that I'll need for developing a new analysis of levels. This is straightforward enough to see. We can envisage an account of methodological reduction according to which science is organised around distinct disciplines, whose overall goal is to develop their part of a Nagelian bridge-law for the purposes of reduction.

Each discipline churns out level-bound theories and explanations, which are made available to a meta-scientific reductive project like the unity of science project. The switch in focus to methodological accounts may not instantly afford an incompatibility with the LCM, but as I'll demonstrate throughout this chapter, it will provide us with a full array of tools to generate it. To get there we'll have to take two steps. First, to change the unit of analysis from *disciplines* to *strategies*; giving epistemic practices a specific interpretation through a *practice-oriented* approach to philosophy of science. This will provide accounts of reductive and non-reductive practices that avoid the 'property-first' approach that was symptomatic of the ontological and epistemological accounts considered in the last chapter. The second step is a slightly bigger one: to move from a distinction between reductive and non-reductive practices to a fluid *continuum* between more and less reductive practices. In doing so, reduction will be presented as an *essentially* relative concept. No practice is reductive or non-reductive *simpliciter*, it is only more, or less, reductive than another according to a frame of reference. Having a framework in which inter-level relations are essentially comparative notions, will go a long way towards facilitating the pluralistic conception of levels of organisation developed in the next chapter.

3.2.1 Epistemic Practices

As I noted in the last chapter, methodological reduction focuses on the epistemic practices of science. But what are epistemic practices supposed to be? In the philosophy of science, we can observe a trend that grows steadily from at least as far back from Kuhn's (1970) seminal *The Structure of Scientific Revolutions*. The trend in question, which I'll call a 'practice-oriented approach', seeks to leave behind the early to mid 20th century preoccupation with highly idealised and rationally reconstructed views of the aims, development, and outputs of scientific inquiry. The approach is an about turn away from seeking to explicate scientific outputs codified into theories and laws, bound within a logical structure of deduction. In short, a rejection of precisely the sort of analysis of science offered by the Nagelian framework I discussed at length in the previous chapter.

Hasok Chang's (2012, 2014) framework of 'systems of practice' provides a clear example of this approach. Chang offers a framework that co-ordinates what he calls 'epistemic activities' into a coherent 'system of practice'. An epistemic activity can be a simple as lighting a match, but it is defined essentially according to its *aims*. Such activities can have an 'inherent purpose' – the activity as an end, getting the match lit – and an 'external function'– the activity as a means to achieving a further goal, perhaps to light a Bunsen burner. Epistemic activities can be assessed individually in terms of their aims, the rules that govern achieving those aims, and the actions taken to do so. Eventually, these epistemic activities can be co-ordinated into more complex collections called 'systems of practice'. Systems of practice can be understood as aiming for coherence between the external function of each epistemic activity in itself.

Chang's framework highlights the typical modus operandi of practice-oriented approaches: understanding the aims of the agent, and their actions performed attemptting to realise those aims, is central to the analysis of science. Chang's notion of 'coherence' between activities is supposed to directly contrast with the method of analysing the *logical* consistency between epistemic outputs codified in terms of theories and laws. Coherence is a measure defined in terms of the aims of the scientist and evaluated in terms of how effective the constituent epistemic activities were in realising those goals.

Chang suggests that this framework can be put to work in variety of ways including descriptive historiography (evidenced by his own 2012); normative evaluation; and bridging gaps between scientific practice and other kinds of practice through the overlapping analysis of 'epistemic activities'. The most salient suggestion Chang makes for my purposes is that his approach can be used to replace abstract philosophical analyses of scientific processes and evaluation. Examples include defining concepts, the nature of explanation, and modelling practices. Whilst I won't be operating within Chang's specific framework, I'll be taking him up on the suggestion to apply the tools of a practice-oriented approach in this latter respect.[1]

1 Chang's analysis is probably at the less sociological end of what has been referred to as the 'practice-turn' in science studies more generally. By 'sociological' I refer to a more involved analysis of what 'practice' means in a scientific setting and the role of wider sociological structures as constitutive of scientific knowledge (e.g., Barnes, Bloor, & Henry, 1996; Latour & Woolgar, 1979). See Solar et al., (2014) for an overview of, and investigation into, the practice-turn in science studies.

In Fig. 2 I linked methodological reductionism to both ‘explanations’ and ‘strategies’. This short introduction to the practice-oriented approach helps me to explain why in more detail. Certain analyses that take explanations to be a primary unit of analysis are developed very much in-line with the practice-oriented approach precisely by placing the activities of scientists at the centre of the analysis. For example, when outlining the framework within which her account of reductive explanation will be developed, Kaiser notes that, “I start my metaphilosophical analysis by pointing out what it means to attempt to understand reduction in current biological research practice (Sect. 1). One of my main theses is that this aim commits you to focus on cases of reduction that actually occur in biological practice (reduction in practice).” (Kaiser, 2015, p. 8). Of course, Kaiser’s development of ‘practice’ and criteria for evaluation are different to Chang’s, but they share the practice-oriented ethos. When understood through the lens of a practice-oriented approach, there is only a fuzzy boundary between accounts of *explanatory* practices and *methodological* practices, with fruitful overlap between the two.

3.2.2 Step One: From Disciplines to Strategies

The first import from the practice-oriented approach to an account of reduction is to change the unit of analysis from disciplines to a more practice-based unit of analysis. I will choose to build on the concept of a research strategy for this purpose. Classical conceptions of research strategies include Kuhnian (1970) paradigms or disciplinary matrices; Lakatos’s (1970) research programmes; and Laudan’s (1977) research traditions. Since these classic accounts, views have been developed that are narrower in scope and aims. For example, they do not attempt to demarcate science from non or pseudo-science, nor are they designed to tackle questions of realism or incommensurability between different strategies throughout the history of science. However, even Kuhn’s account was not primarily designed to provide a global view of the organisation of scientific inquiry like the methodological LCM. A disciplinary matrix or ‘paradigm’ applies to a small community of researchers of, “perhaps one hundred members, occasionally significantly fewer” (Kuhn’ 1970, p. 178) defined in reference to their shared goals, linguistic & conceptual tools, technical literature & education, and formal and informal communication networks (including citation circles).[2]

Let's take a quick look at one of these accounts which is explicitly narrow in scope in a way that is less open to interpretation than Kuhn's; Darden & Maull's (1977) account of 'inter-field theories'. Darden & Maull characterised a 'field' of science as containing the following elements:

a) A central problem
b) A domain consisting of items taken to be facts related to that problem
c) General explanatory factors and goals providing expectations as to how the problem is to be solved
d) Techniques, methods, and, sometimes but not always, concepts laws, and theories which are related to the problem and which attempt to realise the explanatory goals
e) A special vocabulary associated with characterising elements of the field. (Darden, 2006, p. 128)

Darden and Maull utilised the organisation of inquiry into fields in order to show that progress in science often takes place through the construction of *inter-field* theories. The concept of an inter-field theory is functionally analogous to Nagelian bridge-laws as they both serve to 'bridge' gaps between areas of inquiry. Inter-field theories, however, highlight the differences between methodological and epistemological approaches to reduction and emergence. Firstly, fields are first and foremost *problem-oriented*, rather than scope or entity-based like traditional disciplines. Secondly, there is a role within the analysis of fields for epistemic outputs, like theories and laws, but these epistemic items are not *central* to the characterisation of a field. Thirdly, the identification and evaluation of fields is relative to the explanatory aims of the field; which I take to mean the explanatory aims of the scientists involved in the research.

2 Kuhn admits that the label of a community can apply at several levels including that of all scientists, and to groups that roughly align with traditional disciplines. 'Paradigms' can be shared by all manner of groups but I take it that Kuhn's focus in primarily on this narrower sort of community seeing as they are, "the producers and validators of scientific knowledge" (Kuhn, 1970, p. 178). See Hacking (1983, p. 10*ff)* for further discussion.

Darden & Maull's account is a large step away from theory-oriented analyses of scientific organisation. Although theories are not central to characterising fields, their account still relies on theories as a core component in the progress of scientific inquiry (i.e., the construction of inter-field theories).[3] One further step away from theory-orientation is the accounts that place *heuristics* at the core of characterising a research strategy. Such accounts draw large inspiration from Herbert Simon's (1996) work and, in the context of research strategies, are probably developed in the most detail by Wimsatt (2006), as well as Bechtel & Richardson (2010). Wimsatt (2006, pp. 464–465) provides an overview of the characteristics of heuristics, the following is a truncated version of that overview:

Heuristics:

I. Are not *truth-preserving* algorithms; they make no guarantees that they will produce a correct solution to the problem.
II. They are more cost-effective (efficient) than a truth-preserving algorithm
III. Produce *systematic* errors; thus are effective diagnostic tools
IV. Transform the problem into a non-equivalent but related problem
V. Are always purpose-relative
VI. Are commonly descended from other heuristics.

For Wimsatt, then, the development, application, and modification of heuristics to a problem (and of a problem) constitutes the core of a research strategy. Wimsatt's heuristic-based strategies are about the furthest away we can get from Nagelian theory reduction as the methodology of scientific practice and progress. Like Darden & Maull, Wimsatt sets up his approach in direct contrast to the Nagelian-style formal systems of deductive logic that governed theory reduction by explicating heuristics in contrast to 'truth-preserving algorithms'.

3 This is a potential criticism of Darden & Maull's view, although they are very explicit that by 'theory' they do not mean a deductive system, rather they see inter-field theories as solutions to theoretical problems (Darden, 2006, p. 128). Either way Darden (2006) has since embedded her account within a larger 'mechanistic' framework, which affords a full exorcism of Nagelain-style 'theories' as a primary unit of analysis.

Simon, Wimsatt, and Bechtel & Richardson are all interested in the idea that humans *reason* using heuristics and thus scientific activity (done, of course, by humans) should be characterised according to the way *we* problem-solve, store, and transform information.[4] Thus, these are analyses of scientific practice *par excellence:* "We operate on assumptions that psychological constraints are important in understanding scientific change, confirmation, discovery, and that the sorts of explanatory strategies employed in scientific problem-solving are analogous to the strategies employed elsewhere in human problem-solving (Bechtel & Richardson, 2010, p. 7).

Wimsatt, as well as Bechtel & Richardson, developed a *reductionist* account of heuristics-based research strategies by focusing on the heuristics of *decomposition & localisation* (see 3.3.1). Similarly, the extremely prominent 'mechanistic' approach to philosophy of science, for which Wimsatt, Darden, and Bechtel & Richardson laid the foundations, can be characterised as type of reductive research strategy.

In addition to these reductionist views of scientific practice, an increasing amount of attention has been turned to 'non-reductive' or 'system-level' research strategies. These accounts attempt to grapple with strategies that develop from failures in reductive heuristics such as decomposition and localisation, as well as strategies whose aims and objects of inquiry seem less amenable to reductive approaches (Brigandt, 2013; Green, 2015; MacLeod & Nersessian, 2015). Additionally, practice-oriented analyses have yield new insights on previously marginalised forms of explanatory practices such as topological explanations (Kostić, 2016) and constitutive explanations (Kuorikoski & Ylikoski, 2013).

Broadly, the practice-oriented approach has opened up new ways of thinking about 'reduction' and 'anit-reduction', by associating those terms with *actions* rather than *logical relations* between propositions. This has afforded the reassessment of cases classically considered to be exemplars of reduction in the history of science, as well as highlighting the importance of different kinds of reductive and non-reductive practices in order to solve problems across (and within) scientific fields (Andersen, 2017; Chang, 2015; Kaiser, 2011; O'Malley et al., 2014).

4 This idea explored in depth in Kahneman & Tversky's work (1972; 1982) who purported to demonstrate that humans do not always reason according rules of propositional logic (i.e., 'rationally'), and instead the errors in their problem-solving abilities can be traced to systematic heuristics and biases. Their 'Prospect Theory' framework (1979) went on to be a cornerstone of behavioural economics.

In what follows, I want push this reassessment a step further; beyond the distinction between reductive and non-reductive practices altogether. I'll argue that no practice is reductive or non-reductive *simpliciter*. Rather, such practices are essentially comparative – a practice is only more reductive or less reductive than another. Research strategies can be placed on a continuum for the purposes of comparison and analysis.

3.3 Step Two: from a Distinction to a Continuum

To take this second step I'll start by providing an account of research strategies as a unit of analysis for understanding reductive and non-reductive epistemic practices. It is not intended as a rival to other accounts of research strategies considered above and there will inevitably be some overlap. However, it is designed to draw attention to the key elements of a strategy that shape its character such that it can be considered more, or less, reductive. There are four such key elements that comprise my view of a research strategy. The first is a well-defined research question – a clearly spelled out question or target of inquiry that serves to motivate and focus epistemic inquiry. When attempting to compare and contrast research strategies, the well-defined research question gives an indication that the respective strategies in question are dealing with, at least approximately, the same target phenomenon. This question will often be formed in terms of a particular behaviour of system. 'Behaviour' could – but need not – be understood as static features of a system; not wildly dissimilar to a 'property', in the sense seen in the previous chapter (mental 'properties' for example). However, particularly as the system itself is conceptualised as more integrated (3.3.1), system 'behaviour' will be understood as certain patterns of interaction that map inputs and outputs of the system, perhaps that result in stable and predictable oscillations between different states of the system (3.3.2). I think a comfortable way to think about behaviour is simply that a 'behaviour' is something a system *does*, rather than something a system *has*.

Related to the research question is the second element: a level of abstraction at which the inquiry is to begin. Fixing the level of abstraction at which the system-level behaviour or output emerges (the boundary of the target system of inquiry) is crucial for the purposes of comparing approaches on the continuum of research strategies I develop below. Specifically, it is important in order to assess the respective strategies in terms of component or interaction dominance (see 3.3.2).

In contrast, it is the *differences* in the proposed level of abstraction at which the inquiry begins (where the system boundary is drawn) and how the inquiry *proceeds* (where the functional parts responsible for the system's behaviour are posited) that are reflected in the relative positioning of research strategies on the continuum. This element features *levels.* As noted in the previous chapter the concepts of 'levels of organisation' and the inter-level relations that constitute accounts of reduction and emergence are intimately interwoven, so it shouldn't be surprising to find 'levels' being referenced in my account of research strategies. However, at this stage I mean 'level' in a very minimal sense. This second element of fixing a 'level' for the inquiry to begin only means *drawing a system boundary.* That is, carving out the system of inquiry from its environment. The proposed level of abstraction for the inquiry to develop refers only to whether (and to what extent) certain heuristics will be applied to the system once the boundary has been drawn, for example decomposition into discrete modules.[5]

The third element is methodological assumptions. These may include equipment used and certain experimental techniques, as well as the setting and mode of inquiry for the experiment(s). For example, will the experiment take place in the field or in the laboratory? Are the experiments to be conducted *in vitro, in vivo,* or *in silico*? If computer simulations are required (*in silico*) how will these be set-up and what procedures will be adopted?

The final element is theoretical assumptions. These assumptions form the basis of what point on the continuum a research strategy is positioned relative to another by making assumptions about the structure and features of the target phenomena. These assumptions may comprise a wealth of different commitments about the causal features of a system, the organisational and constitutive features of system, and the relation of the target phenomena to other systems. Overall the theoretical assumptions form a core part of the *conceptualisation* of the target system.

5 It is difficult to articulate the precise relationship between research strategies and levels of organisation (under my framework) at this juncture in the process of introducing these concepts. However, as we'll see in the next chapter, aspects of what I will argue constitute a levels of organisation concept are distributed amongst the components of a research strategy. From that perspective it would be no surprise that 'levels' shows up in the process of boundary carving, I'll discuss this issue in detail in 4.2.3.

Each of these elements is drawn from the practice-oriented approach. The aims of the inquiry are central to the analysis of the strategy (first element). The identification of a system boundary (level of abstraction at which the inquiry is to proceed) as defining of the strategy places an emphasis on modelling practices over the discussion of free-floating ‘properties’ (second element). Methodological assumptions are important components in the strategy (third element), that, under further investigation could situate a strategy within its broader sociological and technological context. Finally, the fourth element seeks to explore the problem-solving strategies that are involving with building, refining, and comparing models of target systems of inquiry.

All four of these elements constrain each other and are influenced by one another. Rather than a hierarchical structure, they are mutually dependent and dynamically responsive to changes.[6] Nevertheless, my contention is that the conceptualisation of the target system shapes the overall character of the research strategy. I will focus on two broad aspects of system conceptualisation and show how they inform a strategy of research: (1) the system’s susceptibility to *decomposition* and (2) the key *dynamics* of the system. These aspects of system conceptualisation will serve as the landscape of the continuum.

3.3.1 Decomposition

Decomposition is the process of breaking down a system into its component parts. The process includes carving out the boundaries of the system from the environment (Wimsatt, 2006, 2007), as well as the internal break down of the system’s components structurally, functionally, or both (Bechtel & Richardson, 2010; Kaiser, 2015, p. 75). Bechtel and Richardson’s (2010, p. 26) analysis of decomposition begins with the division of systems into three types: aggregative, component, and integrated.

6 For example, often practical decisions or practical obstacles will strongly affect system conceptualisation (theoretical assumptions). Simplifying assumptions are often required given limitations of experimental tools available or a lack of knowledge about the system *at a given stage of inquiry* e.g., in the search for an appropriate unit of intervention for the system. Assumptions such as studying parts in isolation from their environment context (Kaiser, 2015, pp. 225–229), or homogenising or fixing environmental factors in space and time (Wimsatt, 2007, pp. 347–352) sometimes (but not always) follow from these sorts of practical considerations. Here, my analysis is focused on the theoretical assumptions that result in system conceptualisation as a starting point for a practice-oriented account of reduction. Should this analysis be convincing, then important specificity will surely be added with a thorough treatment of how and why those given theoretical assumptions are adopted.

These systems are primarily divided by their susceptibility to decomposition: aggregative systems are simply decomposable; component systems are nearly decomposable; and integrated systems are minimally decomposable. I introduced the idea of aggregativity in Chapter One, in a fairly simply way so let's add some detail here. Aggregative systems are simply decomposable in the sense that they meet several conditions. The system continues to function and its behaviour remains invariant: first, regardless of the rearrangement or interchanging of parts. Second, when the number parts are added or taken away (that is, the quantity of parts does not affect the system, the qualitative features of those parts may do so). Third, regardless of any particular operation of decomposition and re-aggregation of parts (you can take the system apart and put it back together in any number of ways without affecting the system-level properties). Finally, there are no co-operative or inhibitory interactions among parts. These conditions capture the strength of dependence of the system on its parts and, due to their severity, it is argued that all of these conditions are hardly ever concurrently met, particularly in natural systems (Wimsatt, 2007, pp. 279–281).

Component systems are nearly decomposable in Herbert Simon's (1996, pp. 197–198) sense of the concept when interactions within a component (or part) are stronger than interactions between components (or parts). Near decomposition of systems into sub-components takes place by mapping the strength of interactions within the system in order to calibrate relative frequencies of interaction. Where clumps of interactions can be detected, decomposition into subcomponents can occur. This will give a more fluid decomposition than simple decomposition, the boundaries of each component will not be so discrete and some organisational constraints on the system may arise. Nevertheless, each component in the system will "operate primarily according to its own intrinsically determined principles" (Bechtel & Richardson, 2010, p. 25). That is, for both simply and nearly decomposable systems the components are relatively independent of one another – parts will function relatively independently of the whole system within which they are embedded.

For integrated systems, susceptible only to minimal decomposability, the situation is different in that "systemic organization is significantly involved in determining constituent functions" (Bechtel & Richardson, 2010, p. 26). In other words, the component parts fail to be functionally discrete from each other and system-wide interactions will affect the functioning of individual components.

In such systems it is neither possible to functionally decompose the system into sub-components nor functionally localise the system-level properties to physical sub-components. What this means in practice is that the research strategy will be forced to abstract away from internal structure and construct models with the aim of capturing the operations required to produce and maintain the system-level behaviour; a strategy that Bechtel and Richardson label as 'synthetic' in contrast with the 'analytic' strategies of decomposition. For integrated systems, then, rather than a focus on decomposition, the strategy will attempt to identify organisational constraints, sometimes called organisational principles (Wolkenhauer & Green, 2013), or design principles (Green, 2015) that arise from interactions between the components of the system.

The conceptualisation of the target system as aggregate, component, or integrated in virtue of its susceptibility to decomposition gives the first measure for a continuum of research strategies seen in Fig 4. Decomposition thus affords the first general feature available for the comparison of research strategies: a research strategy will be more reductive the more it conceptualises its target system as susceptible to decomposition (the further to the left on the continuum it positions its target system). A mechanistic system, for example, should be placed roughly around the centre of the map – it is near decomposable into relatively discretely functioning components. Thus, a research strategy embedded in the mechanistic approach is reductive relatively speaking i.e., it is more reductive than a strategy that conceptualises its target system as comprised of highly dependent components that fail to be functionally discrete (towards the right-hand side of the map) but it is not as reductive as strategies that conceptualise their systems further to the left of the map – fully decomposable, aggregative systems (I'll return in significant detail to mechanistic systems in Chapter Five).

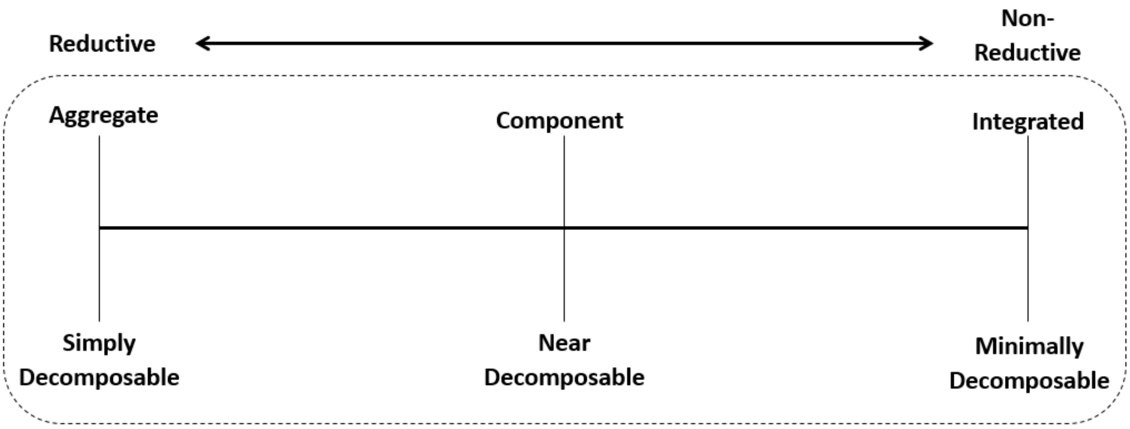

Fig. 5 A partial continuum of research strategies. The double-arrowed line outside the box represents the scale of how reductive a strategy is; more towards the left, and less towards the right. The labels at which the arrows point represent a *never reached value*—i.e., a strategy is, say, *more* reductive but there is no, final, ultimately reductive strategy. The dashed box represents the deconstruction of that scale. At this stage there is one dimension inside the box; decomposability, again, more to the left and less to the right. The straight vertical lines on the dimension of decomposability represent three idealised *paradigm points*—i.e., they do not (necessarily) represent real system conceptualisations, rather they provide fixed points around which an analyses of a research strategies can be oriented. Figure from Baxendale (2018, p. 6).

3.3.2 Dynamics

Another measure of the conceptualisation of a target system is the prevalence of certain key system dynamics; dynamics that are present only across the system as a whole, rather than any of the individual components. Many sorts of interaction fall under the label of ‘system dynamics’ but here I want to focus on a few key dynamics that seem to be most indicative of the sort of strategy adopted by researchers who focus on them. Furthermore, it makes sense to cluster these concepts together because (a) they often arise together with the presence of one strongly suggesting, as well as explaining, the presence of others and (b) because whilst none of them alone tend to be considered as ‘sufficient’ for a system to be integrated, the recognition of several of them in one system usually results in the conceptualisation of the system as such (Ladyman, Lambert, & Wiesner, 2013, pp. 36–40; Mitchell, 2009, pp. 35–

44).[7] The system dynamics I refer to are non-linearity, feedback loops, and self-organisation (understood as a result of distributed control). Furthermore, I will group the presence or absence of these key dynamics in the conceptualisation of a target system under two broad categories that I borrow from dynamic systems theory (1) component dominance and (2) interaction dominance (Holden, Van Orden, & Turvey, 2009; Favela, 2015; Favela & Martin, 2017) .

These features can be outlined and explained through a tried and tested example – the flight of a flock of birds. Starting with feedback, for the flock of birds this can be understood as the movement of an individual bird being affected (or constrained) by the movement of the flock as a whole. More abstractly, feedback results in dependence loops, such that the function of a component can be affected by the functioning of another component that it sent signals to at a previous time-step. A related feature of the flock is that the interactions between the individual birds – the components of the system – are non-linear. Non-linearity refers to the idea that the output effects of a system are not straightforwardly proportionate to their inputs; non-linear interactions are not additive. Additive interactions mean that the output of the interactions will be nothing more than the summation of the interactions of the components. In contrast, the output of non-linear interactions cannot be ascertained by summing (adding up) the values of the components. In the case of the bird flock, this means that understanding how it moves will not be ascertained by assessing the function of each individual bird and adding them together as linear sequence of events.[8]

Within the flock of birds there is no one bird (nor a specific sub-collection of birds) in the flock that is responsible for measuring, controlling, or implementing changes to the synchronised movement of the flock. Thus, the flock of birds can be understood as a system with distributed control. Distributed control means that relatively stable system-level behaviour is generated solely by the system-wide interactions of its components, rather than as the result of any sort of blueprint or directed series of events; resulting in the appearance of what is often referred to as 'self-organisation' within a system.

7 Sometimes these analyses are cast in terms of 'complex' systems rather than integrated systems. I purposefully opt for 'integrated' in order to remain neutral on what a 'complex' system *is*. There's an active debate ongoing concerning whether the concept of a 'complex' system denotes any one type of system in particular. Sometimes the sort of dynamics I am discussing here feature centrally in those analyses (e.g., Mitchell, 2009) but sometimes they do not (e.g., Zuchowski, 2017).

8 For a detailed but clear and accessible explication of non-linearity see Favela (2015, pp. 45–47).

The structure of the relationship between these kinds of dynamic interaction is often precisely the question at issue when studying systems that are conceptualised as integrated. Presumably there will be many ways in which such dynamical patterns of interaction jointly result in system-level behaviour in such systems. Furthermore, my exposition here is in intended as a typical, but not exhaustive, list of the sorts of dynamics that may be of interest to researchers. Accordingly, this brief explication of the dynamical interactions is merely supposed to highlight some key patterns of interaction that are often invoked together when exploring certain kinds of system; resulting in certain sorts of research strategies.

Finally, I borrow an addition from dynamic systems theory – namely a distinction between component dominance and interaction dominance (Holden et al., 2009; Favela, 2015). For my purposes, if a system is conceptualised as having these key system dynamics then the system will qualify as interaction dominant; if not then it will be component dominant. These categories also sit on a continuum and are thus not supposed to be absolute, i.e., systems will be relatively more or less component dominant.[9] The presence and prevalence of those dynamics will be gradual and relative. I use these terms because their characterisation in the literature nicely corresponds with the processes of decomposition – linking the two measures together on the continuum of research strategies. Recall that under Herbert Simon's definition near decomposition occurs when the strength or frequency of interaction within parts in stronger than between parts. Consider that component dominance is characterised as when, "dynamics *within components* dominate interaction among components" (Holden et al., 2009, p. 319, emphasis added), with interaction dominance being when, "dynamics or properties of the interactions among parts supersede those that the parts would have had separately" (Favela, 2015, p. 44). Thus, when considered in these terms, the conceptualisation of the dynamics of a system constrain the appropriate degree of near-decomposability for a system and vice-versa.

9 I should be clear that these usages of interaction dominance and component dominance are inspired by, but not necessarily the same as, their usage in dynamic systems theory. However, I think my use of them as container terms for the dynamics of a system is certainly still faithful to what they represent in dynamic systems theory.

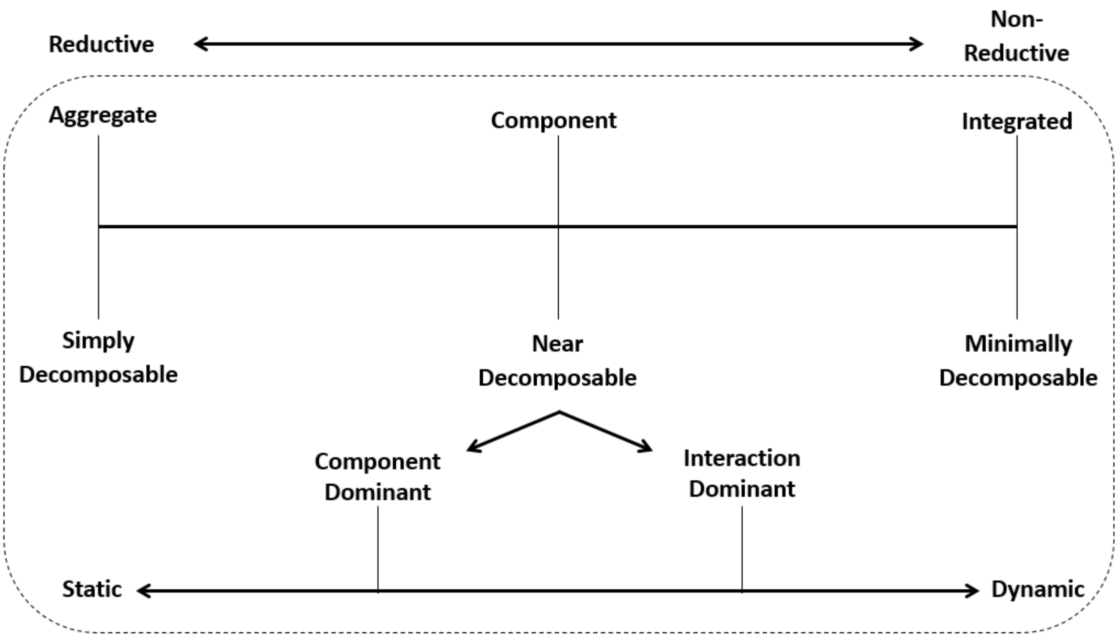

Fig. 6. The full continuum of research strategies. Inside the dashed box we now have a second dimension representing a focus on system-wide dynamic interactions; less towards the left and more towards the right. As in Fig. 4, the labels at which the arrows point represent *never reached values* – no system is thought to be *the* ultimately dynamical system. Like the dimension of decomposability, this second dimension also has vertical lines representing idealised paradigm points – component and interaction dominance – round which to orient analyses of research strategies. Further, these points serve to link the dimensions of decomposability and dynamics through the concept of a near decomposition. Figure from Baxendale (2018, p. 9).

Putting this all together, we get the full continuum of research strategies shown in Fig. 6. The spike of ‘interaction dominance’ denotes the point at which the key system dynamics of feedback, non-linearity, and distributed control become of central importance to the conceptualisation of the target system. As the system moves further to the left it begins to background these dynamics and move towards component dominance. The two spikes also link the dynamics of the system directly to its susceptibility to decomposition. On the map, both forms of dominance (both collections of system dynamics, or lack thereof) are linked to types of near-decomposition. As system conceptualisations move further to the right, the decomposability becomes ever more minimal. Moves to the left will mean increasingly straightforward decomposition strategies. Hopefully, the notions of interaction and component dominance demonstrate the strong link between the measures of decomposability and dynamics in system conceptualisation.

It doesn't seem plausible to claim that one of these measures is a result of the other, rather they are intimately linked. A system will be conceptualised as susceptible to less decomposition because it displays interaction dominant dynamics and concurrently the prevalence or importance of interaction dominant dynamics will strongly suggest that strategies of decomposition and localisation may fail to capture the target phenomenon.

One way to understand why such a strategy may fail is to consider what the appropriate manipulation of a target system might be. If a system is mechanistic, in the sense of near decomposable and displaying minimal key dynamics, one could manipulate any one of the system components to observe a direct effect on the functioning of the whole systems and thus, presumably, ascribe the system-failure to that one (now) malfunctioning component. In contrast, a minimally decomposable system displaying the dynamics of feedback loops, distributed control, and non-linearity, manipulations on a single component – say a single bird – may have no effect at all on the functioning of the system. Alternatively, it may have completely unexpected downstream effects due to its involvement in a feedback loop and its participation in the distributed control of the whole system.

Similarly, often mathematical modelling and *in silico* experimentation become more prevalent when systems are conceptualised as exhibiting system-wide dynamic properties. This is because (1) the modelling of such dynamics involves analysis of computationally expensive high-throughput data,[10] and (2) capturing dynamic interactions that occur only across the system as whole, such as non-linear interactions that exhibit disproportionate relationships between collective variables and order parameters. (Brigandt, 2013; Issad & Malaterre, 2015; MacLeod & Nersessian, 2015).

10 The phrase 'high-throughput' is widespread in contemporary life sciences research. At root it simply refers to the ability to process large amounts of data at once. In the life sciences, at least, it is perhaps most commonly associated with the rise of 'omics' research such as genomics; probably boosted by its role in the completion of the Human Genome Project. In that context, 'high-throughput' refers to the ability to sequence massive quantities of DNA samples at once. I'll give an example of this – cancer genomics – in 3.4.1 below.

3.4 Testing the Analysis: The Somatic Mutation & Tissue Organisational Field Theories of Carcinogenesis

The foregoing introduction to research strategies allows me to draw a nice landscape of possible system conceptualisations (Fig. 6), but we'll need a concrete case study to see how this continuum can be used to identify *comparatively* reductive strategies. Further, a case study will allow me to bring out the full utility of the continuum by showing how strategies that do not neatly fall into a 'reductive' or 'non-reductive' categories, traditionally understood, can be accommodated within this framework. The case study I'll work with is recent developments in the field of carcinogenesis; the study of the origins and development of cancer. A common distinction is made between hereditary & sporadic cancers with the latter being loosely defined simply as *not* the result of inherited genetic mutations and makes up between 95% and 98% of known cancers (Sonnenschein & Soto, 2008, 2011)/ Technically, then, I will be focusing on research into the origins, causes, and development of sporadic cancers. That being said, in either case, the primary goal of cancer research is to understand the complex interplay between a variety of factors – including environmental, developmental, and genetic – that identify the origins, and result in the development, of cancer.

To date, the two most influential – or at the least most discussed and analysed in philosophy of science – theories that focus on sporadic cancer development are the Somatic Mutation Theory (SMT) and the Tissue Organisational Field Theory (TOFT). Saliently, these theories have been pitched as taking reductive and non-reductive approaches to cancer respectively, not only by philosophers discussing them, but also by scientists. For example, the originators of the TOFT, Ana Soto and Carlos Sonnenschein have published several papers in which they explicitly argue that their approach is 'organicist', 'integrative', and generally 'non-reductive', and stand their strategy of research in explicit contrast to what they perceive as the 'reductionist' approach embodied in the SMT (Soto & Sonnenschein, 2005a, 2005b, 2011). The relationship between these two theories of carcinogenesis has become of increasing interest to scientists (Baker, 2015) and philosophers alike (Bedessem & Ruphy, 2015; Bertolaso, 2009, 2011). Here I will briefly lay out their key points of difference.

I need to strongly stress that my intention is neither to weigh in on the relative merits of either approach nor to claim that these quick characterisations fully capture their nuances. My focus is on the research strategies pursued by contemporary researchers in carcinogenesis and demonstrating that the framework I offer can afford an analysis of those strategies that fully reflects the complex, multifaceted, and dynamic processes in scientific practice.

The classic SMT focuses on DNA mutation as the primary cause or origin of cancer. I am labelling what follows as the 'classic' SMT to differentiate it from the most recent incarnations of the SMT (more on this below). The classic SMT can be elucidated by articulating its three main principles:

(1) Cancer is derived from a single somatic cell that has successively accumulated multiple DNA mutations (monoclonality).
(2) Those mutations occur on genes that control cell proliferation and the cell cycle.
(3) Implicitly the default state of cell proliferation in metazoa is quiescence (quiet, still, or inactive) (Sonnenschein & Soto, 2008, p. 3).

The core of the SMT theory states that cancer is to be explained by articulating the mechanism that starts with multiple DNA mutations on certain genes (oncogenes). These genes affect the cell cycle and cell proliferation, and results in uncontrolled cell proliferation and tumour growth. A specific example of a form of SMT is presented by Hanahan & Weinberg (2000) and Weinberg (1998). Hanahan & Weinberg (2000, p. 57) argue that, "cancer cells have defects in regulatory circuits that govern normal cell proliferation and homeostasis." They identify six 'essential' alternations in cell physiology that dictate malignant growth. These are: self-sufficiency growth signals; insensitivity to growth inhibitory signals; evasion of programmed cell death (apoptosis); limitless reproductive potential; sustained angiogenesis (the development of new blood vessels; and tissue evasion. These six alterations fill out (2); they spell out exactly what the genetic mutations affect in the cell proliferation mechanism and cell cycle. They also implicitly adopt (3) by indicating that 'normal' cells (cells without these altered mechanisms) are subject to intra-cellular constraints on growth.

In other words, it is the processes within the cell itself for both normal and abnormal cells that dictate the default state of the cell; that state being quiescence. Finally, by placing the focus of research firmly within the cell they also adopt (1).

The TOFT is an alternative to the SMT, chiefly promoted by Soto and Sonnenschein (2005a; 2008; 2011). Soto and Sonnenschein characterise their theory in direct contrast to the SMT. Firstly, they deny that quiescence is the default state of cells, rather they argue that proliferation is the default state of cells. They note that, “among microbiologists, it is axiomatic to accept that proliferation is the default state of prokaryotes and unicellular eukaryotes” (Soto and Sonnenschein, 2011, p. 2). They argue that researchers seemed to have ignored this issue in cancer research and for multi-cellular eukaryotes more generally rather than providing any argument or conclusive experimental data against it. They have demonstrated their hypothesis in experiments on estrogen target cells (Soto & Sonnenschein, 1987) and others have attempted to corroborate this hypothesis with research on the active maintenance of quiescence in lymphocytes (Yusuf & Fruman, 2003).

Secondly, they argue that carcinogenesis originates in a failure of tissue organisation rather than as the result of a single somatic cell with mutated DNA. To be more specific, carcinogens disrupt the organisation of the tissue including the morphogenetic fields and various components in the extra-cellular matrix. The concept of a ‘morphogenetic field’ was developed by the Organicist Joseph Needham (1931) and Soto & Sonnenschein, at least, explicitly align themselves with an ‘organicist approach’. The interference of carcinogens to the organisation of the tissue causes a break-down in negative controls on epithelial cells, which in turn enable them to express their constituent property of proliferation, a process known as hyperplasia. Thus, for the TOFT, the proliferation of cells responsible for tumour growth is to be explained by the break-down in tissue organisation rather than the uncontrolled proliferation of a single malignant cell with mutated DNA.

3.4.1 Plotting the Approaches

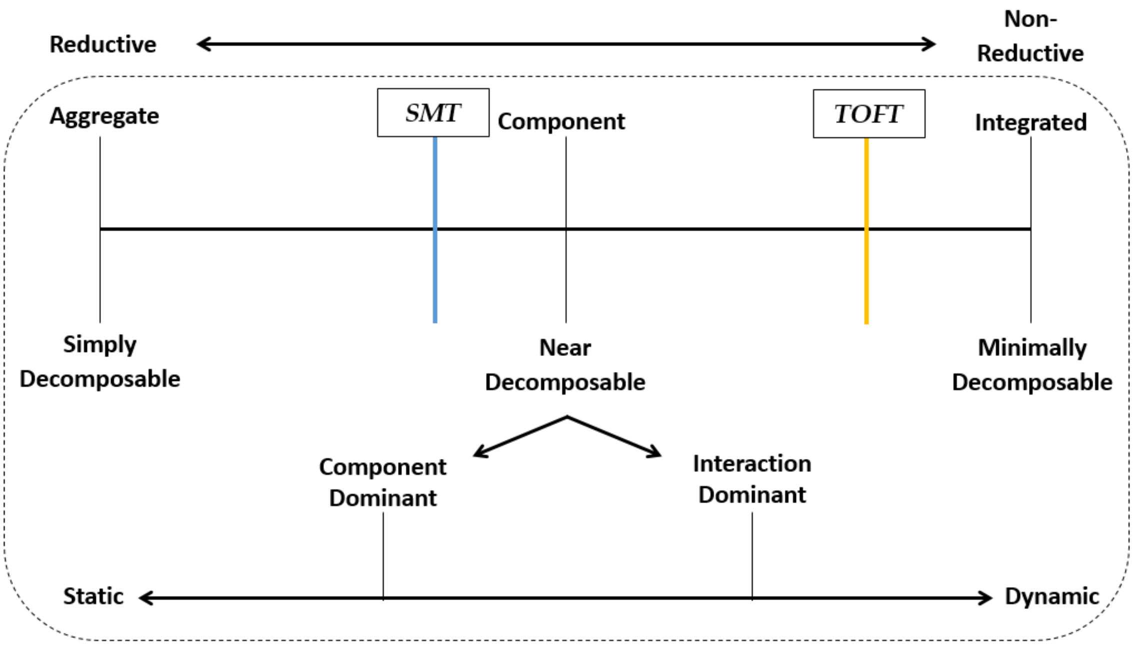

Fig. 7 The SMT and the TOFT plotted on the Continuum. Figure from Baxendale (2018, p. 12).

The classic SMT and the TOFT can be plotted on the continuum as shown in Fig. 7. For both approaches the well-defined research question is the genesis and development of sporadic cancers. The level of abstraction at which the system level behaviour emerges is the tissue, or organ, level. Cancer is usually taxonomically classified by reference to the organ or tissue mass it is found in; pancreatic cancer, bowl cancer, breast cancer, and so on. For both approaches some decomposition is of course required such that the organ-level presentation of the disease is broken down into the cellular components of the organ i.e., the epithelial cells.

The epithelium is the collection of epithelial cells that line and cover all organs in animals (as well as their skin). The stroma is the collection of cells (such as fibroblasts, adipocytes, and mast cells) that make up the main tissue of the organ and are embedded within an extracellular matrix (i.e., they are not necessarily tightly packed together). As far as susceptibility to decomposition goes, the two approaches begin to differ after this initial system boundary demarcation.

The classic SMT postulates that exposure to carcinogens will result in a DNA mutation in a single epithelial cell affecting key aspects of its cell-cycle resulting in hyperplasia (uncontrolled cell proliferation). Thus, the SMT takes the tissue to be decomposable into functional subcomponents – first epithelial cells, then nucleus of a single cell, and ultimately DNA being identified as the subcomponent responsible for the development of cancer in the tissue as a whole. This picture of decomposition and localisation is why I place the classic SMT just over the left-hand side of the continuum.

The TOFT treats the tissue as a less decomposable system. The focus of the TOFT is on epithelial-stroma interaction rather than decomposition and functional localisation within subcomponents. For the TOFT cancer occurs only at the scale of tissue organisation, thus decomposition of, for example, the epithelium cell, would be an inappropriate strategy. For the TOFT, very little relevant information could be gained from further decomposing and localising the system (the tissue) to the extent that the SMT suggests (Sonnenschein and Soto, 2011). The goal of the TOFT is to investigate the various interactions between the epithelial and stromal elements of the tissue that (a) enable the tissue to maintain its structural and functional integrity throughout constant change, and (b) identify those key pathways that, when disrupted, result in the controls and constraints on epithelial cells being removed resulting in hyperplasia. In other words, the TOFT focuses on identifying organisational constraints on the system that, when disrupted, result in a loss of functional capacity and ultimately failure to maintain regular control (a synthetic strategy).

As noted above, the measures of decomposition and dynamics are intimately linked. I have placed the SMT just short of the 'component dominant' spike. This is because, whilst dynamic interactions have a role to play in the SMT, the approach focuses on component-specific dynamics in favour of system wide dynamics – namely interactions between genes within the DNA of the cell and the intra-cellular machinery involved in their expression, rather than across the tissue as a whole. For the SMT, feedback mechanisms and perhaps even nonlinear dynamics play a role in the theory but only in the sense that gene regulation networks (GRNs) and normal processes of cell growth signalling and apoptosis, for example, exhibit these properties. The quest to understand gene regulation networks and their relationship to carcinogenesis comprises a substantial area of inquiry in contemporary cancer research – cancer genomics.

Projects such as the NIH's The Cancer Genome Atlas (https://cancergenome.nih.gov), have sought to taxonomise different cancers by compiling their complete genome, with the hope of developing novel treatments, detection methods, and prevention strategies (Stratton, et al., 2009; Tomczak, et al., 2015). Considered in isolation GRNs are highly dynamic and integrated systems. However, when placed in the context of the SMT, they instead highlight the component dominance of the SMT.

Recall that, the level of abstraction must be fixed for comparison on the continuum and for approaches to carcinogenesis this is the level at which different cancers are categorised – the tissue or organ.[11] Given the fixity at the tissue-level, the SMT is component dominant because the study of highly dynamic GRNs represents a focus on dynamics within components (i.e., within the cell nucleus) rather than between the components (epithelium-stroma interactions).

The TOFT ends up past the interaction dominance spike, in part, because it places the nonlinear feedback interactions between stromal and epithelial components at the forefront of the conceptualisation of the target system as a whole. For the TOFT, the primary focus of inquiry is on the cell-wide mechanisms rather than individual cells. For the TOFT, then, DNA mutations are an effect of serious disruptions of feedback mechanisms in the tissue architecture whereas for the SMT, DNA mutations are the cause of such disruptions.

The case of cancer genomics highlights an important point when thinking about system-wide dynamics more generally. The term 'system-wide' is, of course, relative to the level of abstraction at which one is working and thus an informative use of the term requires the level of abstraction to be fixed before any comparisons can be made between approaches. If GRNs were defined *as the system* of inquiry for cancer research, then that system would indeed be highly dynamic and integrated. GRNs would sit far to the right of the continuum of research strategies; further perhaps than the TOFT itself (see 4.4 for a discussion of MAPK Cascades, a kind of intra-cellular protein signalling pathway that belongs on the far right of the continuum).

11 As testament to this claim, even for the Cancer Genome Atlas project almost all 33 types of cancer investigated are classified according to the tissue/organ in which they present, and then further sub-classified at the molecular level. Squamous Cell Carcinoma is an exception, although the project focused only on squamous cells collected from the mouth, nose, and throat. Full list is available here: https://cancergenome.nih.gov/cancersselected.

However, in cancer research GRNs are a subcomponent in the overall system and thus their role in the overall study of cancer must be understood as tracking within component dynamics. GRNs may well be studied in isolation from the rest of the system but their contribution to the system will ultimately have to be re-contextualised (put back into the context of the system) in order to understand any given type of cancer. Studying subcomponents as if they were isolated from the system is a hallmark of reductionist strategies (Kaiser, 2015, pp. 225–229).

The fact that projects such as the Cancer Genome Atlas have risen to prominence in the field is indicative, not that the genome is the level at which cancer ought to be characterised, but that GRNs are excellent subcomponents to isolate from the system and study. The differences between approaches to cancer research concern what role the subcomponent of GRNs plays in the system. If we start from the left – the SMT – they play a larger role and as we move to the right – towards the TOFT – the output of that subcomponent is considered to be of less importance to the production of the system-level development of carcinogenesis, than the interactions between cellular components of the system.

As for control mechanisms, for the classic SMT DNA is pulling all the strings. Cancer is the result of the functional activity of an individual cell, itself driven by a specific mutation(s), thus the development and spread of cancer throughout the tissue is centrally controlled and organised by DNA expression. Recall that for the SMT the default state of the cell is quiescence. So, cancer development is a gain of function process for a single epithelial cell – that cell has gained control of a process that will ultimately lead to the destruction of the tissue. In contrast, the TOFT places focus on the breakdown of system-wide interactions that normally maintain the integrity of the tissue architecture. The failure of system constraints leads the epithelial cell to return to its default state of proliferation – no gain of function, just loss of control. In other words, the control mechanisms that are relevant to the development of cancer in the tissue are distributed across the tissue-wide constraints that are normally in place and not functionally localisable to particular sub-components within the tissue.

3.4.2 The Direction of Cancer Research: Hybrid Approaches

So far, I have given an analysis of the research strategies of the SMT and TOFT according to which the research strategy associated with the classic SMT has a more reductive character than the one associated with the TOFT. Importantly, this is not the same as claiming that the SMT takes a reductive approach where the TOFT is holistic or non-reductive in terms of its research strategy. Rather, ‘reductive’ and ‘non-reductive’ are, under this framework, essentially comparative terms. This is a subtle but important point. It affords a move past a strict standoff between the two theories. Rather than envisioning cancer research as comprised of two entrenched approaches – one reductive, one non-reductive – we can use the markers of the classic SMT and the TOFT to map the width of approaches taken in the field to understanding the target phenomenon of carcinogenesis. The classic SMT and the TOFT mark boundaries on the continuum within which a whole plethora of research strategies are utilised. The classic SMT marks the edge of the more reductive approaches, whilst the TOFT carves a boundary at the less reductive end. We can expect most of the approaches in the field to fall within this boundary, and fewer to fall outside of those markers. Thus, marking the SMT and the TOFT on the continuum has two purposes, firstly to illustrate how the continuum can be used to analyse and compare different approaches to the same target phenomenon, and secondly, to chart the territory of approaches in the field using the SMT and the TOFT as boundary markers within which most research in the field will fall.

The framework can be used to do more work than analyse the research strategies employed by these two prominent theories and mark the boundaries of the approaches taken in the field. It also can be used to track, or give an indication of, the direction in which the field is developing. Recently, some philosophers have advocated for an integrated approach to carcinogenesis (Bedessem & Ruphy, 2015; Bertolaso, 2009; Morange, 2007; Plutynski, 2013). Concurrently, scientists have been developing ‘hybrid’ or integrated approaches to cancer research – hybrid in the sense of combining important insights from the classical SMT and the TOFT (Oh et al., 2015; Rosenfeld, 2013; Rückert, et al., 2012). My claim is that the framework I offer can help make sense of these developments in terms of the character of the research strategies they adopt.

It can demonstrate how the field is moving in the direction of exploring non-reductive strategies (moving to the right of the continuum) whilst still maintaining a central role for the important insights gained from more reductive approaches like the classic SMT.

Calls for an integrated approach are motivated by a continuing series of landmark experiments that seem to show that both approaches have serious merits. On the TOFT side of things, Maffini et al., (2004) showed that whilst stromal exposure to the carcinogen N-nitrosomethyluera (NMU) resulted in tumour growth in the mammary gland of mice, NMU exposure to the epithelial cells did not result in tumour growth, strongly indicating the primary role of stromal organisation in tumour growth and casting doubt on the role of DNA mutation in the epithelial cells. Greenman et al. (2007) found isolated so-called 'zero-mutation tumours', tumours in which no DNA mutations occurred. Zero-mutation tumours are not easily accommodated on the SMT, but the TOFT can do so for similar reasons to the Maffini study. Furthermore, Mally and Chipman (2002) demonstrated the possibility of 'non-genotoxic carcinogenesis', tumours that developed with chemicals known not to affect DNA such as chloroform and p-dichlorobenzene. Once again, the TOFT can accommodate such a possibility by noting the effect that these chemicals have on gap junctions. Gap junctions are intercellular pathways through which various cellular components and electrical signals are transferred (electrical synapses, for example, transmit signals over a narrow gap, i.e., a gap junction). Clearly gap junctions will be crucial to the maintenance of tissue organisation, so disruptions that affect the flow of signals through gap junctions which result in the development of cancer will be a central aspect of the TOFT's research, but will be neglected and possibly missed by the SMT (See Baker, 2015 for a rundown of seemingly TOFT supporting experiments).

On the SMT side of things, some studies have shown that certain mutations have a very high prevalence across cells in certain types of tumour. Vaux (2011) uses such studies to argue that these systematic, rather than random, correlations are left unexplained by the TOFT but align well with what the SMT predicts.[12]

12 On a cautionary note, this issue remains contested. On the one hand intra-tumour heterogeneity (ITH) is a major area of research, with some claiming that genetic (as well as epigenetic and phenotypic) ITH is now widely recognised across many major tumour types; putting some pressure on Vaux's argument (Gay, et al., 2016; Marusyk, et al., 2012). On the other hand, see footnote 13 below.

Stephens et al., (2011) present data consistent with the SMT's claim that carcinogenesis is the result of a single catastrophic genetic event. On the clinical side of things, whilst new target therapies are on the increase that take a view more akin to the TOFT (Baker, 2009; Bissell & Hines, 2011), current clinical successes based on SMT influenced approaches give large weight to the plausibility of the theory given what is ultimately the main goal of cancer research. Probably the most famous cases of clinic successes so far were the development of two drugs that specifically target individual protein kinases to inhibit their activation of pathways that result in uncontrolled cell proliferation: imatinib (often marketed as the drug 'Gleevac') and Trastuzumab (usually marketed under the name 'Herceptin').[13]

The combination of successes and shortcomings in both approaches motivates moving beyond a dichotomy between the SMT and the TOFT – beyond a stand-off between reductive and non-reductive approaches – and towards hybrid or integrative approaches to carcinogenesis. It is worth noting that the classic SMT (as I have called it) has of course continued to develop. Robert Weinberg's book *One Renegade Cell* (1998) served as the textbook for the classic SMT but, of course, Weinberg has continued to update the theory in response to some key experiments (Hanahan & Weinberg, 2000, 2011). In a reflective recent article Weinberg concludes:

> "So, perhaps ironically, we have come full circle, beginning in a period when vast amounts of cancer research data yielded little insight into underlying mechanisms to a period (1980–2000) when a flurry of molecular and genetic research gave hope that cancer really could be understood through simple and logical reductionist thinking, and finally to our current dilemma. Once again, we can't really assimilate and interpret most of the data that we accumulate" (R. A. Weinberg, 2014).

13 Gleevac targets an overactive tyrosine kinase protein (a signalling molecule); caused by a 'fused' *bcr-abl* gene, which is formed through the fusion of two chromosomes (22 & 9). This gene is present in chronic myelogenous leukemia (CML) white blood cells (Pray, 2008). The drug has been widely used against CML ever since a trial conducted by Druker et al., (2001) in which the drug was found to be effective in staggering 53 of 54 participants. A five-year follow up found an 89% overall survival rate after 60 months of use (Druker et al., 2006). Herceptin targets a receptor protein for the growth hormone HER-2, which can be overactive in HER-2 Positive breast cancer. A more recent development has been BRAF inhibitors for the treatment of metastatic melanoma (Hodis et al., 2012; Bruno et al., 2016).

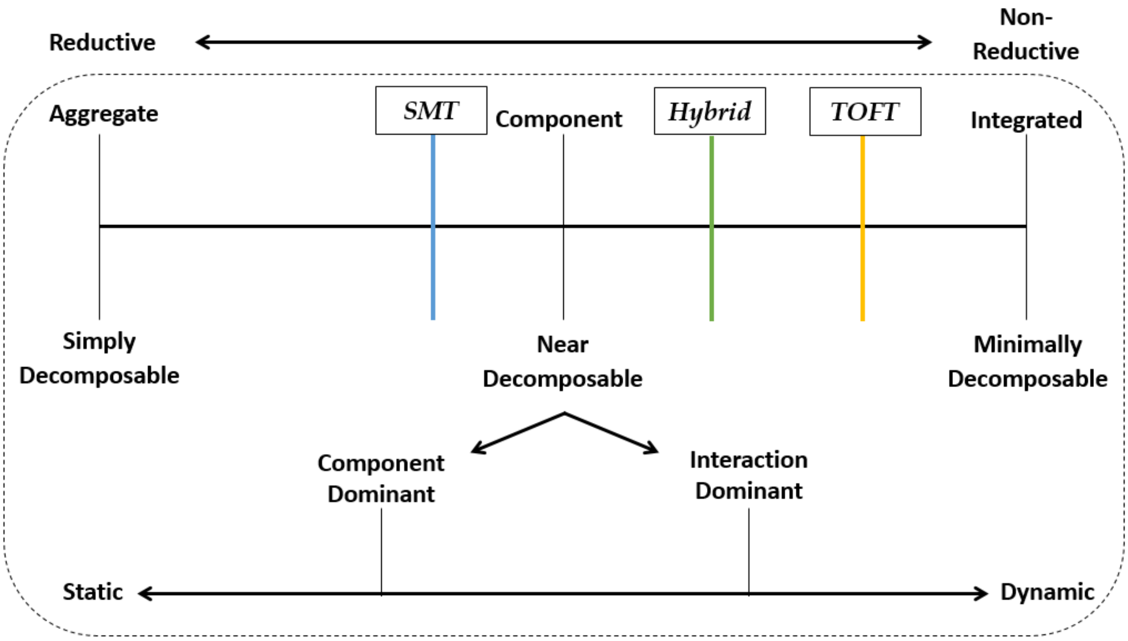

Fig. 8 Hybrid theories marked on the continuum. Figure from Baxendale (2018, p. 18).

Weinberg's message is strongly indicative of a move away from more reductive methods but also comes with a cautionary message to newer 'systems biology' (less reductive) approaches not to fall fowl of overconfidence in a single approach to the complex phenomenon of cancer development. In other words, at least as I interpret Weinberg, to pursue concurrent multiple strategies and integrative approaches where possible. Indeed, it is plausible to claim that contemporary versions of the SMT (Hanahan and Weinberg, 2011) have shifted towards the right of the continuum as interaction between neoplastic cells and the tissue architecture within which they proliferate has a bigger role to play in the theory.

Weinberg's message is heeded in new attempts to integrate the classic SMT and TOFT. Here I will give two examples. The first is Rosenfeld's (2013) use of 'Self-Organised Criticality (SOC)' as a framework within which to embed both approaches. In short, 'criticality' are situations in which small changes to a system may result in changes or avalanches of all magnitudes. SOC corresponds to a feedback mechanism that ensures a system is, and remains, in a state of criticality.

Common examples for explicating SOC include phase transitions in physics, wild fires, landslides, crowd stampedes and, saliently, bird-flock organisation (Rosenfeld, 2013, p. 223). Rosenfeld's approach conceptualises the precancerous tissue as a system in a critical state on the verge of an avalanche in the sense meant by SOC. To understand how metastable states – precancerous tissue – succumb to critical collapse we need import from both the SMT and the TOFT: "the mutant cell capable of starting the domino-effect of subsequent failures should be able to overcome the tissue's natural defences; this may happen only if the tissue is already preconditioned for failure and resides on the verge of systemic collapse" (Rosenfeld, 2013, p. 228). So, under SOC, genetic mutations can be a driving force for carcinogenesis (SMT) but they cannot do anything unless they occur in a tissue architecture already on the verge of collapse (TOFT).

A further example of new approaches includes the Feedback Model (FBM) of carcinogenesis. The FBM seeks to analyse the origins of tumour growth by showing that a normally negative feedback loop between the epithelium and stroma is transformed into a positive feedback loop, allowing for the maintenance of an inflammatory cellular environment, which in turn explains the uncontrolled proliferation of epithelial cells (Rückert et al., 2012).

The FBM is certainly towards the non-reductive end of the spectrum but it is not explicitly committed to the TOFT. Interestingly, its proponents do set up the FBM in contrast to the SMT, but also recognise the important role that genetic mutations may play in intercellular signalling (Rückert et al., 2012, pp. 2-3). In a similar vein Oh et al. (2015) (2015) sought to investigate breast cancer in terms of abnormal 'co-expression networks' mitigated by feedback mechanisms between the epithelium and stroma otherwise known as 'epithelial-stromal cross talk.

On Fig. 8 I have marked these hybrid approaches as falling between the classic SMT and the TOFT. When it comes to conceptualising the decomposability of the system, they both take the tissue-system to be certainly more decomposable than the TOFT on account of making room for genetic mutations as a driving force. Thus, a certain amount of decomposition and localisation is involved here. However, the localised functions of these cellular subcomponents are not sufficient to give an understanding of how cancer begins and develops.

For the SOC view, features of the environment are crucial to understanding the tipping point at which meta-stable states (normally functioning tissue) suffer catastrophic collapse.

As for dynamics, I've placed hybrid approaches just in and around the 'interaction dominant' spike. By placing normal and abnormal feedback dynamics at the centre of its approach, the FBM is certainly interaction dominant, as is SOC given its focus on the dynamic relationship between mutated cells and the environment in which they find themselves. However, they both sit a little further towards the reductive end of the spectrum than the TOFT. Whereas the TOFT black-boxes the dynamics of GRNs, these integrated approaches take the insights of research on GRNs, and subcellular interactions more generally, to be of vital importance to understanding carcinogenesis. Furthermore, for the SOC control is not as clearly distributed as for the TOFT, given that it affords a role to both DNA and environment.

Similarly, whilst feedback indicates nonlinearity in the FBM, this doesn't preclude genetic mutations from being a more centralised control mechanism given their role in intercellular signalling. This is only a broad-stroke analysis of these two approaches but it suffices to show that the continuum can be operationalised to provide analyses of different sorts of strategy in the field, and that it can be used to map the direction of research. I suspect that a more fine-grained analysis of individual hybrid approaches would place them at different positions within the boundary marked out by the SMT & TOFT. For present purposes, when placed within the framework offered here hybrid approaches are shown to be less reductive than the classic SMT but more reductive than the TOFT.

3.5 The Continuum of Research Strategies & the LCM

Developing the continuum of research strategies as a framework for thinking about reductive and non-reductive strategies has taken us far afield from the sort of discussions seen in the previous chapter concerning bridge-laws and identity relations, supervenience and multiple realisation, as well as anti-reductionism regarding explanations. To see where we've got to, I'll return to the lessons picked out at the end of the last chapter and elaborate on how the continuum of research strategies accommodates them.

> Lesson (1): Rather than relying on an account of reduction *or* emergence as a strategy to avoid the LCM, we need a new account of reduction *and* emergence that is explicitly incompatible with the LCM.

One of the main goals of constructing the continuum has been to dissolve precisely this distinction. Understanding that terms like ‘reduction’ and ‘emergence’ are *essentially comparative* concepts means that we jettison the need for definitive identification criteria for ‘reductive’ and ‘non-reductive’ strategies and in its stead we open up a whole spectrum of inter-level relations upon which to build a pluralistic conception of levels of organisation.

An example of this is the kinds of reduction that the continuum can accommodate. In the previous chapter I considered only *synchronic* reduction and emergence. The continuum can actually capture two senses of reduction and a further conception of emergence. First, it can capture synchronic reduction as system conceptualisations that are situated towards the left-hand side of the continuum – i.e., relatively more aggregative systems. Second, it can capture the *historical* sense of reduction in the sense of ‘theory replacement’ or ‘succession’. The example of the hybrid approaches falling in-between the canonical boundary markers of the SMT and the TOFT exhibits this sort of ‘reduction’. Perhaps the specific label of ‘reduction’ is no longer the most appropriate way to capture the idea, but the process which it stood for – i.e., theory ‘development’ – can certainly be captured on the continuum. The hybrid approaches explore the space in-between the boundary-marked approaches and may serve as successors to either view in the future. Finally, it can also capture the *dynamic* sense of emergence that has been the subject of recent developments in discussions of reduction (Bedau, 2008; Humphreys, 2016; O’Connor & Wong, 2005; Silberstein, 2006). Dynamic accounts of emergence focus on of a system property that *develops* over time and as the result of complex patterns of interaction at the lower-levels of the system. These sort of system conceptualisations sit further towards the right-hand side of the continuum.

On the whole, emergence is a much less controversial issue within the framework I’m offering here. Much like no strategy can be reductive *simpliciter*, neither can a strategy, or system behaviour, be ‘emergent’ *simplicter*.

There is no fixed point on the continuum at which system-level behaviours become emergent. Even if we came to some decision about where a threshold point was fixed, that system conceptualisation would still only be emergent *relative* to some other strategy that fell on the other side of the threshold on the continuum. This is not a set back for the continuum, far from it. Consider, for example, Mark Bedau's (2008) often-cited account of 'weak emergence'. The fulcrum of Bedau's account is that emergent properties of a system are of the sort that can only be investigated using complex epistemic tools like running simulations and observing the system evolve over time (Bedau, 2008, p. 161). On this account, weakly emergent properties are defined according to their accessibility to *us*, and the ways in which it we must investigate them. The continuum can be the basis of an analysis into precisely what features of the system result in these epistemic issues and suggest strategies for tackling them (for example, hunting for design principles that regulate robust state changes in a system). I don't think the continuum is in tension with Bedau's view. Rather, my point is that we don't lose anything by giving up the ability to strictly classify something as 'emergent' because we can still fulfil the *aims* of accounts like Bedau's, without that ability.

In sum, the views considered in the last chapter may have provided good reasons to be *anti-reductionist* about system structure, but they fell short of dismantling the system structure itself. The continuum of research strategies does precisely that by developing the framework for situating a plethora of different system structures relative to others. This will provide the underlying base for picking out different deployments of a levels of organisation concept that represent system structure differently. The continuum gives us a way to think about reduction *and* emergence concurrently on a sliding scale. To see how that account is also explicitly incompatible with the LCM we'll need to see how it has incorporated the other three lessons picked out at the end of the last chapter.

> Lesson (2): The account of reduction/emergence must provide a clear understanding of the role that aggregation plays in inter-level relations and a clear sense of how and when aggregation fails.

Aggregativity was a main component in the LCM and it proved a stumbling block for anti-reductive accounts, both ontological and epistemological.

This was not so much because they *were* aggregative, but because they couldn't be used to articulate how and when aggregation fails. I diagnosed this issue as stemming from a 'property-first' approach to inter-level relations. With the continuum of research strategies we move to a 'system-first' approach. The primary unit of analysis here is not metaphysical dependence relations between properties but modelling assumptions about building target systems of inquiry. This approach means we can both understand what sort of inter-level relation aggregation is and, by comparing it to other strategies on the spectrum, see progressively how and why systems will be conceptualised as less and less aggregative.

Specifically, the utility of an aggregative inter-level relation is inversely proportionate to the prevalence of system-wide dynamics and the break down of decomposition relations in the conceptualisation of the target system. The continuum allows us to see that this is a *gradual* process. Accepting the continuum as a framework means giving up on being able to designate an inter-level relation as 'aggregrative' or 'non-aggregative' without qualification (remember that the points on the continuum are only idealised paradigm points used to position strategies relative to one another). This is an important marker of incompatibility with the LCM; even if it was possible to apply a levels concept at a global scope – as per the LCM – we cannot designate that resulting system as *aggregative.* It can only be relatively aggregative in comparison with other possible conceptualisations of the target system, in this case, the world. Taken separately, this is perhaps a minor point of incompatibility, but presently it will be combined with what my account has to say regarding the issues of monism, distinctness of levels, and scope. At that stage a sharp incompatibility will have emerged.

To wrap up the issue of aggregation, it can also be noted that the continuum gives us the tools to recognise the role that relatively aggregative systems can play in scientific inquiry – that is, utilising an inter-level relation of aggregative part-whole decomposition – without falling back into a LCM-syle framework. This is because firstly, a 'fully aggregative' system will be a rare find in scientific inquiry, if ever utilised in system conceptualisation. Secondly, whilst *relatively* aggregative systems will figure in scientific inquiry, even they will surely not constitute a large swathe of system conceptualisations, let alone be the *dominant* way that inter-level relations are deployed.

Of course, I haven't actually plotted a relatively aggregative system on the continuum yet, but I will in the next chapter. Even without such a system plotted on the map, the analysis of the SMT and TOFT gives enough of a feel for how that process unfolds such that we can understand how a system might end up towards the left-hand side of the spectrum and thus be relatively more aggregate, component dominant, and reductive than strategies plotted to its right.

Finally, as just noted, considering the issue of aggregation lead me to reject a property-first approach and background causal relations between properties. An upshot of this was to focus instead of system dynamics, which I claimed to be more metaphysically neutral. Of course, this switch in focus doesn't make metaphysical issues of causal relationships between components of a systems simply disappear – granted. But my intention has not been to make such a claim. Rather, I've argued that if we're interested in developing a scientifically plausible account of levels then it is more fruitful to *start* with an account of inter-level relations that puts system structure at the foreground rather than metaphysical connections between properties. We saw in the last chapter how the latter approach (property-first) was, at best, inconclusive for rejecting the LCM and the purpose of this chapter has been to provide evidence for utility of the former approach (system-first).

There's a plethora of options available for exploring issues of causation after the framework has been established using a system-first approach. For example, one could investigate causal frameworks developed to illicit causal claims from a variety experimental procedures and modelling practices, such as the philosophically popular account of 'interventionism' (Pearl, 2000; Spirtes, Glymour, & Scheines, 2000; Woodward, 2003). Another option is to embrace the idea that the epistemic outputs that result from conceptualising the target system towards the right-hand side of the continuum, aren't best understood in terms of causal relationship at all. For example, an increasing topic of interest in philosophy of science is to explore *non-causal* modes of epistemic outputs – such as the system-wide dynamical patterns of interaction we've been looking at in this chapter – and their relationship to more traditional understandings of causal explanation (Reutlinger & Saatsi, 2018). Given the primary goal of developing a scientifically plausible account of levels of organisation – pursued within a practice-oriented framework – I will continue to background weighty metaphysical claims about the causal relationship between properties.

However, my hope is that the work done in this chapter and the remainder of the dissertation serves as strong evidence of the fruitful pay-off of approaching levels of organisation from a systems-first perspective.

Lesson (3): Monism & Distinctness of levels are mutually reinforcing aspects of the LCM

As with issues of aggregrativity, the anti-reductionist views considered in the previous chapter also struggled to distance themselves from the distinctness of levels and entities, theories, and explanations remained level-bound. As I noted there, demonstrating the failure of distinct levels would be an excellent way to avoid monism about levels of organisation concepts. To that end the plurality of inter-level relations displayed on the continuum will go a long way to avoiding distinctness of levels. It is true that in order to *compare* research strategies on the continuum it is necessary to *fix* the level at which the inquiry is to proceed. This was particularly important for identifying interaction or component dominance, as illustrated by the example of GRNs. But we do so solely for the purposes of comparing those strategies. This resultant levels in those system conceptualisations have no effect whatsoever on the relative structuring of objects in different systems. In effect, what the account of inter-level relations offered here does is flip the distinctness of level concept completely on its head. Rather than trying to show that objects appear on more than one level in a system, the framework suggests that it would be utterly implausible to suppose that objects *cannot* end up at more than one level. This is because in order to plausibly maintain that objects cannot be at more than one level, you'd have to assemble all system conceptualisations of all target systems of inquiry (maybe even possible conceptualisations too) and check to see if objects ever end up on relatively different levels. To put this a bit more abstractly, in order to maintain that distinctness of levels is true you'd need to show that for all systems S, in which collection of objects *X* is at level Ln; nothing other than collection of objects *Y* is at level Ln-1. Apart from being a presumably endless task, It seems staggeringly unlikely that you won't find a system in which collection of objects *Z* is at level Ln-1 precisely because there is a plurality of inter-level relations available to build system structures.

This point will become clearer in the next chapter when I delve into more detail about how those different inter-level relations can be used to generate different levels of organisation structures. This leads us directly into issues of scope.

Lesson (4): The account must provide an explicit position on the scope of inter-level relations.

Global scope was a largely implicit issue for anti-reductionist positions. There didn't seem to be any aspect of the critiques that spoke directly to issues of scope, let alone rejecting the sort of global scope offered by the LCM. In the previous chapter I stated that we should want a view about inter-level relations to explain why thinking about such relations in a global way is highly problematic, but without recourse to 'in principle' arguments that put strong limits on how models might develop in the future. I think that the continuum of research strategies affords a nuanced position on issues of scope than allow us to meet this aim.

Firstly, the continuum is explicitly local. The relative positioning of a strategy on the continuum is always relative to the carving of a system boundary and, as demonstrated by the SMT and TOFT, there are a plurality of ways to apply heuristics after that boundary has been fixed. Secondly, given that we have a framework for situating different kinds of inter-level relation it's hard to see how these could be combined and unified under one overarching system conceptualisation. It may well be possible to integrate one or two different system conceptualisations under a larger system. In fact, I'll argue that this is an important function of levels of organisation concepts on a *local scale* (i.e., between a few systems for a particular purpose). But, to show global scope we'd have to do this with every single system conceptualisation that pertains to every single target system of inquiry and that seems like a complete non-starter. However, thirdly, whilst it's very difficult to envisage a global overarching system capable of nesting all others within it, it's not *in principle*, impossible that such a model could be constructed. The continuum gives you the tools to see how that might take place, while concurrently demonstrating why, in practice, it seems highly implausible that systems could be integrated on a global scope.

Rejecting global scope is the last piece of the puzzle required to demonstrate that my account of reduction and emergence displays explicit incompatibility with the LCM. Let's recap. Firstly, I've not only rejected applying a levels concept at a global scope, but the continuum offers an explanation as to why global scope does seem so implausible, without relying on in principle arguments to do so – only local applications now seem plausible. Could the local applications nevertheless retain a *consistent* application: i.e., can we still maintain monism about levels? Well, if a failure of distinctness of levels is indeed the roadblock to monism that I've suggested it is, then pluralism is the only remaining route to take. This is because the continuum shifts the burden of expectation regarding distinctness of levels. Instead of assuming it to be a plausible feature of an account of levels, all possible system conceptualisations would have to somehow be compared in order to *check* that no individual ended up at two different levels; a result that seems highly unlikely. Finally, within this pluralistic setting, the continuum allows us to see that system conceptualisations that are indeed *relatively* aggregative will have a role to play in thinking about levels of organisation, but far from it being the *only* inter-level relation offered by the account, we now have a clear sense of how and when that aggregation will fail.

3.6 Conclusion

At the outset of this chapter we had a collection of lessons that needed to be addressed in order to avoid LCM-style thinking about inter-level relations. By the end I've developed a new framework for conceptualising inter-level relations, which will form a core part of my characterisation of levels. Evidently, a lot has happened in-between so let's pause here to take stock. First off, I focused attention on the right-hand branch of Fig. 2 where levels of disciplines connected with methodological approaches to reduction via the overlapping unit of analysis of epistemic practices. To clear the path for my own account, two steps were required from here. The first was to utilise recent work on epistemic practices to shift the focus from disciplines to *research strategies;* providing my own four-part account of what they consist of. The second was to move from a distinction between reductive and non-reductive strategies to a continuum between the two. To develop a continuum I focused on the application of heuristics during the conceptualisation of a target system of inquiry.

Specifically I focused on decomposition and system dynamics, as well as the relationship between the two. This provided a framework of system conceptualisations ranging from simply decomposable, non-dynamic (or relatively static) systems at one end, to minimally decomposable and dynamic systems at the other. A large chunk of the chapter then consisted in demonstrating the operation of this continuum using the detailed case study of cancer research. The continuum of research strategies offers a framework in which a plurality of inter-level relations can be understood and compared. I wrapped up by returning to the lessons we began with and discussing how the continuum meets these challenges. In the next chapter, the continuum will be embedded within a new characterisation of levels of organisation that further develops the aspects of research strategies outlined here.

4

Levels of Organisation as Tools for System Building.

4.1 Introduction

After navigating the broad features of the LCM, learning lessons about inter-level relations from some canonical discussions regarding reduction and emergence, and developing a new account of inter-level relations through a practice-oriented lens, I now arrive my attempt to put all of this work to use by building, explicating, and demonstrating a new characterisation of levels of organisation. It's been a long road since the General Introduction in which I provided the dimensions of evaluation for this characterisation of levels. Accordingly, I'll restate them here. The first dimension evaluates how well the characterisation captures deployments of the concept in scientific practice, that is the *scientific plausibility* of the account. Two important assumptions are embedded in this dimension. First, that we should expect variation in deployment of the concept in scientific practice and so the characterisation must be pluralistic. Second, that case studies drawn from the life sciences will be sufficient to establish scientific plausibility, particularly because they will be based on system *types*.

The second dimension evaluates how fruitful the characterisation is for the analysis of those practices. That is, the *analytic utility* that the characterisation can offer. In this context, analytic utility refers to what aspects of scientific practice an account of levels of organisation can help to illuminate and ultimately how such an account can contribute to a better understanding of the complex processes of knowledge production found in scientific inquiry.

As you might expect, my account of levels is designed to explicitly address these dimensions, and hopefully prove convincing relative to them. To do this, the characterisation comes in two parts, the first concerns scientific plausibility. Scientific plausibility will be tackled in two ways. Firstly, the account will be embedded within the practice-oriented approach developed in the previous chapter.

What this means specifically is that levels of organisation will be positioned as *representational tools* deployed as part of a research strategy (4.2.3). As was the case for developing my account of inter-level relations in the previous chapter, I'll jettison issues of epistemic outputs in favour of a focus on how the ideas embedded in 'levels-talk' are put to work in developing the structure of a system conceptualisation that reflects the goals of the overall research project. Secondly, in order to capture the wide-variety of deployments of the concept, I'll be utilising the continuum of research strategies (Fig. 6). My strategy will be to take a look at how the characterisation applies to inquiries that work with systems roughly around the three paradigm points on the continuum – an aggregative (4.3.1), component (4.3.2), and integrated system (4.4). These paradigm points provide the different system *types* I rely on for the generalisability of the characterisation.

The second part of the characterisation concerns the ways in which it can serve as a fruitful concept for analysing scientific practice. Specifically, I will be looking at the ways in which levels concepts can help us understand the integration, manipulation, and transfer of information of, and between, different system conceptualisations.

The chapter is broken down into two main chunks. The first will deal with carefully outlining as explicating each aspect of the characterisation I offer. I'll break the account down into it's simplest components starting with the concept of an organisational principle (4.2); and then vertical (4.2.1) and horizontal principles (4.2.2) respectively. In the second part I will apply the account to the case studies just outlined above and trace features of the account that overlap between these very different examples. In doing so, I'll be aiming to show both aspects of the characterisation *in situ* – it's scientific plausibility and it's fruitful role in analysing scientific practice.

4.2. Levels of Organisation: A Characterisation

The characterisation of levels of organisation that I will outline and defend as a scientifically plausible alternative to the LCM is as follows:

> Levels of organisation are collections of vertical and horizontal organisational principles deployed as representational tools for (a) conceptualising a target system of inquiry and (b) the manipulation, integration, and transfer of information, of and between systems of inquiry.

For the remainder of this section I will provide details about key aspects of this characterisation. There's a lot packed into this characterisation but we can start on familiar ground given the work done in the previous chapter: conceptualising a target system of inquiry. Conceptualising a target system was an important aspect of a research strategy that I introduced there. Let's quickly recap that notion as it's elements will be utilised again here. A research strategy, in my sense, contained four elements: (i) a well defined research question, (ii) a level of abstraction at which to pursue this research question (carving a system boundary), (iii) methodological assumptions, and (iv) theoretical assumptions. I suggested that two key theoretical assumptions (decomposability and dynamics), and the relationship between them, constituted a conceptualisation of a target system and, in turn, used these measures to develop a continuum of research strategies on which different system conceptualisations could be plotted (Fig. 6).

The characterisation of levels offered here broadens out and provides more substance to the way in which systems are conceptualised for the purposes of inquiry. It does so in two ways. Firstly, corresponding with (a), it groups heuristics into two broader categories: vertical and horizontal principles. Secondly, corresponding with (b) it, offers insight into how whole systems are utilised in progressing scientific inquiries. This aspect attempts to capture the *purpose-relativity* of levels concepts: they are always deployed relative to a purpose, which is derived from the aims of the inquiry (i.e., the well-defined research question).

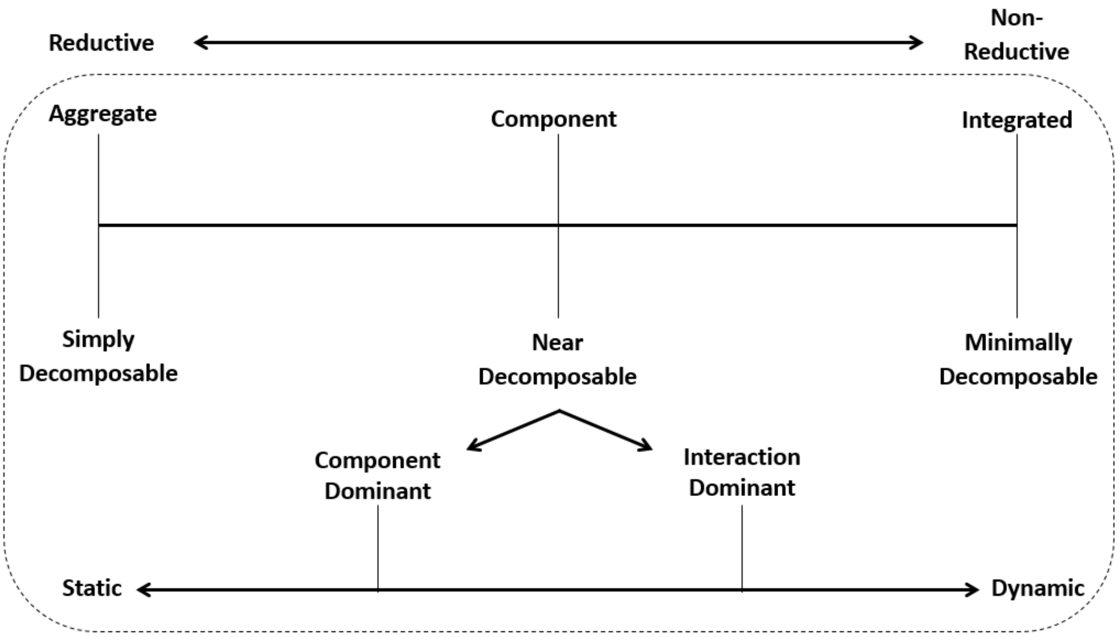

Fig. 6 The full continuum of research strategies. Inside the dashed box we now have a second dimension representing a focus on system-wide dynamic interactions; less towards the left and more towards the right. As in Fig. 5, the labels at which the arrows point represent *never reached values* – no system is thought to be *the* ultimately dynamical system. Like the dimension of decomposability, this second dimension also has vertical lines representing idealised paradigm points – component and interaction dominance – round which to orient analyses of research strategies. Further, these points serve to link the dimensions of decomposability and dynamics through the concept of a near decomposition. Figure from Baxendale (2018, p. 9).

In sum, then, levels of organisation concepts play a crucial role in how a system is represented and how different systems can be integrated towards a refined picture of the target system of inquiry. It follows from this that the characterisation of levels I offer here is purely epistemological/methodological, in line with the practice-oriented approach outlined in the previous chapter. In this sense, levels of organisation concepts are part of a range of epistemic practices that are applied to both build a picture of a target system and subsequently investigate it. This is all I mean by ‘representational tool’. Whilst the term ‘representation’ is a highly loaded one in philosophy, I intend it only to mean a device or strategy deployed to construct systems for the purposes of inquiry; that is, to represent them.

Practically speaking, that is often going to amount to ‘representational tools’ being objects, materials, principles, or rules, that are central to the construction of a *model* of a system and the subsequent manipulation of that model.

This could be understood in a very concrete sense like the sticks, balls, and pieces of metal that Watson and Crick used to literally build their model of the double helix structure of DNA. These tools can also be more abstract – as will be the case in the examples below – in such cases representational tools are going to amount to collections of principles, rules, or instructions, that are used to build, structure, and allow for the manipulation of, systems of inquiry. Simply put, if you want to investigate a system you'll usually have to find a way to represent it and in order to do that you'll need some tools. The vertical and horizontal principles that constitute a levels of organisation concept are important tools for representing a system of inquiry as exhibiting key structural and contextual features. For the remainder of this section I'll get clear on the basic aspects of the characterisation before illustrating them in operation through examples drawn from different points on the continuum of research strategies – i.e., different types of system.

The most basic idea in the characterisation is a *principle of organisation*; that is, after all, what I am suggesting that levels are. The most elementary description I can offer of a principle of organisation is that a principle of organisation imposes some form of *order* onto at least two individuals relative to a *purpose*. It might be thought of as a set of instructions for the relative positioning of individuals to one another, where 'positioning' can be taken literally, as the placement of individuals within a defined spatio-temporal boundary, or more conceptually as the placement of abstract relata in a conceptual space. In the previous chapter I used the language of 'partitioning' to describe this sort of process. A principle of organisation is, in that sense, simply a procedure for partitioning a system. That partitioning must be relative to at least some very minimal purpose, otherwise it would be hard to make any distinction between a random collection of individuals, and individuals organised according to a set of instructions.

Consider the way in which you should organise the contents of your fridge from a food-safety point of view. There are basically two principles that regulate fridge-organisation with the purpose of maximising food-safety: temperature and cross-contamination. Items are placed at different sections of the fridge depending on what temperature is best for them to be stored at *and* in order to optimise the avoidance of potentially harmful cross-contamination. Fridges are warmest in the door shelves and on the top shelf so cooked items, preserves and the like can be kept there.

The middle shelf maintains the most stable temperature so items more susceptible to temperature fluctuations should be place there such as dairy products. Finally, the bottom shelf is the coldest area of the fridge making them ideal for raw meats and so on. In addition to temperature considerations it is optimal to keep cooked and raw meats as far away from each other as possible, pushing them higher and lower respectively in the organisation of the fridge and never on the same shelf. So, here we have two principles of organisation – temperature of storage and cross-contamination – that regulate fridge organisation relative to the purpose of optimising food safety.

If one changes the purpose and principles the organisation can drastically change. Suppose the purpose of fridge organisation was based on aesthetic reasons such as *colour co-ordination*. In that case the principles would concern the grouping together of objects based on their colour and perhaps relative to other colours. So, for example, one could implement a 'traffic light' system such that not only were all green, orange, and red items organised on the same shelf respectively but also from top-to-bottom, resembling traffic lights. There seems no obvious reason to go for a traffic light organisational scheme in a fridge, but the point is that different principles of organisation can radically change the way in which objects are organised.

Of course, a standard fridge is a very different sort of system than the ones of interest in scientific inquiry. For a start, the principles of organisation under discussion here are concerned only with the best way to position individuals in an artificial system given a goal, such as food safety. For natural systems of interest in scientific inquiry the goal is very different. Most often it concerns understanding the functioning of the system or any number of its components, as well as understanding how that system relates to others in a larger net of interactions. The components of food items in the fridge contribute in no way to the functioning of the fridge. Similarly, the organisational principles do not reveal any particular interesting insights on the relationships between, for example, eggs and chickens stored on different shelves, nor between those items stored in the fridge and those stored in the dry cupboards (a different ‘system’).

The example is thus only supposed to illustrate some very basic ways of thinking about the concept of organisation. More specifically that (1) an organisational principle is nothing more than a way of positioning two individuals with respect to one another within some sort of structure, and (2) that an organisational principle is illuminating only insofar as one understands the *purpose* of its deployment.

4.2.1 Vertical Organisational Principles

Vertical organisational principles are, as you might expect, ones of *level differentiation*. They specify how two individuals end up at two different levels of organisation and in doing so, actually constitute the structure of the system conceptualisation – i.e. the relative positioning of individuals in a system such that one is 'higher' or 'lower' than another creates the levels in the system. So, vertical organisational principles are *inter-level relations* precisely of the sort we've been concerned with for the past two chapters. Accordingly, the account of inter-level relations provided in the previous chapter (partly) constitutes these vertical organisational principles. In developing that account – manifested as the continuum of research strategies – I specified three different decomposition relations, aggregative (which we'd seen from the LCM); near decomposition; and minimal decomposition. These were adequate for the purposes of getting the account of inter-level relations off the ground, but they do not exhaust the category of 'vertical organisation principles'. Vertical organisational principles is a general term for a plethora of relations that are used to position individuals in a system at relative 'heights' to one another and, in doing so, to develop a stratified system structure. To get more detail on how this works I need to pause here and comment on the relationship between three closely-related concepts that pertain to system structure: decomposition, hierarchy, and part-whole relations.

Probably the most ubiquitously associated concept with 'level of organisations' is that of *hierarchy*. So far in this dissertation, the concept of hierarchy has been conspicuous by its absence in any sort of substantive sense. I didn't use it to characterise the LCM (i.e., not in the four distinctive aspects of the LCM in box 2), to develop a new account of inter-level relations, and it doesn't appear in the characterisation of levels provided here.

Understanding why it doesn't appear is key to understanding what I precisely mean by vertical organisation principles. First, let's take a look at a well-known definition of 'hierarchy' from Herbert Simon, according to which:

> "By a *hierarchic system* or hierarchy, I mean a system that is composed of interrelated subsystems, each of the latter being hierarchic in structure until we reach some lowest level of elementary subsystem" (Simon, 1996, p. 184).

At first glance this explanation of 'hierarchy' seems familiar enough: a system with nested subsystems bottoming out at a lowest subsystem. But on closer inspection, this bare notion of hierarchy reveals very little about the structure of the system. Without knowing exactly what it means for a system to be composed of an 'inter-related' subsystem, we're left without much more than a description of individuals ordered from lower to higher. I think this point applies much more broadly across several familiar 'vertical' relations including part-whole relations and decomposition. It is often unclear what precisely 'hierarchy', 'part-whole', and 'decomposition' mean individually and importantly their relationship to one another is far from obvious: does each have a distinct and well-established usage, are they in fact just synonymous with one another, and is any one of these notions more 'basic' than the others? My answer to all three of these questions is no. There is not a distinct usage of each available and their meaning so opaque that we can't claim that they are synonyms. Instead, I suggest that you can take any one of these notions as a *primary unit* and sub-categorise the others *in reference* to it. We've seen this already in Chapter One. There, Oppenheim & Putnam described two different 'senses' of 'decomposition': wide and narrow. The wide one turned out to be aggregation or spatial containment, and the narrow one turned out to be mereological part-whole relations. I think this does capture an orthodox way to think about part-whole relations i.e., either as spatial containment or in the technical mereological sense. But it would have been equally coherent for O&P to take part-whole relations as a primary unit and develop a wide and narrow notion of 'part-whole'. i.e., the wide sense of *part-whole* means 'aggregative decomposition' and the narrow sense means 'formal decomposition' (mereological). After all, this is what the project of mereology is all about; taking 'parthood' as basic and analysing other kinds of relation in reference to it.

The same point applies to 'hierarchy'. In fact, Simon himself (1962) defined a broad and narrow sense of 'hierarchy', with narrow sense coming from engineering (control hierarchies) and the broad sense was his own relation of 'near decomposition'.

There is no fact of the matter about what these notions mean and without specification in a context, they don't mean much at all. 'Part-whole' doesn't *mean* spatial containment, it can pertain to a range of inter-level relations or it can pertain only to spatial containment. Similarly, as seen in Simon's quote above, the concept of 'hierarchy' alone doesn't mean much at all, not until we specify what kinds of inter-level relations fall under it. I contend that these the three notions are completely inter-changeable as a *primary unit* of analysis, at least conceptually. The important point is that one specifies explicitly which are them is taken to be the primary unit of analysis and what kinds of relations fall under them.

The lack of specificity around these terms is why I've chosen the label 'vertical organisational principles'. Seeing as these notions have little content before specifying the range of inter-level relations they cover, I might as well go for the broadest and least loaded label possible as a primary unit: hence, vertical organisational principles. The label of vertical organisation principles is also explicitly pluralistic, it's clear from the label that it applies to different kinds of relation, whereas it is often assumed that 'hierarchy', for example, refers to a much more specific relation (even though on closer inspection, it doesn't). As seen from the continuum of research strategies in the previous chapter, I've chosen to sub-categorise these principles mainly in terms of different kinds of decomposition, although I'll show in 5.4 that decomposition doesn't exhaust vertical organisation principles by considering inter-level relations of 'control'. I have chosen decomposition simply because it seems most apt when discussing the process of applying heuristics in *system conceptualisation*, which forms the basis of my account of levels. Perhaps only because it is already a prevalent term in literature on the topic (e.g., Bechtel & Richardson, 2010; Wimsatt, 1994).

With the triumvirate of hierarchy, part-whole, and decomposition, negotiated, I can introduce the vertical organisational principles that will be the focus this chapter. Broadly, we'll be looking at vertical principles of *decomposition* and *control*. By 'decomposition', I refer to the range of inter-level relations that were discussed in the previous chapter and that constituted the dimension of 'decomposability' on the continuum of research strategies (Fig. 6). Further, thanks to the work done in the previous chapter we can understand decomposition not in isolation but as it relates to system-wide dynamics. Doing this generated three kinds of system conceptualisation that drive a research strategy: full decomposability with minimal system dynamics (aggregative system), component-dominant near decomposition (component system), and interaction-dominant near decomposition (integrated system). In what follows I'll be looking at two examples that utilise a vertical principle of decomposition: an aggregative system (4.3.1) and a component system (4.3.2).

Although not always operationalised under the labels I'm using here, vertical principles of decomposition dominate discussions of levels of organisation in contemporary philosophy of science. As we'll see in the next chapter, many contemporary discussions of levels take it as an assumption that 'levels of organisation' must refer to decomposition relations between wholes and their parts (see 5.2.3 in particular). My account does not make this assumption and instead offers a broader understanding of levels of organisation as a concept utilised in scientific practice. Accordingly, we'll look at an example that exhibits an alternative vertical principle to decomposition: control (4.4).

Decomposition and control are not intended to be an exhaustive collection of vertical organisation principles. Indeed, in Chapter One, I discussed an interpretation of the LCM which understood the inter-level relations as *mereological* part-whole relations (1.4.2). As just mentioned, mereology is a formal philosophical project in which the 'parthood' relation is taken to be basic and an attempt is made to analyse other relations in reference to it. This might seem like a very technical and abstract sense of 'vertical organisation' but, as we'll see, the mathematical modelling of control relations can be similarly technical and result in a fairly abstract representation of the target system.

Another vertical principle that has played an important role in some classic notions of levels of organisation is s*ize* – that is ordering objects from smallest to largest. We find this most perhaps prominently in Churchland & Sejnowski's (1992) account of levels of organisation in cognitive neuroscience and Wimsatt's account of levels of organisation. Wimsatt (2007, p. 204), for example, argued that relative size ordering was a very good indicator of levels when combined with near decomposition. Recall that near decomposition identifies subcomponents in a system by mapping the strength or frequency of interactions across the system. Subcomponents are formed where interactions are stronger, or more frequent, within than between; i.e., forming clusters of interactions that can be identified as subcomponents.

Wimsatt suggested that size will play an important role in such decompositions because, whilst not always true, it is a good heuristic to suppose that individuals of a similar size will interact more frequently with one another than individuals of different sizes.

In sum, vertical organisation principles are heuristics that, when applied to a system, partition that system into two *levels* one ‘lower’ and one ‘higher’. The type of structure that develops depends on the type of vertical organisational principle utilised. Here I’ve introduced a non-exhaustive collection of such principles: decomposition, control, mereology, and size. From here on, I’ll be focusing specifically on decomposition and control. The application of such a principle (usually iteratively) forms a central part of the conceptualisation of a target system of inquiry.

4.2.2 Horizontal Organisational Principles

Vertical principles capture the core idea contained within levels of organisation concepts: the partitioning of a system into vertically arranged levels. However, this does not exhaust the content of levels of organisation concepts. To gain a full understanding of how systems are structured into levels we need to consider what I am labelling ‘horizontal organisational principles’. Very simply, just as vertical principles concern the partitioning of systems into vertical stratified levels, horizontal levels concern the partitioning of individuals *at a given level*. There are three senses in which an organisational principle can be ‘horizontal’ that I will consider here:

i. Intra-level partitions
ii. Specification of intra-level interactions
iii. Boundary carving.

I’m gathering all three of these conditions under the label ‘horizontal’ chiefly because I think that they play a crucial role in representing system structure but they are decidedly not of the same kind as the vertical principles just considered. The first sense is perhaps the most aptly named ‘horizontal’ principle as it captures the process are carving intra-level partitions. The idea of intra-level partitions is just to mark distinct groups of individuals at the same level of organisation.

This will crucially affect the further vertical partitions in the system as intra-level partitions effectively create new distinct 'branches' in the system structure, each of which may potentially be subject to further vertical partitioning (see Fig. 9 in particular). The carving of intra-level partitions is almost always going to be strongly linked to the purpose of the inquiry. Accordingly, it's difficult to say much about these in the abstract, rather we'll see them in operation in the examples below. What can be said, is that intra-level partitions further underline the idea the vertical organisational principles do not exhaust levels of organisation concepts. Investigating a system by applying vertical heuristics is going to give clumps of individuals at different levels of organisation. But, given the purposes of the inquiry there will be additional reasons why one would want to make demarcations *within* those levels and not just between them. These may be highly pragmatic reasons, such as only having the resources to investigate certain kinds of individuals within the level. They may also be more theoretical reasons such as isolating certain individuals that play very different roles within the system as a whole but that cannot be differentiated vertically from other individuals. The second type of horizontal condition concerns the patterns of interactions between individuals at the same level, which practically speaking will consist in specifying the techniques used to track those interactions. This condition will play a more prominent in role in cases of *system integration* and we'll see this principle in action in 4.3.2.

Boundary carving for a system is an extremely important aspect of any system conceptualisation. Essentially, the process of boundary carving involves deciding what processes, interaction, or individuals are to be taken within the system and which belong to the environment in which the system in embedded. No system boundary is absolute and a large part of the progression of an inquiry will consist in renegotiating this delicate balance between system and environment. One further consideration might be that the less that is included within the system boundary, the better. This is simply because the difficulty in isolating and controlling the variables responsible for the system behaviour of interest will increase in proportion to the amount of variables already operating within the system. If you can afford to push a variable outside the system boundary – that is, if you can remove variables without affecting the system behaviour – then it will yield a pay-off to do so.

Of course, a lot of the time scientists won't be able to push a variable out, but it is at least an example of the sort of considerations in play when thinking about drawing system boundaries. Decisions regarding boundary carving are often driven by pragmatic and creative considerations, rather than principled criteria, as well as entrenched carvings of systems that have been long established within a particular field of research. As such I think that claims regarding system demarcation need to be made on a relatively local and context-sensitive scale, with detailed input from a given inquiry into a system. Once again, then, the purpose-relativity of the deployment of a levels concept will provide more illumination on this point.

In lieu of a detailed and empirically developed account of the process of boundary carving, the main point I want to underline is that the process of demarcating system boundaries is not straightforwardly a matter of applying a vertical organisation principle to the environment. This seems neither an empirically nor conceptually adequate way of thinking about system boundaries. Empirically, it doesn't look as though systems boundaries are always demarcated in the same way as they are vertically partitioned. We'll see examples of this below, but again we can recall the SMT and TOFT to see this. Both systems were demarcated according to the presentation of uncontrolled cell proliferation (cancer). However, the two paradigm examples of the SMT and TOFT proceeded to apply different vertical heuristics to that same system, leading to markedly different system conceptualisations.

Conceptually, the idea would lead to some unusual claims about the relationship between system and environment. If system demarcation consisted in the application of vertical principles *alone* this would require a pre-existing system to partition in the first place. Presumably this system would be 'the environment' and applying a vertical heuristic would generate a *lower-level* subsystem. In such a picture, 'environmental variables' and the 'target system' would stand in a 'higher-lower' level relation (or vice-versa). But this seems odd. If this was the case, then all processes that are exogenous to the system would be at a 'higher-level' than the system itself (or vice-versa). The issue of the relationship between exogenous processes and levels in the system is complicated and potentially hazardous for the plausibility of any account of levels. It's an issue I'll return to in more detail in the next chapter, specifically with regards to the relationship between boundary carving and the *scope* of application of a levels concept (5.3.2).

For now I suggest that it makes more sense to separate the process of carving a system boundary from the process of partitioning the system itself into levels. I recognise that the specific label of a ‘horizontal’ condition might be less apt for boundary carving than it is for intra-level partitions; boundary carving is less obviously ‘horizontal’. However, my main point it that to understand boundary carving as a horizontal condition is to suggest it is a procedure whereby one clusters together processes, interactions, and individuals into one system *independently* of that system being partitioned into a stratified structure using vertical organisational principles.

4.2.3 Levels of Organisation and Research Strategies

According to the characterisation offered here, levels of organisation are representational tools and as such they are part of the epistemic activities involved in a research strategy. Of course, they do not exhaust such strategies but they are linked to each aspect. I’ll conclude the explication of my characterisation by making those links explicit. Firstly, the well-defined research question, which encapsulates the aims of the inquiry or at least operates as a proxy for the wider aims of the project, governs the deployment of levels concepts. Understanding the relationship between the aims of the project and the choice of levels concept will be a central part of my analysis below. The second aspect is straightforwardly a horizontal organisation principle as it pertains to carving a boundary for the system. A deeper issue for analysis, which I must bracket here, is the relationship between this aspect and others in the strategy. For example, how will the availability of data – including test subjects for example – and experimental techniques (methodological assumptions) impact on decisions regarding drawing the system boundary? Placing the carving of a system boundary into the broader context of a research strategy, opens up the links between it and a range of other epistemic activities.

The third aspect – methodological assumptions – was underdeveloped in the previous chapter’s work on building the continuum of research strategies. However, we can now see at least one role that methodological assumptions play in the conceptualisation of a target system by picking out the experimental techniques used to track intra-level interactions: a horizontal organisational principle.

Finally, the fourth aspect – the theoretical assumptions – continue to play the role of picking out inter-level relations that form a central part of the conceptualisation of the target system. In the context of levels of organisation, these assumptions now provide the basis for vertical organisational principles.

4.3 Levels of Organisation in Decomposable Systems

Now we have a broad (and rather abstract) characterisation of levels of organisation in place, the remainder of this chapter will consist in applying this characterisation across three examples. Let's recap that characterisation in here:

> Levels of organisation are collections of vertical and horizontal organisational principles deployed as representational tools for (a) conceptualising a target system of inquiry and (b) the manipulation, integration, and transfer of information, of and between systems of inquiry.

This work has a few purposes. The first is to demonstrate precisely how the different principles operate in different system conceptualisation and in doing so to illustrate clearly what the characterisation is supposed to consist in. The second is a direct consequence of the first: to demonstrate the scientific plausibility of the characterisation; showing that the account I have offered applies to a range of different types of system conceptualisation. Finally, the examples will show the utility of analysing levels of organisation as certain kinds of epistemic activities situated within the broader context of a research strategy. This will be particularly salient with regards to the second part of the characterisation, i.e., (b) the manipulation, integration, and transfer of information, of and between systems of inquiry. We'll see how levels concepts can contribute to the analysis of not only building one system of inquiry, but how they are involved in the integration of different systems and the transfer of information from one model to another. Each of the three examples below will follow a similar structure. I'll start by highlighting the purpose of the inquiry within which the structured models are developed.

I'll then pick out the vertical and horizontal organisation principles used to construct the model and, finally, demonstrate how these particular organisational principles (this particular levels concept) contributes towards the realisation of that broader aim of inquiry.

4.3.1 Aggregative Systems: The Human Microbiome Project

When discussing different kinds of systems Wimsatt (2007, p. 280) and Bechtel & Richardson (2010, p. 26) (2010, p. 26) argue that very few interesting systems are aggregative or composite ones, in the sense that Wimsatt's four conditions for aggregativity are rarely even met in natural systems.[1] However, *conceptualising* a system as at least sitting towards the aggregative end of the spectrum can, in fact, be a very useful strategy for the organisation of large amounts of new data into a structured system to be explored in further inquiry. A perfect illustration of such a case was the organisation of data in the Human Microbiome Project (HMP).[2] The project's goal was to map the genomic make-up of the microbiota living on and within the human organism, in order to establish possible new understanding of crucial processes within the human organism – digestion being of chief interest due to the quantity and diversity of human gut microbiota – as well as treatments for a range of illnesses and diseases, from skin disorders, to antibiotic resistant bacteria (e.g., MRSA), and even prenatal and neonatal diseases (e.g., neonatal necrotizing enterocolitis).

The most salient aspect of the HMP for present purposes was the sheer volume of data collected. If we include the ongoing successor to the HMP – Integrative Human Microbiome Project (iHMP) – a rough total of 14.23 terabytes of data has been gathered from gene sequencing of the microbiome, and researchers claim to have encountered between 81%-99% of genera of the microbiota to be found living symbiotically with healthy western humans (The Human Microbiome Project Consortium, 2012). How did the researchers refine and structure all this data into systems that could be clearly understood, shared across projects, and manipulated for a wide variety of experimental purposes?

1 I've discussed these conditions in several places already (1.4.2. and 3.1.1). To recap, these conditions were (1) the intersubstitutability of parts, (2) qualitative stability regardless of quantitative changes in parts, (3) stability under reaggregation of parts, and (4) only linear interactions among parts.

2 I refer specifically to the NIH Common Fund Project: the Human Microbiome Project (HMP) (https://commonfund.nih.gov/hmp/index) which took place from 2008 to 2013. It's successor the 'Integrative Human Microbiome Project (iHMP) (http://ihmpdcc.org/) began in 2014 and is ongoing. Or course the HMP and iHMP are not the only projects/researchers to investigate the microbiome, but it is useful to take these large projects as a point of reference.

My claim here is that, at least partly, this was achieved by the creation of systems structured using a clearly defined levels of organisation concept.

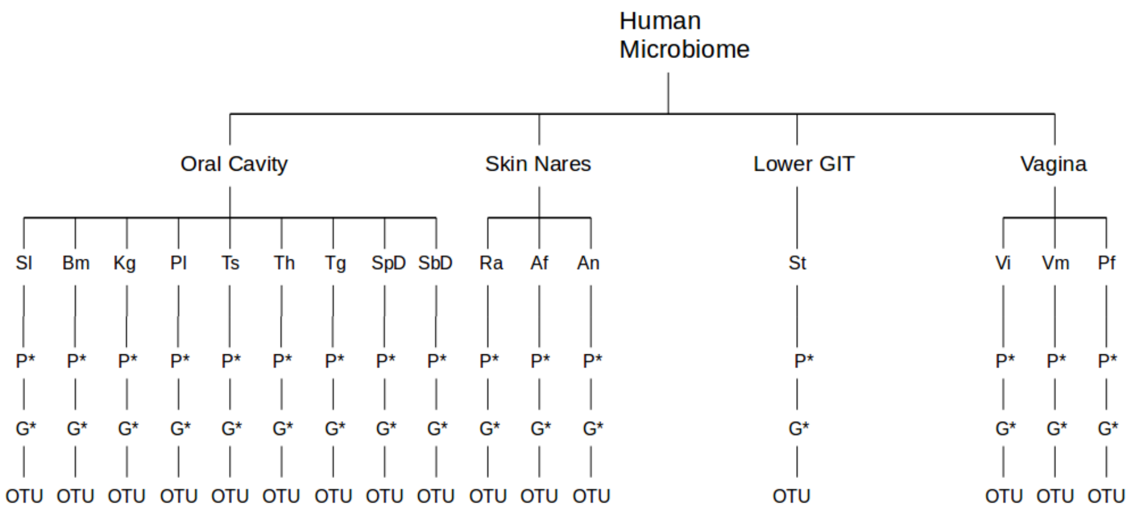

Fig. 9 The Human Microbiome Master Tree. A scaffold of the entire human microbiome dataset collected by the NIH common fund project 'the Human Microbiome project', divided by areas of the human body known to be home to a range of microbiota. P* stands for a collection of phlya, G* for a collection of Genera, and OTU stands for 'Operational Taxonomic Unit', which roughly corresponds to a species-level classification.

The HMP imported ecological taxonomic categories to organise the data yielded from 16*S* rRNA sequencing.[3] The initial data gathered by the HMP Consortium can be structured into a model with levels of organisation (Fig. 9), a structure I label the Human Microbiome Master Tree (HMMT). The first level is populated by major skin regions. One level down we find more specific skin sites within those regions. Then the ecological criteria kick in. Firstly, into a subdivision into phyla marked by P*.

[3] The technique consists in identifying 16*S* rRNA genes in a sample using relevant primers. Primers mirror the highly conserved nucleotide sequences allowing for researchers to 'match' primers to 16*S* rRNA genes; identifying and isolating them from a sample. Once copies of the 16*S* rRNA genes have been multiplied (using a technique known as PCR) the hypervariable regions can be analysed in order to identify different taxonomic clusters such as phyla, genera, and species. There are no universally accepted similarity thresholds for identifying different taxonomic groups but usually between 97% and 99% similarity is required for species identification, >95% for genus identification, and >80% for phylum identification (Janda & Abbott, 2007; Woo, Lau, Teng, Tse, & Yuen, 2008).

P* is a *set* containing any number of distinct phyla and each P* on the tree contains a different collection of phyla. The same notation applies to genera, i.e., G*, which are found at the next level down from phyla. I have collected the phyla and genera into sets for the purposes of fitting everything into one diagram but each branch will contain different quantities of these groups. Finally, 'OTU' stands for Operational Taxonomic Unit. An OTU is any unit prescribed for the organisation of individuals into a discrete unit according to a set of specific criteria for the purposes of taxonomy and classification. In the microbiome literature OTU is generally taken to represent 'species-level' classification. The terms OTU and 'species' are used interchangeably in the literature. Once again, each OTU on the tree represents a set of OTUs, each of which is different in constitution.

This is, at least roughly, an *aggregative* system in the sense that the vertical organisational principle used to partition the system is simple decomposition – i.e., that individuals at level n are fairly straightforwardly aggregates of individuals at level n-1. Take, for example, the branch that starts with 'Skin Nares'. Important to understand when reading this tree is that 'skin nares' does not denote swatches a skin, as if *skin* was a part of the human microbiome. Rather 'skin nares' delineates a spatial region of the human microbiome and it is the microbiota that reside in that spatial region that qualify as 'parts' in a constitutive sense. Each member of the level immediately below the skin nares – Ra, Af, An – collectively constitute the skin nares. At the next level down we see the broadest classification of microbiota – phyla, which in turn are comprised of more fine-grained clusters of bacteria: genera and OTUs respectively. One can take any vertical branch of the HMMT and see the structure of that branch understood as vertically organised levels, in which each level is an aggregate of the level below, for example:

(L6) Human Microbiome
(L5) Skin Nares Microbiome
(L4) Antecubital Fossa (Af) Microbiome
(L3) Phyla of Af Microbiome
(L2) Genera of Af Microbiome
(L1) OTU (species) of Af Microbiome

More broadly, one could classify the levels of the whole HMMT as follows:

(L6) Human Microbiome
(L5) Microbiome of Major Regions
(L4) Microbiome of Sub Regions
(L3) Phyla
(L2) Genera
(L1) OTU (species)

What about the horizontal principles at work here? As far as boundary-carving goes, this was done both with the most inclusive approach possible and to ensure that the areas included in the microbiome were those of most interest to studying the links between populations of microbiota and specific diseases or developmental issues. Evidence for this can be seen is the follow up project iHMP of which one strand in particular focuses on pregnancy. Specifically, the relationship between the microbiome of the mother and child throughout pregnancy, a project known as MOMS-PI (Multi-Omic Microbiome Study: Pregnancy Initiative).[4] Having the vagina included as a major region on the HMMT was crucial to lay the foundations for this project and related projects focusing on STIs and cervical cancer (Fettweis, et al., 2012; Huang, et al., 2014).

The upper levels of the HMMT perfectly illustrate horizontal partitions in operation. Vertical principles of aggregation will help to build the structure of the system but they cannot account for the selection of different skin sites intra-level, particularly at the highest level. These skin sites, whilst covering a broad range of sites, do not cover every possible skin site on the human body. Rather, the four sites are selected in reference to those areas in which the relationship between the population of microbiota therein is most strongly linked to site-specific diseases. As we’ve just seen the vagina is a salient examples of this, as is the lower gastrointestinal tract as evidence by the iHMP follow-up project investigating irritable bowel syndrome (known as The Inflammatory Bowel Disease Multi'omics Data project – IBDMDB). Finally, the horizontal principle that applies to lower levels is similarity of gene sequence at each level.

[4] Project overview available here: http://vmc.vcu.edu/momspi

Remember that P*, for example, represents a *collection* of phyla, for which *intra-level* partitions will be made based on genetic similarity via 16*S* rRNA gene sequencing. The horizontal principle of intra-level interactions plays less of a role here as it will in the next two examples. This is to be expected in an aggregate systems where, after all, *interaction* is not particularly relevant to the decomposition of groups of individuals from one level to groups at a lower level. Nevertheless, shortly we'll see why understanding that the taxonomic structure of the HMMT is developed through gene sequencing techniques is important for subsequent developments and refinements made to the system as a result of further investigation into the blueprint it provides.

This brings to a close the first part of the characterisation i.e., (a) the role of vertical and horizontal principles in conceptualising the target system and their links to the broader aims of the project. The second part of the characterisation, (b), concerns the role that these principles play when that system is further investigated and it's information is utilised and refined to create models that further the aims of the project. We've seen this already in the example of the two iHMP follow-up projects MOMS-IP and IBDMDB. To make this point clearer I'll draw attention to two quick further examples here.

As noted above, the purpose of the HMMT is to provide a blueprint for the project. Accordingly, individual inquires can be understood as interrogating the fine-grain details of the structure – investigating *one level* of the structure. For example, Conlan et al. (2012) conducted an inquiry at (L2) – attempting to provide a fine-grained mapping of the *Staphylococcus Genus* throughout the entire microbiome *(*Fig. 10A). Conlan et. al's study is salient here for two reasons. Firstly, it highlights the role of the HMMT as a scaffold that can be zoomed in on, and then zoomed out again, to refine the original model. *Staphylococcus* is a population of microbiota that is of particular concern in contemporary medicine give the rise of Methicillin-resistant Staphylococcus aureus (MRSA) infections. Zooming in to get a fine-grained analysis of the Genus across the entire microbiome is data that can be subsequently used to identify areas of the body particularly at risk given abnormally diverse or abundant populations of S. Aureus.

Secondly, Conlan et. al also used their study to compare different hyper-variable regions of the 16S rRNA gene to see which gave the highest resolution taxonomy. Thus, the study highlights how the methodological techniques used to construct levels can themselves be the objects of inquiry as well as the constituents of those levels.

Another example can be found in a study by Grice et al., (2009). Their interest lay much higher up the system than Conlan et al.'s Genus-level focus. They aimed to investigate different skin regions, thus setting their focus at (L5), and provide a detailed analysis from a large array of skin sites classified by environment type (dry, moist, and sebaceous). Both these analyses, shown in Fig 10A and Fig 10B, also produced stratified structures, adding ever more detail to the blueprint of the HMMT.

These two additional examples point towards the role that levels can play in (b) the manipulation, integration, and transfer of information, of and between systems of inquiry. The HMMT serves as a scaffold for a large project and individual inquiries interrogate aspects of that model at different levels creating refinements and specificity to the original large scale blueprint. In other words, the HMMT can be seen as a *structured repository* for the progression of a large-scale project – the HMP. The structure is provided by a levels of organisation concept, one which employs vertical principles of aggregation and full decomposition, and horizontal principles that including selecting sites for investigation and intra-level partitioning, as well as the specification of methodological techniques.

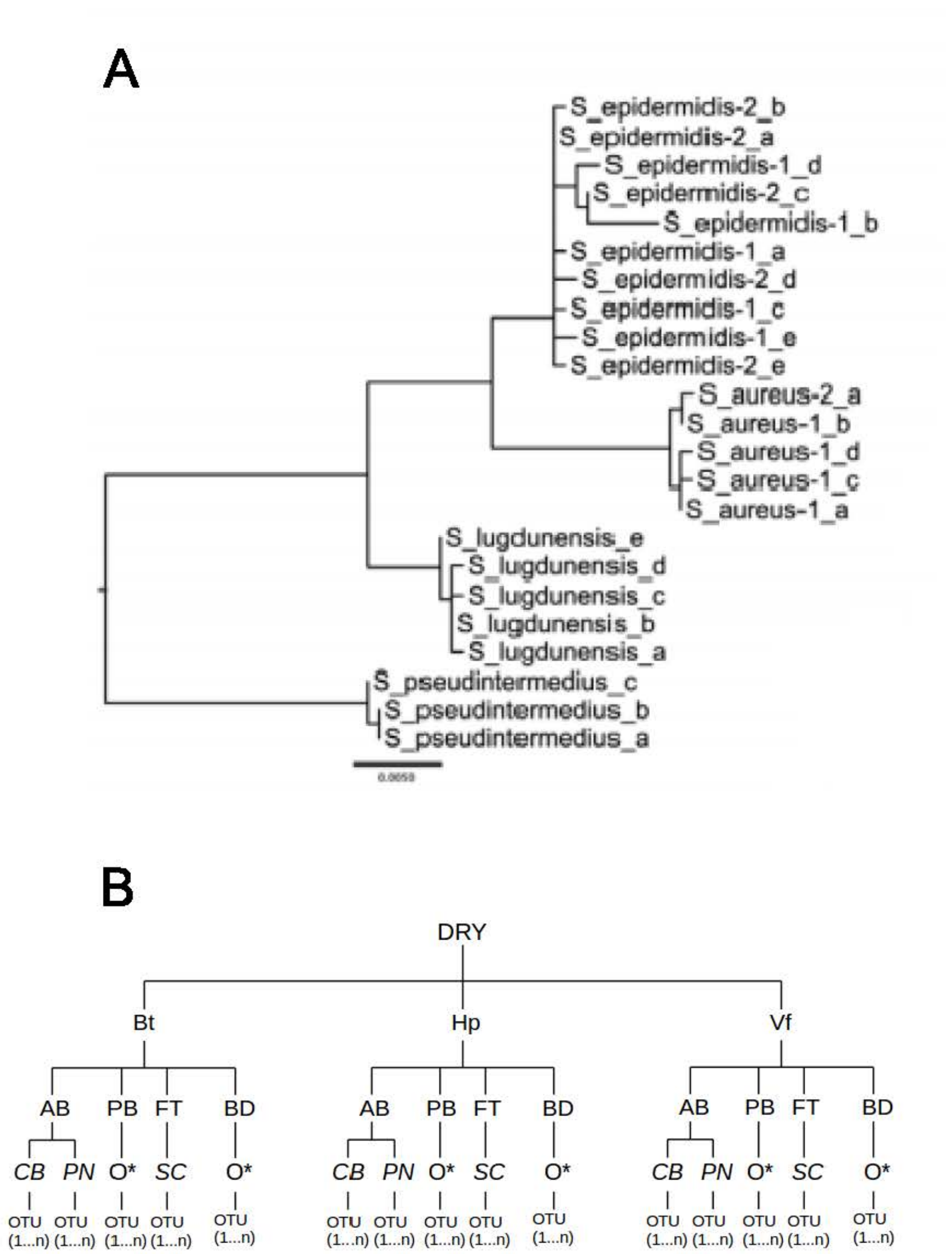

Fig. 10. Investigating the HMMT at One Level **A**: The full-length sequence of 16S rRNA gene for the Staphylococcus genus. Four species are shown on this tree for ease of presentation. Conlan et al., distinguished forty different species in the their analysis. Figure From Conlan et al., (2012 p. e47075). **B**: A Taxonomic tree for dry skin sites representing a study conducted by Grice et al., (2009). The four levels correspond to a specific skin site, microbial phyla, genre, and then species (OTUs). This figure shows only the ‘dry’ skin sites, a full representation could be constructed with all three sites at the highest level: dry, moist, and sebaceous.

At the end of Chapter One (1.5) I identified combination of a global scope, step-wise decomposition, and an aggregative system type as the essential tension at the heart of the LCM. However, I was careful to point out that this essential tension does not speak to these aspects of levels concepts taken individually; each needs to be assessed on its own merit. Now, within a local and pluralistic setting, we can see that aggregative system conceptualisations do indeed have an important place in an analysis of levels of organisation in scientific practice. Of course, in the examples just outlined we've been looking at systems of *taxonomy* and perhaps building a taxonomy is the only useful application of a levels concept which utilises aggregativity as a vertical horizontal principle. Even if this is the case, taxonomies are clearly extremely valuable for the organisation of large-scale data for huge research projects like the HMP and iHMP. Additionally, we've seen the role of horizontal principles and how they are often closely governed by the aims of the project, e.g., the selection of skin sites for sampling and their relation to prevalent site-specific diseases.

One final point to make regarding this example is that it begins to exhibit an important point that will be developed in more detail below. Within a single system the vertical organisational principle remains the same throughout each level, i.e., each inter-level relation is a result of the same vertical principle reiterated to create a new partition. On the other hand the horizontal principles can vary throughout the system, i.e., intra-level relations can be generated via the application of different horizontal principles within the same system representation (apart from carving the system boundary of course). This will be much more evident in the next example but it is still exhibited in the HMMT as the horizontal principles governing the top two levels (spatial) are different from those governing the lower levels (similarity of gene sequence).

4.3.2 Component Systems: Glucose Homeostasis

All the systems considered from the HMP were aggregate systems, conceptualised as simply decomposable without reference to dynamic interactions. As noted, they are systems of taxonomy and as such they served as clear examples of the features of my characterisation of levels. However, a characterisation of levels that claims to be scientifically plausible must surely be applicable beyond the construction of taxonomies and to systems that involve more complex processes of decomposition and intra-level interactions.

With that in mind, let's move on to consider a system that occupies the centre ground of the continuum (Fig. 6), namely a component system. The healthy human body is able to maintain a glycaemia (glucose concentration in the blood stream) basal range of between 0.8g and 1g per litre. The insulin molecule plays a large role in this homeostatic process. This process and the mechanism responsible for it, provides an excellent example of a system conceptualised as roughly a component one, displaying some system dynamics and subject to near rather than simple decomposition.

In a recent paper Tarik Issad & Christophe Malaterre (2015) use this example to argue that whilst a mechanistic understanding of this process can be illuminating for certain purposes, a full understanding of the phenomenon has required a mathematical understanding of key dynamic interactions within the cell. They propose a new framework called 'Causally Interpreted Model Explanations' that acts as a continuum upon which mechanistic explanations, dynamic mechanistic explanations, and mathematical derivations have relative positions and are operationalised for different purposes.[5] For my purposes, Issad & Malaterre's paper contains excellent diagrams of the processes involved in glucose homeostasis at different levels of organisation, which they themselves adapted from the multi-level models developed by Nyman et al. (2011). My presentation of the example draws largely from these authors, in which the relevant scientific sources are detailed in full. Here I use Issad & Malaterre's diagrams as they are slightly less technical than Nyman et al.'s and so easier to follow. As further reference I draw on details from a widely-cited annual review paper by C.R Kahn (1985) entitled 'the molecular mechanism of insulin action'. Let's take a look at the multi-level model.

[5] I should note that Issad & Malaterre's (2015, p. 245) model contains a fourth, lower-level which I have omitted here. The lower-level illustrates the extreme complexity of interactions at the molecular level and is utilised by them to argue that there is distinctive information at this level that will provide different kinds of explanatory import than can be gained from only the three levels above (i.e., information gained from the mathematical derivations section of their continuum rather than mechanistic end). Issad & Malaterre's additional level highlights how vastly difficult it is to integrate *all* the information available into one multi-level model. This mirrors the situation we saw in the previous chapter in which hybrid approaches to cancer research attempt to reconceptualise the target system so that this integration problem can be overcome. Thus, Issad & Malaterre's point illustrates the limits of integrating systems via near decomposition inter-level relations. My final example below will consider models that begin where these limits end, towards the 'integrated' end of the continuum of research strategies.

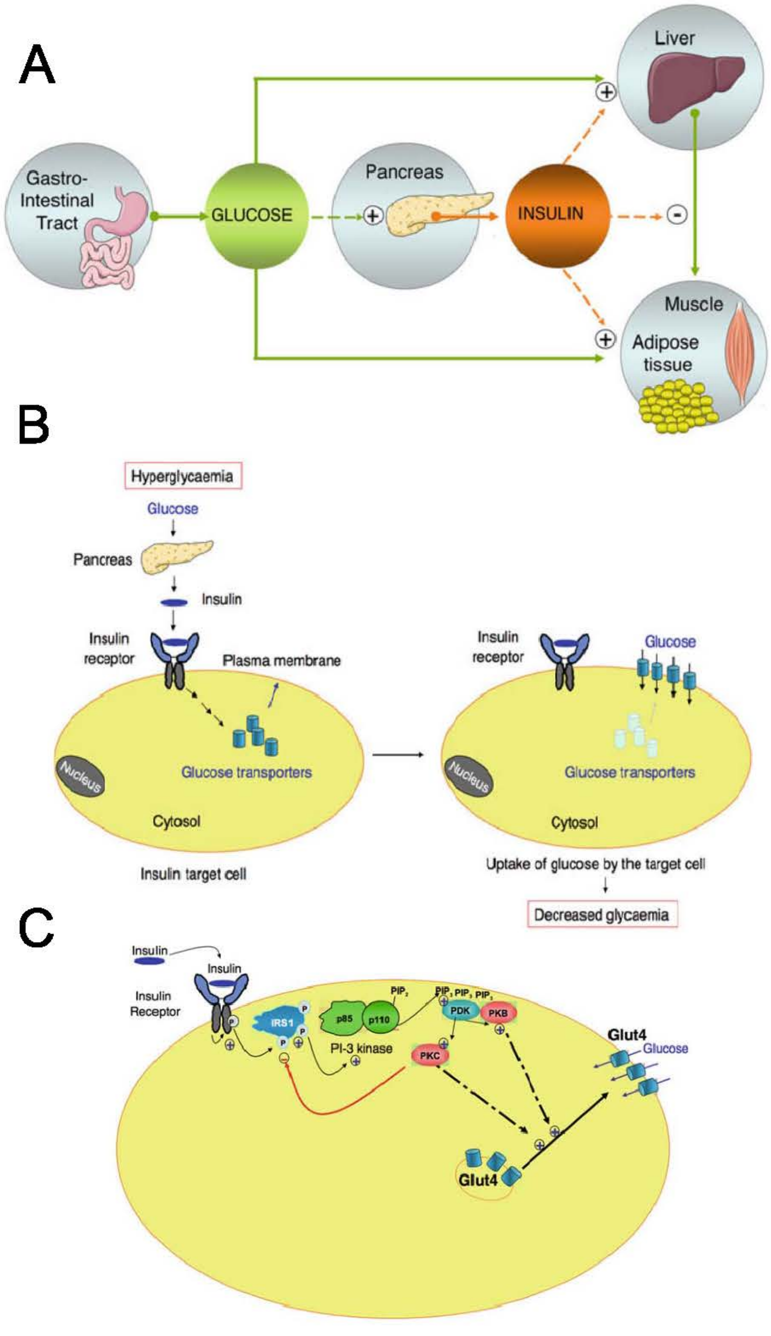

Fig. 11 The Whole Body Mechanism of Glucose Homeostasis. **A**: Highest level of the model showing the digestive system response to increase concentration of glucose in the system. The pancreas produces insulin which stimulated uptake of glucose in either the adipose tissue for storage or the liver to be metabolised. **B**: Intra-cellular interactions within the adipose tissue. Insulin molecules bind to the insulin receptor which activates, sending a signal to stimulate movement of glucose transport to the cell membrane. **C**: The insulin action mechanism. A protein signalling pathway showing the transmission of the signal from the activated insulin molecule to the glucose transporters . Figure combined from several from Issad & Malaterre (2015, pp. 271-273).

Fig. 11 shows the whole body mechanism of glucose homeostasis. Fig. 11A describes the highest level of model in which increased amounts of glucose coming from the GI tract and into the pancreas stimulate production of insulin. Insulin molecules stimulate uptake of glucose into the liver and adipose tissues, which results in a decrease in the concentration of glucose in the blood stream. The glucose is either converted and stored as glycogen, or metabolised. Fig. 11B shows the middle level of the model; zooming in to the activity occurring at individual cells within the adipose tissue. Here, an insulin molecule binds to a receptor, which in turn stimulates the movement of glucose transporter from the cell plasma to the cell membrane at which point glucose can be absorbed into these transporters and brought into the cell, decreasing the concentration of glucose in the blood stream. Finally, Fig. 11C shows the insulin action mechanism and describes the exact process by which the insulin receptor, complete with insulin molecule, stimulates the movement of glucose transporters via a complex chain of protein interactions. We can abstract away from the details of these mechanisms to see the multi-level structure of the whole process of glucose homeostasis:

(L3) The Digestive System
(L2) The Glucose Transport Mechanism
(L1) The Insulin Action Mechanism

The vertical organisational principle is decomposition but not in terms of the full decomposition seen in the trees of the HMP, rather it is the functions, or activities, of the parts that are relevant for the decomposition into lower levels. Rather than simply aggregating the individuals at one level in order to generate a higher level, the inter-level relations are built by considering what is responsible for a process, or set of interactions, at one level. On the highest level, we can see that increased insulin production ‘stimulates’ glucose uptake in the liver and the adipose tissue, but what is responsible for this uptake? Answering this question is a process of decomposing the process of ‘simulating uptake’ into the parts and their functions that afford this, namely, the binding of the insulin to the receptor and the subsequent activation of the glucose transporters. This process is reiterated to get us another level of structure – what is responsible for the ‘activation’ of the glucose transporters? This takes us into the lower level structure of protein interactions.

The vertical principle at play here can be understood as near decomposition – tracking strength or frequency of interactions takes us down a level at each key process and this is iterated to build levels of organisation. More specifically, in terms of the continuum of research strategies, we're talking about *component-dominant near decomposition*. Recall that component-dominant near decomposition refers to the localising of system behaviour to the functioning of discrete components. Whilst these components interact in ways that are essential for the production of the system behaviour, each component can perform relatively independently of the other components in the system.

Each level of the model features interactions between components and it is the processes that result from those interactions that are the subject of decomposition. The decomposition is still component – rather than interaction – dominant precisely because the level building process consists in decomposing the interactions *within* key components at the level above. Take for example the decomposition of Fig. 10B into Fig. 10C. In Fig. 10B we find an insulin receptor (component) sending a signal (process/interaction) to glucose transporters (components), who then move to the cell membrane (process/interaction). The *decomposition* that takes place is not of, for example, the insulin receptor itself, nor the glucose transporters, but the *process* of a signal being sent between the two: an an interaction. This yields Fig. 10C which consists in detailing the components and interactions that are responsible for the signal being sent from the insulin receptor to the glucose transporters. This is *near decomposition* because the strength or frequency of interactions *within* that signal sending process (the interactions in Fig. 10C) are of a greater magnitude than those between the insulin receptor and the glucose transporter (Fig. 10B). It is component dominant because, for example, the insulin receptor does not rely on the glucose transporter in order to send the signal: that part of the process is linear in Fig. 10B. As long as a threshold is passed in the receptor (i.e., it is activated by insulin molecules), it will send the signal, even if that signal breaks down along the way and fails to activate the glucose transporters.

Before adding additional points regarding vertical principles as well as considering the horizontal conditions, it's crucial to get a grip on the purpose of this multi-level model. Nyman et al., (2011) focus their research on understanding the failures of glucose homeostasis that result in Type 2 diabetes which "is due to both insulin resistance in insulin-responding tissues and to insufficient insulin release by the pancreatic β-cell."

This resistance is thought to be a result of insulin signalling failures. Most research into insulin signalling network has been conducted *in vitro (*cells isolated from their natural context) and *in silico* (computer-based) modelling. Whilst progress has been made in this area, Nyman et al., highlight that researchers have faced large difficulties in setting realistic parameters for these models such that using them to intervene on natural systems might be possible. Specifically, researchers have struggled to ‘scale-up’ the model of insulin action (Fig. 10C) to, match the glucose uptake of the adipose tissue *in vivo* about which high quality data has been obtained (Fig. 10A). As they note:

> “Such simple scaling was precluded because the in vitro cell-based data and the in vivo whole-body data had been obtained under fundamentally different conditions, such as the addition of insulin to cells versus consumption of a meal, with very different time scales and insulin concentration profiles over time” (Nyman et al., 2011, p. 26039).

Furthermore, from a clinical perspective – that is, diagnosis and treatment – Type 2 diabetes is a complex disease involving malfunctioning of a process of energy homeostasis across key components of the digestive system and any clinical applications of a molecular understanding of the insulin action mechanism must be understood in relation to these key components. Thus, Nyman et al.’s purpose in constructing the model is precisely to *bridge the gap* between in two different well-studied process extracted using different techniques and under different experimental conditions.

This purpose drives the integration of two different systems via the vertical organisational principle of component-dominant near decomposition. This is a subtle point because of course – in some sense – every system is the ‘integration’ of several systems that *could be* investigated in isolation from one another. However, there does seem to be a difference in purpose between developing a model of a system like the HMMT; a structure which represents one system at the broadest scale possible, and developing a multi-level model with the explicit goal of integrating two independently well-studied systems, like the multi-level model of glucose homeostasis.

Let's turn to the horizontal principles which are, as with the HMMT, strongly governed by the purpose of the inquiry. The boundary of the system is dictated by all those components relevant to a clearly identifiable homeostatic process: the maintenance of blood glucose levels between 0.8g and 1g in the human bloodstream. If an individual is relevant to that homeostatic process then it comes within the boundaries of the model. Input is excluded. For example, exactly what raises the glucose level is not important (eating a large portion of white rice, for example). Rather, what is important is to understand whatever quantity of input arrives into the system, the baseline levels of glucose are restored (presumably within certain thresholds for a functioning digestive system).

The intra-level interactions is were things get interesting for this particular multi-level model. The methodological techniques involved at different levels vary. At the highest level we have *in vivo* data about the absorption of glucose into adipose tissue. At the lowest, *in vitro* and in *silico* modelling yield an abundance of data about the complex intra-cellular interactions that facilitate uptake of glucose into the cell. As noted above, this lower-level data needs to be connected up to the higher-level mechanism if it can yield 'actionable' information: by which I mean information that is relevant to clinical diagnosis of Type 2 Diabetes and subsequent intervention. Thus, understanding the different horizontal principles used to develop these two levels is crucial to understanding how the model serves to transfer information that can be used to refine knowledge of the system as a whole and contribute towards manipulation of the system – i.e. (b) in my characterisation of levels of organisation. This further underlines why horizontal principles can vary within one system whilst the vertical principles remain consistent throughout. The variation in horizontal principles can track different kinds of interaction-patterns at different levels. Being able to accommodate varying horizontal principles *within* a consistent vertical organisational structure is precisely the creative and complex challenge facing researchers who construct such multi-level models.

An additional point to note about this model is that the purpose that L2 serves as a bridging-level between L1 and L3, means that individuals from both L1 and L3 end up on L2. Notice that the insulin molecule, as well as the receptor, is on both L2 (Fig. 11B) and L3 (Fig. 11C), whilst the pancreases appears on both L1 (Fig. 11A) and L2 Fig. 11B). This, I contend, is not a mere aesthetic choice to make the diagram easier to understand.

Rather, it reflects the *horizontal* condition that individuals can be brought into L2 if they are necessary to bridge the gap between the levels developed using different methodological techniques. If that means that some individuals appear on more than one level – going beyond mere decomposition to build levels, then so be it.

I'll wrap up this example with one final point. The vertical principle that applied to the multi-level model of glucose homeostasis was component-dominant near decomposition. It's worth using this to highlight an important point that may have been lost in the previous chapter. There, I argued that cancer research seemed to be moving towards exploring hybrid models that fell somewhere between the component-dominant SMT and the interaction-dominant TOFT. This claim applies *only* to that particular case study and I did not intend for it to be a generalised claim. The model of glucose homeostasis illustrates the importance of using component-dominant near decomposition to build inter-level relations in a system. Whether this vertical principle will be appropriate will depend entirely on the case at hand. In this case, it served the purpose of attempting to integrate information relevant to a process gained using varying experimental techniques into a multi-level model. It is may well be that such integrative modelling will outlive its utility when different aims are in play. If my discussion in 4.4 below is correct, then it certainly will. This only further underlines the requirement for a context-sensitive and purpose-relative pluralism in thinking about levels of organisation as representational tools for system conceptualisation and reinforces why it is central to the characterisation of levels I'm offering here.

4.4. Considering alternatives to Decomposition

To pick up where I just left off, let's consider a case where decomposition relations do seem to outlive their utility; a case from the furthest end of the continuum (Fig. 6). We'll be looking at systems that sit a little further towards that end of the spectrum than perhaps even the TOFT did in the previous chapter. For these systems, and integrated systems more generally, their susceptibility to only minimal decomposition means that if levels are to be an important part of their conceptualisation, then the vertical organisation principle will have to be an alternative to decomposition. What we'll see is that the vertical relation of *control* is more appropriate in such cases.

In biology, researchers have been conceptualising target systems as control systems – importing concepts from engineering and mathematics – since at least the mid-20th Century[6], although arguably the advent of the new discipline of 'systems biology' has made that relationship closer than it has ever been before.

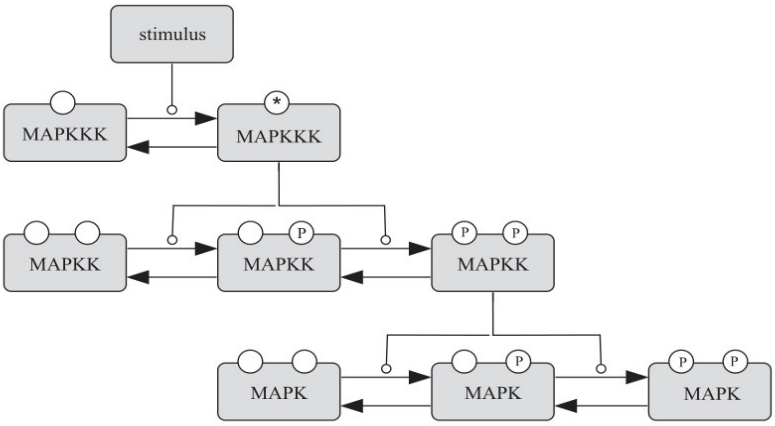

Fig. 12 Schematic representation of a three-level MAPK cascade. Circles with a 'P' inside denote phosphorylation groups. At each level, each individual box represents a stage of activation. At the first level, the kinase is activated by the stimulus (denoted with a '*'). The next two levels show two-stage phosphorylation required to pass the signal on to the the next kinase. Figure from Frey et al., (2010, p. 74).

An example of the sort of analysis that results from this approach is the model of the mitogen-activated protein kinase (MAPK) cascade, shown in Fig. 12. MAPK cascades are intracellular signalling transduction pathways.[7] The activation and deactivation of MAPKs plays a role in crucial cell processes such as a growth, proliferation, differentiation, motility, stress response, and apoptosis.[8]

6 See the preface to Iglesias and Ingalls (2010) for examples of early adopters of these approaches.

7 A 'kinase' is any enzyme (a macromolecule, usually a protein) that transfers (catalyses) phosphate groups from a high-energy molecule to a substrate. In Fig. 12 the 'stimulus' refers to the 'donation' of a phosphate group by a high-energy molecule. Through several stages of catalytic reactions, this group is passed down to a substrate molecule. The whole process acts as a signalling pathway from one molecule to another.

For example, mutations in the *ERK* pathway – a specific type of MAPK cascade – result in permanent activation of the signal that is being transmitted down the cascade, leading to out of control cell proliferation. Mutations on the *ERK* pathway are widely found in a variety of tumours and, as such, the pathway is a key Gene Regulation Network for cancer research, of the sort discussed in 3.4.1.

A simple sketch of what the model shows is that an external stimulus activates the first kinase in the chain – MAP3K. A signal is passed down the chain via a process of phosphorylations (the passing of phosphate groups from a molecule to a substrate). Eventually the MAPK passes the signal to target regulatory proteins that affect a wide-array of cell processes (not shown on the model). Some MAPKs – such as *ERK2* – can bind directly to DNA and thus act as transcription factors on certain genes (Hu et al., 2009).

The cascade is classified according to the lowest level kinase in the chain that passes the signal to the target regulatory protein. There are currently four known mammalian MAPK cascades – the *ERK1/2 cascade, JNK cascade, p38 Cascade,* and the *ERK5 cascade* (Plotnikov et al., 2011, pp. 1620–1621). Let's just focus on one of the cascades to get a sense of how the sequence unfolds. Isolating the ERK pathway, an Epidermal Growth Factor (EGF) can bind to the cell receptor, triggering transient ERK activation (resulting in cell proliferation). This activation consists in passing phosphate groups from one set of proteins to another. Starting at the highest level: Raf-1, and B-Raf; Rafs. Once activated, these proteins activate (via phosphorylation) the MAPKK proteins: MEK1 and MEK2. Finally, the process is reiterated to activate ERK1 and ERK2 at the MAPK level. This short introduction and the model shown on Fig. 12 are simplified versions of a MAPK cascade, but this level of detail will suffice for my purposes here.

Focusing on Fig. 12 specifically, Frey et al., are not seeking to investigate the effects of different MAPK cascades according their *constituents* –e.g., *ERK1/2* or *p38*. Rather, they are searching for (1) the consequences of different quantitative signalling assumptions (time and duration of signal), and (2) designs principles that certain structures realise (Frey et al., 2010, p. 75). Put simply, Frey et al., want to know whether they can illicit *structural* principles about the MAPK cascades, that do not depend on any specific constituent of a MAPK pathway.

[8] My explication of the MAPK cascades is a cumulation of the explanatory sections of Frey et al.'s (2010) paper as well as two review articles by Plotnikov et al., (2011) and Cargnello and Roux (2011). Any individual points are flagged with specific references.

Their method is to mathematically quantify varying degrees of phosphorylation (passing the signal from one group of proteins to the next). As you can see on Fig. 12, the MAPK cascade features activation by *double* phosphorylation of the MAPKK level and the MAPK level (the highest level is denoted simply as being ‘activated’ via the stimulus).

Frey et al.’s quantification of phosphorylation allowed them to compare activation by single, double, and triple phosphorylation whilst keeping the cascade length at three (Frey et al., 2010, p. 84). They found that double phosphorylation had a minimal product of time until activation and signal duration compared with single and triple phosphorylation, meaning that the double phosphorylation of MAPK cascades is well-suited to trigger a fast response to stimuli of a short duration (Frey, et al., 2010, p. 80).

Let’s take a step back from the experiment and consider what role the levels concept is playing here. It’s been a while since we’ve seen the characterisation so I’ll state it here again:

> Levels of organisation are collections of vertical and horizontal organisational principles deployed as representational tools for (a) conceptualising a target system of inquiry and (b) the manipulation, integration, and transfer of information, of and between systems of inquiry.

The most significant difference between a representation like Fig. 12 and Fig. 11 is that the vertical organisational principle is *control* rather than *decomposition*. The enzymes that comprise the MAPKK level are not ‘constituted’ by those at the MAPK level, rather those at the MAPK level are *subordinate* to those at the MAPKK level; their activation depends on the phosphate signal being passed on from the proteins at the level above. Constructing a model with inter-level relations of control is a vastly different kind of representational tool than using decomposition relations. It is a much more abstract kind of representation used to map phases of a complex process. Indeed, as far as Frey et al.’s study goes the qualitative features of the components are minimal at best; the system is almost component-less. They have only one function which is to pass on a signal once an activation threshold has been reached. This process can be modelled entirely mathematically without specification of features of components or different functions they can perform.

Frey et al.'s experiment is typical of a growing area of inquiry across the life sciences so-called 'systems' approaches. The mathematical models that describe the features of a three-tier control system that requires double-activation at each stage is a perfect example of what systems biologists call 'design principles'. (Green, 2015; MacLeod & Nersessian, 2015; Wolkenhauer & Green, 2013). The very purpose of experiments like Frey et al.,'s is to unpack design principles that govern certain kinds of systems.

This brings me to horizontal principles. We can see both the (i) and (ii) sense of horizontal principles in conjunction here. With regards to (i) the intra-level partitions made in Fig. 12 represent the stages of processing required *before* activation to a new lower level. At each level below initial activation there are three intra-level partitions, mapping the requirement to trigger activation by *double* phosphorylation. If the process required triple phosphorylation, then four intra-level partitions would be made in the system. That these partitions reflect stages of processing reflects the techniques used to investigate the system i.e., (ii) tracking intra-level interactions. This highlights the very different roles that horizontal principles can play in the conceptualisation of a target system. Recall that in the first example considered here – the HMP – the intra-level partitions had most significance when thinking of the HMMT (Fig. 9) as a blueprint for the large-scale research project. The partitions made in the first level – spatial skin regions – reflected the different lines of inquiry that researchers focused on as the inquiry developed. Here, however, the horizontal conditions are more indicative of the distinctive methodological techniques used to investigate these sorts of systems.

To summarise, the vertical organisational principle of control, combined with the horizontal principles that represent stages of activation and techniques used to investigate the flow of signalling through the system, nicely illustrate part (a) of my characterisation of levels: conceptualising the target system of inquiry. If we were to plot Fig. 12 on the continuum of research strategies it would be placed somewhere pretty far to the right-hand side of the continuum. Importantly, and unlike the TOFT considered in the previous chapter, it would be placed beyond the spike of *interaction-dominant near decomposition*. The control relation used to develop structure in this model is *not* a kind of near-decomposition. There's no sense in which the interactions (the catalytic reactions of phosphorylation) at the MAPK level are of greater frequency or strength than those at MAPKK level.

Rather the former is *subordinate* to the latter and it's demarcation as a 'lower-level' reflects it's dependence on the level above for activation. Furthermore, with respect to part (b) of the characterisation they help in understanding how *manipulations* of that system will be attempted. This is particularly salient with respect to the horizontal organisational principles. As we've seen, Frey et al., fixed the cascade length at three (vertical structure) and manipulated signal length and duration across single, double, and triple phosphorylation models – i.e., they manipulated both the stimulus and the *intra-level* partitions in the system in order to draw conclusions about the features of three-tiered, double activation systems like the MAPK cascade.

To end on a more speculative note, I think there is an additional interesting role that the levels concept might play in examples like the MAPK cascade. In addition to aiding understanding of the manipulations of the system, I think there's a sense in which the levels concept might play a role in the *transfer of information between* highly abstract models of control systems like Fig. 12 (another aspect of part (b) of the characterisation). Several philosophers have discussed the role of 'horizontal model construction' (Bokulich, 2003; Hesse, 1966; Zuchowski, 2017). Zuchowski, for example, (2017, pp. 6-7) gives two examples of such modelling practices: rule space parsing (systematically varying the dynamics of a system) and using genetic algorithms (meta-rules that adjust the dynamics of a model until it reaches a desired behaviour). Broadly, the idea is to develop a large repository of mathematical models that will describe many possible ways (perhaps all the possible ways) in which the dynamics of a system might unfold. Particular models in the repository can then be applied to a given target system to see if the model can serve as the basis for prediction and control of a behaviour of interest in that system; if not try another model and so on.

Put simply Frey et al.,'s analysis provides mathematical models of a three tier control-systems requiring single, double, or tripe activation. This double activation model that in fact maps Frey et al.,'s MAPK cascade might apply to a range of MAPK cascades (all four mammalian ones, for example), or maybe just one type (like the *ERK* pathway). Further, it might apply to other kinds of protein regulation systems within the human body or other kinds of system that are not protein regulation networks at all. The information being transferred is inextricably tied to the *structure* of the models that comprise this repository and, if what I've been arguing is this chapter has been convincing, the structure of the model is largely informed by the application of a levels of organisation concept.

In the example of glucose homeostasis, we saw how a vertical principle of component-dominant near decomposition could serve the purpose of *integrating* different systems investigated using different techniques.

I'm now suggesting that the sort of inquiry undertaken by Frey et al., gives us a horizontal version of this. Importantly, this sense of 'horizontal' is not itself a *control* relation; that remains a vertical organisational principle alone. Rather it is the *transfer* of information about control systems that is 'horizontal'. I'm not sure whether this kind of horizontal information transfer can apply only between systems that have vertical principles of control. I suspect not. But certainly what matters is that the *dynamics* of the system are the primary focus of the inquiry rather than, say, the discrete functioning of individual components because it is these very dynamical patterns that are the stored in the 'repository' and made available for application to many different models, for example Frey et al.,'s mathematical model of signal flow through a 3-tiered double activation system.

In sum, my point here is that there is a particular kind of modelling practice in which information transfer between systems conceptualisations is not achieved by *vertical integration* of different systems into a wider structure, in which the systems are nested within one another. Rather than integrating vertically, the system is taken to exhibit certain mathematically describable features – design principles or network motifs – that can be systematically manipulated to create a plethora of models that can be imported into models of other systems in an attempt to uncover the underlying patterns of interaction that result in their behaviour. In this sense, horizontal principles can operate between models to facilitate the transfer of information between models as specified in part (b) of my characterisation of levels of organisation.

4.5 Conclusion

The purpose of this chapter was to introduce my characterisation of levels of organisation and illustrate both its scientific plausibility and its role in the analysis of scientific practice through a range of examples.

To quickly recap, the characterisation was as follows:

> Levels of organisation are collections of vertical and horizontal organisational principles deployed as representational tools for (a) conceptualising a target system of inquiry and (b) the manipulation, integration, and transfer of information, of and between systems of inquiry.

The main task of 4.2 was to explicate key aspects of this characterisation, most notably vertical and horizontal organisational principles. The label ‘vertical organisational principle’ acts as an umbrella term for the partitioning of systems into two levels through the application of a heuristic. My main focus here has been on two types of vertical principle: decomposition and control. The horizontal principles came in three kinds: intra-level partitions, intra-level interactions, and boundary carving. The remainder of the chapter consisting in applying these principles to an array of examples.

The three examples took us across very different kinds of modelling practices and saw levels concepts utilised for very different reasons. A large part of my aim in this chapter was to demonstrate the scientific plausibility of my characterisation of levels by showing its applicability across these different kinds of inquiry. Picking systems that spanned the continuum of research strategies was my method for realising this aim. With this diversity in place, I want to conclude by picking out some more common, generalised aspects of the account of levels that we can see emerging across these different sorts of system conceptualisation.

Throughout the three examples we saw that the vertical principle used to construct the model remains consistent throughout whilst differing horizontal principles account for important, often inquiry-specific, features of each level. This was most clearly the case the example of glucose homeostasis but was also seen in the HMP models. I’m not sure it would serve any purpose to argue that this *must* be the case in order for levels to be scientifically plausible but it certainly seems to make a lot of sense that vertical principles remain fixed whilst horizontal principles can differ. The difference between these types of principles brings into focus a central challenge facing researchers across the life sciences: the goal of bringing together qualitatively different datasets into one coherent model. A clear example of this was found in the case of glucose homeostasis. New technology affords the development of increasingly complex models of dynamic interactions in fine-grained detailed. But these models, which often exist only on a computer and/or expressed in abstract mathematical motifs, must eventually be embedded into a biological system and that requires somehow situating that data in relation to much more coarse grain mechanisms populated by tissues, cells, and organs.

The application of levels concepts can be understood as system-specific or as meta-systemic. By system-specific, I mean that the vertical and horizontal principles are applied to build a model of one system, with the specific aim of investigating that target system conceptualised in a specific way, like the HMMT for example. By 'meta-systemic' I mean that they can also be applied with the explicit goal of transferring information across systems. This distinction corresponds to the two-part division of the characterisation i.e., (a) is system-specific, and (b) meta-systemic. For vertical principles this was exhibited in the example of glucose homeostasis and for horizontal principles in the MAPK cascades.

A closely related feature of the levels discussed is their role in building repositories. This was seen in two very different sorts of system: the aggregrative HMMT and the integrated MAPK cascade. With the HMMT, the model served as a blueprint for the aims of the HMP project. Individual inquiries can be understood as zooming in and out of that structure. At the complete other end of the spectrum the MAPK cascade structure in Fig. 12 is the starting point for a very different kind of repository; one of varying system dynamics, through the exploration of all possible parameters for the system (for example).

In sum, there can be little doubt that the characterisation of levels here is a long away from the LCM. Gone is the requirement for distinct levels, the application of the concept at a global scope, and the exclusive reliance on aggregative systems using simple decomposition. Instead, I have offered a characterisation of levels of organisation that is not only scientifically plausible in its pluralistic application, but that can help illuminate the different epistemic practices involved in building *structure* in a model of a target system. Whether that be as a primarily organisational (taxonomic) scaffolding structure (Fig. 10), a functional sense of structure (Fig. 11) or a more abstract structure of iterated control relations (Fig. 12).

5

Levels of Organisation in Philosophy

5.1 Introduction

Oppenheim & Putnam's layer cake model served as the entry point for my analysis of levels of organisation. Throughout the analysis, the views of several philosophers on levels have been touched upon but to this point not explored in any detail. This omission was in service of focusing as clearly as possible on developing a scientifically plausible characterisation of levels of organisation. In this final chapter, my aim is to situate the account I have developed in relation to contemporary discussions of levels of organisation. As it happens, philosophical conversations regarding levels of organisation are on the rise. Whilst, as we've seen, the topic has never been completely abandoned despite the pervasiveness of the LCM (Churchland & Sejnowski, 1992; Feibleman, 1954; Guttman, 1976; Novikoff, 1945), there has been a marked increase in engagement with levels of organisation in philosophy of science over the past ten years. This engagement has tended to come in two flavours. The first has been a new brand of levels scepticism (Eronen, 2013; Ladyman, Ross, Collier, & Spurrett, 2007; Rueger & McGivern, 2010; Thalos, 2013) All of these philosophers have questioned the usefulness of the concept – both in science and philosophy – in the wake of the clear inadequacies of the LCM. Some have sought to explore alternative concepts that might play a more useful role, such as scale (DiFrisco, 2017; Eronen, 2015; Potochnik & McGill, 2012).

The second type of engagement has sought to advance new, alternative accounts of levels of organisation that can both overcome the problems of the LCM and offer a positive contribution to the philosophical analysis of scientific inquiry (Brooks, 2016; Craver, 2015; Povich & Craver, 2018; Kaiser, 2015). In the General Introduction, I labelled this group as 'levels revisionists' and it will be the revisionists that are in the foreground here, although it will remain important to avoid the important problems pointed out by the levels sceptics.

More specifically I will be concerned with the view that has been described by even one of its chief critics as “the most coherent and scientifically plausible account of levels of organization to date” (Eronen, 2015, p. 40): mechanistic levels of organisation.

The primary goal of discussing the relationship between the account of levels I’ve been developing in this dissertation and the mechanistic account is to highlight and clarify distinctive features of my view; those aspects that provide a significant contribution to contemporary philosophical thinking about levels of organisation in scientific practice. Part of achieving that aim is to show how my account offers a fresh perspective on levels of organisation, whilst incorporating key insights picked out by the mechanistic account. In short, this will involve showing that the mechanistic account can be interpreted as analysing the structure of a certain kind of system conceptualisation: a component dominant nearly decomposable system. Of course, this will involve making suggestions regarding the interpretation of mechanistic view, which may or may not be convivial to its original proponents. Specifically, I will argue that the mechanistic account has come a long way from the problematic aspects of LCM-style thinking about levels, but not quite far enough in three respects: (1) in terms of its pluralism about levels of organisation, (2) it’s heavy-handed approach to ‘scope’ and (c) maintaining an *ontological* (rather than representational) understanding of levels of organisation.

Dialectically, I am not invested in the mechanistic framework and so my purpose here is not to *defend* the mechanistic account of levels from criticism per se. However, I take it as a positive result for my framework if I can show a space for other analyses within my own account. Additionally, I’ll argue that the subsumption of the mechanistic account under my own can actually provide compelling responses to the primary criticism raised against the mechanistic view: namely that it is too restrictive and local in its application. If the proponents of mechanistic levels of organisation were so inclined, they could harness my broader framework to resolve a key complaint against their view.

A large part of the ‘reinterpretation’ I undertake will involve disentangling the important claims of the mechanistic account of a levels from a few problematic assumptions that are a result of the account’s embeddedness in the broader mechanistic framework. The ‘embeddedness’ of levels concepts is not a peculiarity of the mechanistic account.

In Chapter Two, we saw the close conceptual ties between the LCM and the unity of science project. In Chapter Three, this was further reinforced by discussing the connection between the LCM, the unity of science project, and Nagelian theory reduction. In a recent paper, Brooks (2017, pp. 150–153) highlights the double-edged sword of the ‘embeddedness’ of levels concepts. On the one hand, understanding the embeddedness of levels concepts within a broader conceptual framework can help to illuminate the purpose-relative use of that particular levels concept. However, on the other hand this very same asymmetrical dependence of mechanistic levels on the conceptual framework of mechanistic analysis (for example), creates obstacles to positioning the account as a *general* and *pluralistic* analysis of levels of organisation in the life sciences. As mentioned in 1.5, this observation forms the core of Brooks’s ‘guilt-by-association’ argument against the levels sceptics, according to which many criticisms of ‘levels of organisation’ turn out to be criticisms of either the LCM or concepts within which it is deeply embedded, such as reduction. Indeed, in Chapter Two we saw just how complex this embedded relationship turns out to be, not only for reduction but also emergence.

My primary objective in this dissertation has been to develop precisely such a general and pluralistic account of levels of organisation. Therefore, in this final chapter, showing how my characterisation has a broader range of application than the mechanistic account *without* losing important insights offered by it, will be significant step in the right direction towards exhibiting the kind of generality and pluralism I am aiming for.

5.2 Mechanistic Levels of Organisation: The New Default Conception of Levels

In the recently published SEP article on levels of organisation in biology Eronen & Brooks (2018) describe the mechanistic account of levels of organisation as, “the standard view of levels in philosophy of neuroscience.” I think it’s fair to suggest that the account has become the standard view of levels – or at least the most discussed view – in the philosophy of the life sciences more generally, at least where explicit discussions of levels occur. Accordingly, I am orienting this discussion of contemporary views of levels around the mechanistic account.

The caveat being that of course it is not the only account one. Alternative accounts are offered by Kaiser (2015, pp. 175–185) and Brooks (2016) to name two. Although I won't be discussing the details of these alternate views, I will be discussing observations about, and objections to, the mechanistic account that stem from Kaiser's and Brooks's work, and in doing so, we'll get a sense of how their views diverge from the mechanistic picture of levels.

Whatever pros and cons the mechanistic account of levels may have it cannot be denied that its dominance in contemporary discussions of levels stems primarily from its association with the mechanistic approach to philosophy of science more generally, which itself has dominated issues in the philosophy of the life sciences over the past fifteen years or so. Indeed, the mechanistic account of levels was primarily developed by two of the most prominent figures within this mechanistic approach Carl Craver (2007, pp. 163–195, 2015; Povich & Craver, 2018) and William Bechtel (2008, pp. 143–148, 2017). With this in mind, it is both important and expedient to start by taking a quick look at the mechanistic approach to philosophy of science, before moving onto the account itself and the broader features of the view.

5.2.1 Mechanisms & the New Mechanistic Philosophy of Science

The 'new mechanistic approach' to philosophy of science is commonly associated with Machamer, Darden & Craver (2000); Bechtel along with various collaborators (Bechtel, 2010; Bechtel & Abrahamsen, 2005; Levy & Bechtel, 2013) Stuart Glennan (1996, 2002), and many others, with roots stretching back at least as far as Bill Wimsatt's (1976) and Lindley Darden's (1991) respective work.

In very general terms the mechanistic approach to philosophy of science posits that scientists (at least in certain disciplines such as neuroscience and biology) are in the business of discovering a specific type of phenomenon – mechanisms. The bulk of scientific practice thus consists in discovering that a certain mechanism is responsible for the production of a certain phenomenon. The bulk of scientific output consists in elucidating how those mechanisms operate in order to produce said phenomenon – that is, giving mechanistic explanations. A 'mechanism' here is a technical term of art. There is an abundance of definitions of this technical term in the literature and a wide-variety of discussions dealing with further tweaking of such definitions in the face of putative counter-examples or problematic cases (see Illari & Williamson, 2012 for a succinct overview[1]). The (at this point) classic definition of mechanisms is as follows:

> “Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions” (Machamer et al., 2000, p. 3).

So, for a given phenomenon there are entities (things) that, due to their interaction with one another (activities) in a specific order (organisation) regularly ‘produce’ (are responsible for) the occurrence of that phenomena within definite boundaries (set-up and termination conditions). A mechanism, then, can be conceptualised as a system consisting in a whole – the phenomenon of interest – and its various parts: the organised interacting entities (Fig. 13). The mechanistic account also provides a condition of *relevance* for components to belong to mechanisms. For a time, the most popular condition of relevance in the literature was Craver’s mutually manipulability condition, which states that a component *X* is relevant to mechanism *M* when a change can be brought about in *M* by manipulating *X*, and a change can be brought about in *X* by manipulating *M* (Craver, 2007, p. 153). More recently, proponents of the view, including Craver (2015, p. 15), have followed Couch’s (2011) suggestion of adopting Mackie’s (1974) INUS condition as a relevance condition between components and mechanisms.[2]

1 There is a staggering amount of literature on mechanisms in philosophy of science. As good a place as any to start is Craver & Tabery’s (2017) Stanford Encyclopedia Article. Glenann & Illari’s (2018) recent collection contains entries on just about every topic discussed within the contemporary mechanistic debate, including a foreword by Wimsatt giving a brief history of the development of contemporary mechanistic philosophy.

2 Nothing I have to say here will turn on whether one adopts mutual manipulability or the INUS condition as the condition of relevance between parts and wholes. But for the sake of completeness, under the INUS interpretation a part is relevant to a whole, when its behaviour is an Insufficient but Necessary part of a collection of parts, which taken together are Unnecessary but Sufficient for the behaviour of the whole. Perhaps of note is that both mutually manipulability and INUS conditions are tied to the idea of tracking – if not causal relations – then at least significant patterns of interaction, leading credence to my claim below that the mechanistic analysis involves the strategy of applying the heuristic of near decomposability to a system.

The mechanistic account can be used to give a narrative about the progression of scientific inquiry as the gradual filling in of mechanism (Craver & Darden, 2013; Darden, 2008). Mechanisms can come in various forms, the most minimal of which are mechanisms schemas: the abstract description of a mechanism. Slightly more detailed mechanisms are labelled mechanism sketches: mechanisms with functional components described but without fully described entities or activities, in other words with ‘black boxes’. Finally, an instantiated mechanism schema in which the system behaviour of interest has been decomposed into specifically organised entities to which the full complement of functions required for bringing about that system behaviour have been localised (i.e., their activities elucidated). The mechanistic picture has thus become a powerful framework through which to discuss the explanatory practices of scientists and to facilitate discussions of these practices with deeper metaphysical questions (e.g., Kaiser & Krickel, 2017).

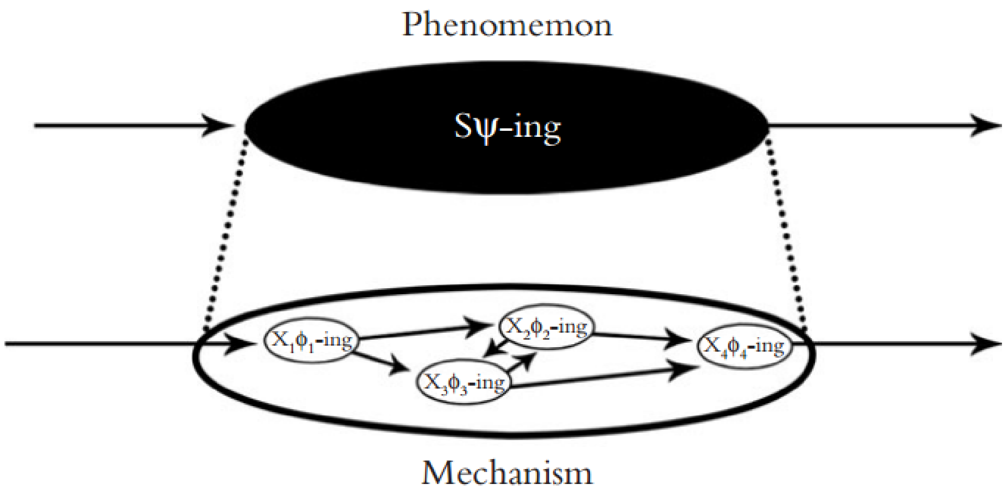

Fig. 13. A Mechanism. The top level represents the phenomenon of interest, which is a process or behaviour (S ψ-ing). The bottom of the mechanism contains the constituent parts (‘working components’) that collectively interact (activities) in a specific way (organisation) to produce the phenomenon of S ψ-ing. Figure from Craver (2007, p. 7).

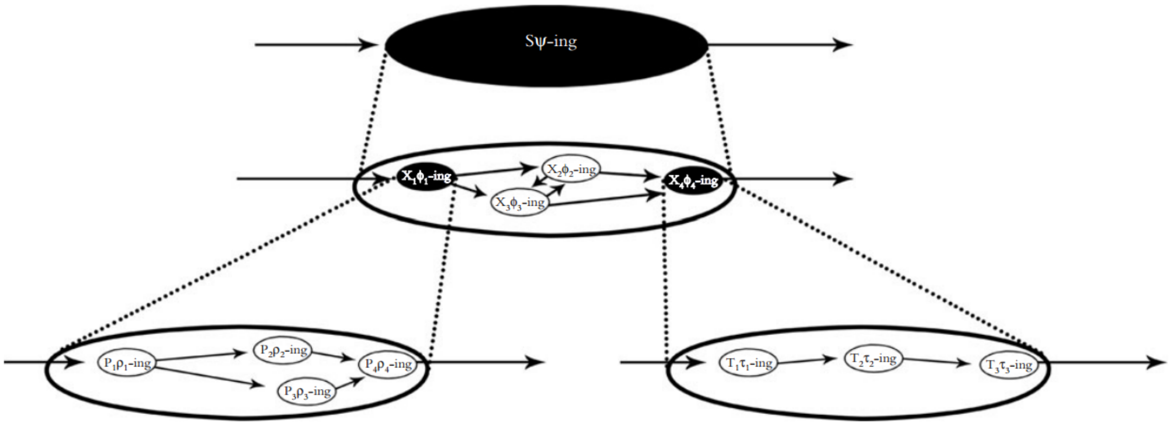

Fig. 14 Levels of organisation in a mechanism. S ψ-ing is a phenomenon and X1..X4 constitute the parts of the mechanism, whose interaction (φ-ing) and organisation (directed edges) is responsible for S's ψ-ing. Both X1 and X4 can themselves be considered as a whole and decomposed to discover what parts are collectively responsible for their φ-ing. This iterated process results in a multi-level mechanism with three distinct levels of organisation. Figure from Craver (2007, p. 194).

5.2.2 Levels of Mechanisms

In Fig. 13 we can already see two levels in a mechanism but for getting the full details of the mechanistic account of levels it will be useful to expand the model into a three-tiered mechanism as shown in Fig. 14. The basic details regarding what a mechanism is and the process of investigating mechanisms provides most of the information we need to understand the mechanistic account of levels but there are a few important additions that arise once we get a mechanism structure like Fig. 14. Craver sets up the account by offering answers for the several questions. Firstly, precisely what is being organised into levels? (the relata question) Second, in virtue of what are two items at different levels? (the relations question), and third, in virtue of what are two items at the same level? (the placement question). For the mechanistic account of levels the answers are as follows:

> *The Relata Question*: some activity or property of a mechanism as a whole, and the activities, properties, or organizational features of its components (its relevant parts and organization)
>
> *The Relation Question*: X's ψ-ing, is at a lower mechanistic level than S's φ-ing if and only if X's ψ -ing is a component in S's φ-ing, that is, if and only if X's ψ-ing is a relevant spatiotemporal part of S's φ-ing
>
> *The Placement Question*: X's ψ-ing and S's φ-ing are at the same level of mechanisms only if X's ψ-ing and S's φ-ing are components in the same mechanism, X's ψ-ing is not a component in S's φ-ing, and S's φ-ing is not a component in X's ψ-ing. (Craver, 2015, pp. 17–19).

These answers look a little more complicated because they are set within the mechanistic framework. Put more straightforwardly, the relata in question are parts and wholes of mechanisms.[3] Something is a part of a whole mechanism if it's behaviour is relevant to the behaviour of the whole mechanism. Finally, two parts are on the same level when (a) they are in the same mechanism, and (b) they are not parts of each other.

So, mechanistic levels are built (discovered, if you like) by reiterating the process that builds the first two-tiered partition in Fig. 13 between the system behaviour of interest and the organised, working components that are collectively responsible for that system behaviour. As you can see on Fig. 14, each working component can *itself* be investigated further in this very same way and in doing so another level is developed below.

The final aspect of the account I want to draw attention to is the envisaged relationship between the mechanistic account and others. Craver describes his view as a descriptive pluralist about applications of 'the levels metaphor'. In order to illustrate what he

3 For the sake of simplicity I'll adopt the following convention for discussing wholes and parts in mechanisms. Where I am referring to a mechanism, or a component as a higher-level process, I'll denote the component and its function, e.g., S φ-ing. When I'm referring to a component as a part or lower-level process I'll just denote the components e.g., X1 or X1...4 (X1, X2, X3, & X4). Under the mechanistic account a component only qualifies *as such* if it has a related activity e.g., X's ψ -ing, but this can be understood without having to write 'φ-ing', 'ψ -ing', and so on each time.

means by this, Craver produces a taxonomy of possible applications of the levels metaphor and isolates levels of mechanisms from other possible uses (Fig. 15).

The most crucial move happens right at the start of the taxonomy where it initially splits into levels of science and levels of nature. Following the 'levels of science' branch, by 'units of science' Craver means scientific disciplines or fields. By 'products of science', Craver means, "epistemic constructs, such as analyses, descriptions, explanatory models, and theories" (p. 171). Products of science are what I have been referring to as 'epistemic outputs'. On the right-hand side we have different levels of nature. I'll return to this issue in the next section in more detail. For now I want to make two points clear about Craver's thinking regarding levels. The first is that whilst Craver is a pluralistic about 'the levels metaphor' he has a more limited pluralism regarding levels of *organisation*. The only part of the taxonomy that concerns levels of organisation for Craver is the lower right-hand side under the branch levels of composition. This identification is further refined when Craver develops a continuum between aggregation and mechanisms: "This organizational spectrum from aggregate to mechanism covers all the relations that go into levels of organization, the superordinate class" (Craver, 2015, p. 16).

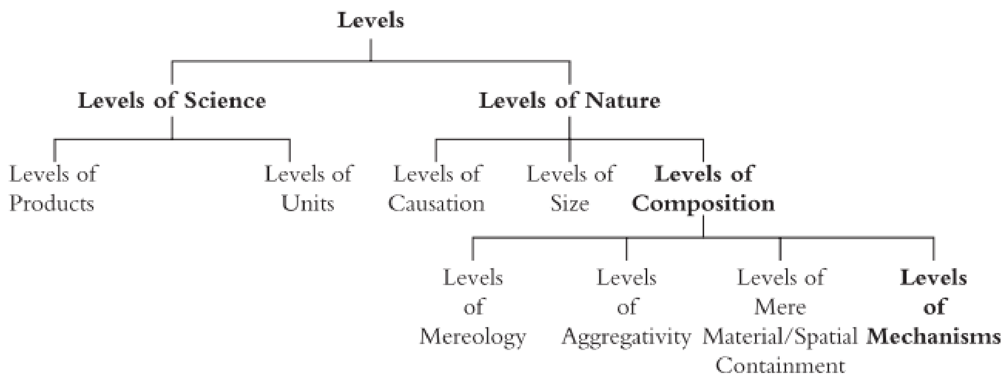

Fig. 15 Craver's Taxonomy of Levels Concepts. The primary condition for forming branches on the taxonomy is what the levels concept is being applied to: the relata. In this first instance, this affords a spilt between, on the one hand, the relata of science: models, explanations, etc., (products) and disciplines, or fields (units). On the other, objects in the world (nature). As the branches develop down the right-hand side, applications of levels concepts are differentiated depending on how they order their relata into higher or lower levels, e.g., by causal interaction, size, or part-whole composition. Figure from Craver (2007, p. 171).

Craver's spectrum of *organisational* relations begins at aggregation and *ends* at mechanisms. By taxonomising these concepts, Craver clearly rules out the other kinds of (what I label) vertical relations from figuring as applications of levels of organisation. For example, mereology is separated from aggregate relations. Additionally, relations of control fall under 'levels of causation' so they are dispensed with at a level above mechanisms in the taxonomy, as are 'levels of size'. Generally speaking, taxonomies serve the purpose of isolating one individual as neatly as possible from other, related individuals and Craver's taxonomy is no different: levels of organisation are distinct from all the other levels concepts on the taxonomy. There is a little wiggle room in the concept in the shape of a range of relations from aggregation to mechanisms. However, as you can probably anticipate given the work done in the previous chapters, I'll be arguing below that this pluralism does not go far enough.

The second, which is tightly related to the first, is that levels of organisation generally, and levels of mechanisms specifically, are about levels of *nature*; part-whole decompositions of stuff in the world. When discussing levels of science, for example, Craver often contrasts these concepts with *ontological levels*, for example: "the take-home lesson: the application of the levels metaphor to fields of science yields a notion of levels only indirectly related to *ontological levels* (as understood in a roughly compositional, part-whole sense)" (2015, p. 7, my emphasis). Povich & Craver (2018, p. 189) even suggest that, "Oppenheim and Putnam could easily have embraced mechanistic levels as the ontic component of their picture; things at higher levels are wholes made up of things at lower levels."

5.2.3 Local Scope: An Issue for the Mechanistic Account

One issue that has been raised against the mechanistic account is that it is *too local* and *restrictive* to account for important aspects of scientific practice. I'll take a look at this issue from two perspectives: (a) in terms of what qualifies as a *component* in a mechanism, and (b) how mechanism boundaries are drawn. My claim in the following section (5.3) will be that when re-interpreted under the pluralistic characterisation of levels I offer, these two issues are, if not ameliorated altogether, then at least made much more palatable for the mechanistic account of levels.

Craver is wary of avoiding what he describes as ‘monolithic’ levels (2015, p. 18); that is, the *global scope* of the layer cake model. Craver is right to be cautious about scope issues. Recall that back in Chapter One (1.5), I argued that it is the combination of step-wise decomposition, aggregation, and global scope that results in the scientific implausibility of the LCM; these elements form the essential tension at the heart of the LCM. I also noted there that step-wise decomposition is a key aspect of vertical system partition. This is reflected in the mechanistic account, both aggregation and global scope are rejected in order to alleviate this tension, while step-wise decomposition remains. As we saw above, there’s a little wiggle room left for aggregation (Fig. 15) but global scope is fully rejected. The extent of this rejection is evidenced in the issue of sameness of levels, i.e., the response to the ‘placement question’.

Although Craver offers an answer to the placement question, the answer is only supposed to indicate that, “sameness of level has no significance within this [mechanistic] application of the metaphor” (2015, p. 16). The way to read the answer to the placement question is not as a condition for ‘sameness of level’ claims, but simply a *negative* claim about what isn’t important for mechanistic analysis of levels. Two individuals, *X* and *Y*, are at different mechanistic levels if, for example, *X* is a component in *Y*. If a further individual, *Z* is *also* a component in *Y*, then it is at the same level as *X*. But if *Z* is *not* a component in *Y* and/or it’s not a component in *X*, then there is simply nothing to be said about *Z’s* relationship to *X;* the two individuals do not have a structural relationship, regardless of what other roles *Z* plays with respect to *X* or *Y*.

The thrust of the objection is that the mechanistic account is correct to be wary of global scope, but it has gone too far the other way, making the resulting analysis too limited in scope to capture important ways the levels concepts are used in scientific practice. The first such objection comes from Eronen (2013, 2015). Eronen offers objections to both a ‘weak’ and ‘strong’ interpretation of the what it is to be a mechanistic component.

Here I'll focus only on the strong interpretation.[4] The strong interpretation identifies a certain subclass of component that is relevant for making levels claims: direct components. A direct component is a component of the mechanism *not* also a (sub)component of any other component in the mechanism (Eronen, 2015, p. 49). In Fig. 14, for example, X1...4 are direct components of S's φ-ing. P1...4 are direct components of X1's ψ-ing, but P1...4 are *not* direct components of S's φ-ing. This reading gets around strange technicality of the weak reading (see footnote 4), and, Eronen (2015, p. 50) suggests, it accurately captures the sense in which Craver intended the placement question to be answered. Under the strong reading, the sameness of level question is pretty much irrelevant: build each level of the mechanism by gathering together direct components. Reiterating the process will develop stratifications in the system one step at a time and if you want to know whether two individuals are at the same level, just look and see whether they are direct components for a specific mechanisms; that's all there is to it.

Eronen argues that as well as capturing the *locality* of levels claims that Craver is after, the strong reading also reveals the limitations of the mechanistic analysis. This is because direct components of two different mechanisms can *never* be on the same level of organisation, by definition. However, when we look at scientific practice, we do not find adherence to this claim, quite the opposite. I'll illustrate this claim with the example provided by Kaiser, so let's take a quick look at that objection first.

4 The weak version is more of a technical issue and turns on the role of 'transitivity' in the sameness of level relation. If transitivity holds across sameness of level relations then Eronen (2015 p. 49-50) is able to generate conflicting statements about the placement of individuals in a three-tier mechanism like Fig. 14. We know that both X1 and X4 are at the same level of organisation because they are components in S's φ-ing. We also know that T1 is at a lower level than X4 because it is a component in X4 ψ-ing. However, T1 can also be a component in S's φ-ing under the mechanistic account as long as it meets the same relevance condition as X4, (mutual manipulability or INUS) which it might well do. But T1 it cannot be at a *lower level* than X1 because it is not a component in X1 and so T1 is in fact at the same level as X1. But, if X1 and T1 are at the same level, and X1 & X4 are at the same level, then by transitivity X4 and T1 should be at the same level. This line of reasoning leads straight to an incoherent outcome: T1 is both at the same level as X4 and at a lower level that X4. Craver (2015, p. 19) also considers such a case (attributing its origins to Lindley Darden) but argues for it as a clear example of the *failure* of transitivity which results only from trying to erroneously apply a uniform 'sameness of level' condition across different mechanism.

The second problem, raised by Marie Kaiser (2015, p. 184) has the same result as the first, but approaches the issue from a different perspective. Rather than looking at what qualify as components in a mechanism, it focuses on how mechanisms are carved from the environment in the first place. According to the general mechanistic framework; a mechanism is carved from its environment according to one *specific* behaviour or function. As we can see on Figs. 13 and 14 the mechanism is *defined* as S's φ-ing, not merely S, nor S functioning in any other way. Accordingly, in order qualify as a component in the mechanism X1 (for example) must play a role in S's φ-ing specifically. As Kaiser points out, this restriction places a limit on levels claims within the mechanistic framework: X can only be at a *lower level* than S if it is a component in S's φ-ing. However, according to Kaiser, scientists do not restrict themselves to building levels structures on the basis of decomposing only one specific function of a system. Rather, their levels claims have a wider scope in virtue of identifying target systems on the basis of *multiple* system-level behaviours.

Kaiser uses the example of the Paramecium to illustrate this claim and it can serve to illustrate Eronen's point about direct components too. A Paramecium is a unicellular (single-celled) organisms, covered in cilia – little hairy structures, which pulse and enable the Paramecium to move and to sweep food in the organism's 'mouth' (technically called its oral groove). The micronucleus is a cell-organelle that divides during asexual reproduction (along with a division mechanism of the macronucleus).

Contracting vacuoles are another type of cell organelle, these can expel water from the cell in order to regulate fluid absorbed by osmosis (they are found mainly in water, so this function is important for the organism). The claim for both Eronen and Kaiser is that under the mechanistic account we can't consider the three components – cilia, micronucleus, vacuole – to be on the same level of organisation, relative to each other and relative to the Paramecium as a whole. They can qualify as *direct components* of processes to which they are relevant – cilia: movement; micronucleus: asexual reproduction; and vacuole: nutrition absorption. However, *qua* direct component of different mechanisms they have no structural relationship to one another. They are as the heart is to the car engine; not comparable in terms of levels.

From a boundary carving perspective it is not clear how the 'Paramecium' can be considered as a whole mechanism anyway given that it does not have just one behaviour that can be used to both carve it from the environment *and* subsume all relevant working components under.

In sum, Eronen and Kaiser's objections highlight gaps in the mechanistic account that emerge as a result of over-restrictive conditions placed on levels claims by the broader mechanistic framework. Specifically, these constraints are the inability to make same-level claims regarding direct components of different mechanisms, and restricting the carving of a mechanism from its environment to *one* system-level behaviour.

Whilst there have been no direct responses from the main proponents of the mechanistic account, a few solutions have been offered that are at least friendly to the mechanistic viewpoint (Bertolaso & Buzzoni, 2017; Kästner, 2018).[5] The gist of these solutions is that both Eronen and Kaiser have missed the *perspective-shifting* nature of mechanistic levels and instead hold the levels *fixed* when considering different 'mechanism-component' relationships in different contexts. For example, Eronen is partially correct to point out that direct components of different mechanisms can *never* be at the same level of organisation, but the 'never' is true only within a given context. What were once direct components of *different* mechanisms in one context, can become direct components of the *same* mechanism in another context. All that matters is that the components meet a relevance condition for some system-level behaviour (mechanism) and they can become direct components, irrespective of their standing in *other* mechanisms in *different* contexts.

If this short dialectic shows anything its that an account of levels needs an explicit way to address issues of *scope*. Craver's path to local scope is to *refuse* answer 'the placement question' in a substantial way, ultimately arguing that claims about 'sameness of level' just don't have a useful application within the framework.

5 I should perhaps differentiate between the two for the sake of accuracy. Bertolaso & Buzzoni directly offer a solution to Eronen's worry and defend the objectivity of 'mechanisms', 'levels', 'components' with a *context-dependence* framework of explanation according to which different explanatory levels correspond to different pragmatic interests and practical possibilities (2017, p. 166). Kästner, on the other hand, agrees that Eronen's objection is a problem for the mechanistic account and proposes to replace the concept of 'mechanistic level' with different *perspectives* on a system (2018, pp. 6-7). Whilst they disagree as to the strength of Eronen's objection, Bertolaso & Buzzoni's notion of a context-dependent explanatory level and Kästner's 'perspectives' on a system, provide a roughly similar way forward for thinking about inter-level relations in mechanistic framework.

This move leaves a lacunae that can be filled with misinterpretation and ambiguity as to the scope of levels claims. Of course, this is an issue that has been central throughout this dissertation and it was something I took care to explicitly address within my own characterisation. Broadly speaking, I agree with Kästner and Bertolaso & Buzzoni that these objections are not damning for the mechanistic view, at least not if we're willing to make some tweaks to the account. To that end, I want to offer what I contend is a neater solution for the mechanistic view: to harness the tools offered by my characterisation to avoid this ambiguity around scope. As we'll see presently, that will involve understanding the mechanistic view as proving an analysis of a certain class of system conceptualisations (Fig. 16). This reinterpretation will involve some minor adjustments to the view but achieved in such a way that the core features of the mechanistic analysis are retained. I'll return to this is issue 5.3.2, for now let's get on with that re-interpretative task.

5.3 Reinterpreting the Mechanistic Account

I will now show how the mechanistic analysis fits into the pluralistic characterisation of levels I developed in the previous chapter. Under my characterisation of levels of organisation, the mechanistic account can be interpreted as picking out an analysis of levels for a *certain kind* of system conceptualisation and/or a certain type of integrative epistemic activity. To see this let's recall the characterisation in full:

> Levels of organisation are collections of vertical and horizontal organisational principles deployed as representative tools for (a) conceptualising a target system of inquiry and (b) the manipulation, integration, and transfer of information, of and between systems of inquiry.

Focusing first on (a), as noted above, the substantial part of the mechanistic account of levels is provided by the mechanistic framework more generally. Building levels in a mechanistic system is mostly an exercise in reiterating the process by which one mechanism is decomposed into its working components.

As it happens, a mechanistic analysis of a system is pretty much an exemplar *component system*. When plotted, such system conceptualisations would occupy the space around of the middle of the continuum of research strategies, but covering more space towards the left than the right. In other words, starting in the area at which the component dominant spike begins and ranging towards, but falling short of, the threshold of interaction dominance (Fig. 16). That is, they are component dominant systems built by near decompositions. Mechanistic systems are far from *aggregative* due to the heavy emphasis placed on *interacting* components *organised* in a specific way. But they are not so interaction-focused that the heuristic of decomposition begins to outlive its utility. Each component in a mechanism retains a degree of independence *relative speaking*, when compared to systems conceptualised such that system-wide interactions (e.g., feedback, distributed control, see 3.3.2) become the primary focus on the system conceptualisation and interaction dominance prevails.

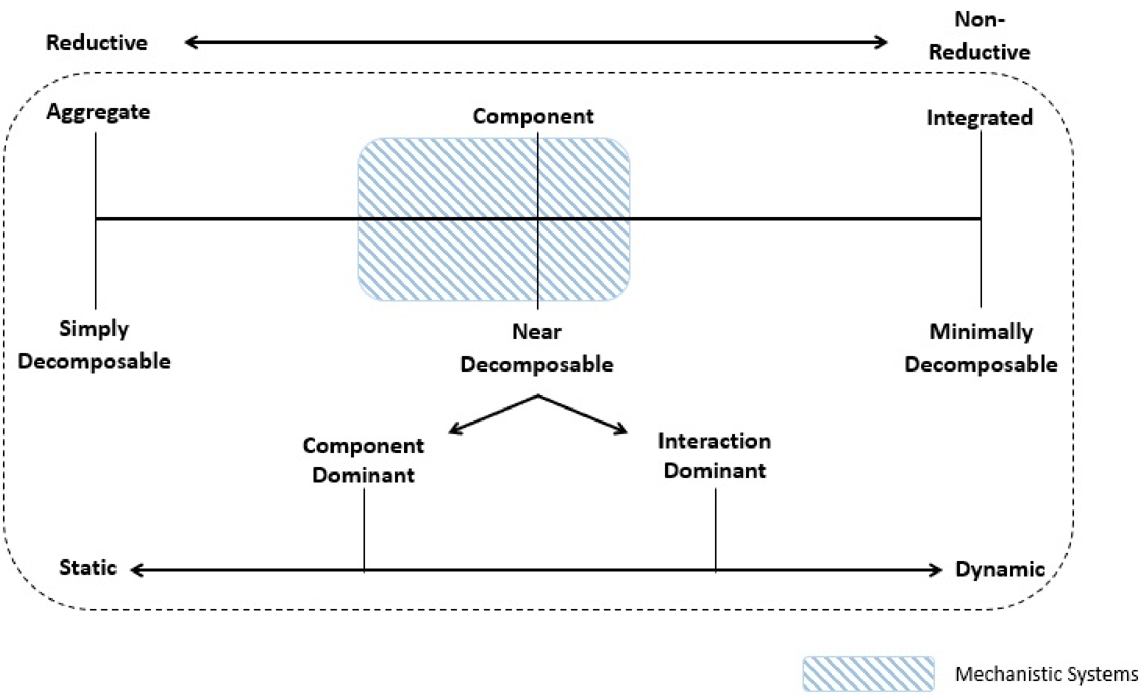

Fig. 16 Situating mechanisms on the continuum of research strategies. The shaded area captures the system conceptualisations that are covered by the mechanistic analysis. These range from the spike of 'component dominance' to just short of crossing the threshold of 'interaction-dominance'.

If mechanistic systems can be characterised this way, then we can understand the levels concept as deploying the vertical organisational principle of component dominant near decomposition to create partitions in the system. Indeed, Craver makes several references to Simon’s notion of near decomposition when explicating mechanistic levels. The association is perhaps at it strongest in the following passage:

> “Indeed, the assumption of near decomposability underlies the strategy of reverse engineering, of discovering how something works by learning how its parts interact. Kauffman (1970) calls this practice “articulation of parts explanation”; Haugeland (1998) calls it “explanation by system decomposition”; Cummins (1975) calls it “functional analysis”; Fodor (1968), Craver (2007), and others (Machamer, 2004; Glennan, 2002; Menzies, 2012) call it “mechanistic explanation’” (Povich & Craver, 2018, p. 186).

In the above quote Povich and Craver *identify* mechanistic explanation with the strategy of reverse engineering, which itself is ‘underlined’ by the idea of near decomposition. Accordingly, I think it is accurate to place mechanisms as a system analysis within the shaded area on Fig. 16 and ascribe the vertical organisational principle of component dominant near decomposition to the mechanistic levels.

That takes care of (a) conceptualising a target system of inquiry. As for (b) the manipulation, integration, and transfer of information, of and between systems of inquiry, I propose that the case study examined in the previous chapter of the multi-level model of glucose homeostasis (4.3.2) serves as a nice illustration of how a mechanistic analysis might operate in this sense. To recap briefly, the intermediate level of the glucose transport mechanism (Fig. 11B) was developed in order to bridge the gap between established mechanisms at a higher scale: the digestive system response to increased glucose concentration (Fig. 11A) and at a lower scale: the insulin action mechanism (Fig. 11C), that were each investigating using different investigative techniques. Leaving aside the issue of different methodological techniques for a moment, the intermediate level served as a bridge precisely by playing the role of a working component of the process in the higher level and playing the role of a mechanism that *itself* was decomposed into working components at the lower level. The claim that each of these levels constitutes *processes* also nicely fits the mechanistic picture. It's not merely that the insulin action mechanism contains a collection of parts that are present in the intra-cellular mechanism of glucose transport, but it shows a process of specifically organised working components that are responsible for the signalling processes that facilitate the uptake of glucose into the cell (by triggering the movement of glucose receptors to the membrane of the cell).

5.3.1 Problems of Interpretation: Levels of Nature

There are two standout obstacles to the reinterpretation of the mechanistic analysis I'm proposing. The first concerns one of the main moves I've made in this dissertation: to switch the focus of analysis from epistemic outputs to epistemic practices. Recall that epistemic outputs are units like theories, laws, explanations involving arguments or the logical structure that holds between propositions and so on. To focus on epistemic practices is, quite simply, to take as primary units of analysis the epistemic activities that *result* in those outputs – however one prefers to codify them (see 3.2). I've elected to develop that idea in terms of a focus on *research strategies* understood as a collection of four key aspects. It is within this context that the characterisation of levels I have developed belongs. Levels of organisation are tools for representing systems via the iterative application of partitioning heuristics, be they inter-level (vertical) or intra-level (horizontal) and be they for the purposes of building system structure (system-specific), or for the integration and/or transfer of information between models (meta-systemic).

As discussed above, the mechanistic account of levels seems to focus neither on epistemic outputs nor epistemic practices but is instead ontological in character, or so it would seem. In fact, at least as far as Craver's exposition goes, the status of levels claims within the mechanistic view invites some serious confusion.

Consider the following passage:

> "This application of the levels metaphor, according to which levels of organization are understood in terms of levels of aggregation and levels of mechanisms, thus offers a no-nonsense ontological picture that comports well with the kinds of explanatory structure one finds in neuroscience and throughout the special sciences generally" (Craver, 2015, p. 23).

This passage gives – at best – very mixed messages about the status of levels. Firstly, the account is the application of a 'metaphor'. So levels are metaphorical. But *the application* of this metaphor seems to give us a 'no-nonsense *ontological* picture'. How we get from using metaphors to gaining an ontology is not clear at all. Finally, the ontology produced by applying this levels metaphor also fits well with the explanatory structure of the life sciences. Levels are nothing but a metaphor to be applied to the world, yet they also yield an ontological picture that accords with a pre-existing explanatory *structure*. To some extent, this could just be an unfortunate choice of words (as a device, metaphors are not usually conducive to 'no-nonsense' pictures, quite the opposite). At the very least I think we can unpack the assumptions that invite this confusion and in doing so, perhaps alleviate it.

The first of these assumptions is the 'relata-first' approach to thinking about different levels concepts. When building his taxonomy of levels concepts – the basis of Craver's descriptive pluralism about applications of the levels metaphor – the first sorting condition is the relata question: exactly *w*hat is being organised into levels. This provides the first big split in the taxonomy between levels of *science* and levels of *nature*. Craver is not alone in making relata an important aspect of differentiating different levels concepts. Daniel Brooks (2016, p. 107) argues that the referent of any levels claim is, at least partly, determined by the entities being structured.[6]

6 For Brooks, the other aspect that determines the referent of the concept is the scope of the entities that the level concept captures on a scale from global to local. In addition to a referent, levels concepts have a 'meaning' which is provided by 'the definitional criteria of content' (which I interpret as something like inter-level relations) and 'the mode of presentation'. Brooks casts levels as a 'fragmented concept' which differs in its deployments based on variations in the referent and meaning of the claims on any given occasion (2016, pp. 107–109).

Accordingly, Craver's taxonomy begins with a division between the two most general kinds of things that levels concepts could be about things in the world – levels of nature – and epistemic items that pertain to those things in the world – levels of science. This initial split already leads to a curious result. Recall that Craver defines levels of products (of science) as, "epistemic constructs, such as analyses, descriptions, explanatory models, and theories" (2007, p. 171). This branch on the very far left of the taxonomy. Levels of mechanisms are literally as far away as they possibly could be on the taxonomy from levels of products. The conceptual space between levels of products and levels of mechanism, gives the impression that when we're talking about levels of mechanisms we really aren't talking about models, modelling practices, or even explanations. Instead we're picking out these genuine features of nature. In levels terms, that implies that we're picking out bona fide, robust ontological levels *in the world*.

But why the explicit commitment to such a strong understanding of levels and where does it come from? As we've seen, Craver is wary of scope precisely because of the difficulties faced by Oppenheim & Putnam's 'monolithic' LCM, so why is he not equally as wary of adopting the ontological perspective on levels that at least one interpretation of the LCM affords (Fig. 1)? The answer to this question brings me to the second problematic assumption at play: the application of the distinction between 'epistemic' and 'ontic' explanations to an analysis of levels of organisation.

Wesley Salmon (1984) made the distinction between epistemic, modal, and 'ontic' explanations. Not worrying about modal explanations for present purposes, epistemic explanations were classified by Salmon as being based on inference or argument, exemplified by the deductive-nomological account (see 2.3). Since the D-N account is out of favour in contemporary philosophy of science, the epistemic view has become associated with a broader class of 'explanatory texts': descriptions, models, diagrams and so on. Ontic explanations, which was Salmon's own view, attempt to capture patterns of regularities in the world, usually interpreted as causal structure. Cory Wright (2015) explicates the standard interpretation of ontic explanation as the commitment to the idea that explanations are both mind-independent objects (they exist 'out there in the world') and are non-representational.

To explain, on the ontic view, is to demonstrate, or show, how explananda phenomena fit into the pattern of regularities in the world.[7]

The issue of whether mechanistic explanations are ontic or epistemic has become an active area of discussion in the mechanistic literature (Illari, 2013; Wright, 2012). Bechtel favours an epistemic view arguing that:

> “The problem with this ontic view is that mechanisms do not explain themselves. They are operative in the world whether or not there are any scientists engaged in offering explanations. Explanation is an activity of scientists who must contribute mental labor in advancing explanations” (Bechtel, 2008, p. 18)

On the other hand, Craver (2007), Glennan (2005), and Machamer et al., (2000) explicitly argue for an ‘ontic’ view: “Objective explanations are not texts; they are full-bodied things. They are facts, not representations” (Craver, 2007, p. 27)

My suspicion is that the disagreement concerning whether mechanistic explanations are ontic or epistemic has seeped into the discussion concerning the status of levels in the mechanistic account. In other words, Craver has this distinction in mind when describing levels of mechanisms *as ontic*. What makes levels of mechanisms robust is that they latch on to real partitions in real systems. These makes sense within the internal logical of the mechanistic account. Mechanistic explanations are ontic because they pick out actively organised components that operate in the causal structure of the world: mechanisms. Mechanisms and the explanations that pick them out are essentially multi-level (Craver, 2007, p. 9-15). So it stands to reason that *levels* within those structures must also be out there in the world. However, I think this terminology, coupled with the relata-first approach, invites a lot of confusion when thinking about the status and role of levels of organisation concepts in scientific practice.

7 Wright’s own argument in this paper is that unpacking the notion of ‘demonstrating’ or ‘exhibiting’ that is required for the ontic view ends up deflating the ontic account into something resembling the epistemic view. At the very least, Wright argues, a fully worked out ontic view has an unavoidably epistemic component.

Under my framework levels concepts are neither ontic nor epistemic. The distinction just doesn't have a significant application to thinking about the role of levels of organisation concepts in scientific practice. Both of these labels belong to assessments of epistemic outputs (theories, explanations etc.) not epistemic activities. Of course, there may well be close connections between the two. The heavy emphasis placed on the link between explanation and scientific practice by the epistemic view sits comfortably with the practice-oriented picture I have been operating within. However, conceptually at least, it seems perfectly consistent to take either view of explanation on my account of levels of organisation. One could adhere to levels of organisation as epistemic tools for system building and still maintain an *ontic* conception of explanations formed on the basis of the resulting model – the explanation can still be mind-independent and non-representational regardless of the analysis of how the model was developed.[8]

The ontological commitments we ought to have regarding certain system conceptualisations, and the different epistemic functions that such models might be used to perform (explaining, controlling, predicting etc.,) are massively complex issues. I have suggested already that the ontological status of the *levels* in any such system is completely tied to the status of the system itself and its role in scientific inquiry. These are two related, but independent conversations. Levels of organisation as epistemic tools belong to the process of developing, integrating, and refining representations of systems as part of a broader research strategy. As such, I suggest that it makes little sense to describe *levels concepts themselves* as ontic or epistemic, in the sense of Salmon's distinction.

Taking these issues together results in three modifications required to integrate the mechanistic account into a broader pluralistic framework. The first is to approach levels concepts not from a relata-first perspective but to differentiate levels concepts primarily based on their *structural features* – the purpose-relative vertical and horizontal principles.

8 Kaiser's view on the distinction is interesting (2015, pp. 242-243). She argues for an *ontic* conception of explanation but is explicitly clear that this does not yield any claims about ontological reduction. Partly this is a result of her 'weak reading' of 'ontic' according to which explanations are epistemic items (contra the mind-independent view) but their explanatory power is derived from the correspondence to actual relations in the world. The relationship between the weak reading of ontic explanations and her own account of levels of organisation is not made explicit, but if we assume that they are connected, then the weaker reading of 'ontic' suggests that Kaiser has a more epistemic understanding of levels than Craver.

After all, if levels concepts are about anything then surely they are about *structure*. Uncovering the differences that result in system structure due to the application of a different set of system partitions (levels), and how these reflect (or are a result of) different purposes of inquiry, seems to be the most fruitful aspect of thinking about the way in which levels of organisation concepts are utilised in scientific practice. This aspect of analysis doesn't turn on the relata of the levels concepts. Consider for example, the protein kinases in the MAPK cascade (Fig. 12). In the example I discussed in 4.5, there were no qualitative features about the proteins that comprised the cascade at all. They were described only as gateways with quantitative thresholds for signal activation and duration. What the example picked out was the vertical *control* relation used to develop the system's structure. But the very same proteins might be modelled as part of a more mechanistic structure in which they form working components that interact in specific ways to be responsible for cellular behaviour. In fact, this is precisely the sort of role that MAPK cascades play in the 'classic' somatic mutation theory discussed in Chapter 3. In both cases we're dealing with the same relata – protein kinases – but analysing the different levels of organisation concepts in play can help to illicit the structural differences in the system conceptualisation and how these differences are linked to the purpose of the inquiry. If we take a structure-first approach to differentiating levels concepts, the mechanistic view sits nicely within the pluralistic framework I offer by providing an analysis of the structural features of a certain kind of system conceptualisation: component dominant nearly decomposable systems (i.e., mechanisms).

Highlighting the tendency of levels talk to take a relata-first approach also uncovers a second, subtle, modification: a difference in the use of the concept of 'organisation' between my account and the mechanistic view. Again, I do not think they are incompatible. The mechanistic view places a lot of emphasis on the fact that working components are *organised* in a specific way. It is this organisation that strongly differentiates mechanisms from mere aggregates (Craver, 2007, p. 190). In the mechanistic account, *each level* exhibits a specific organisation and so the phrase 'levels of organisation' means nested hierarchies of specifically organised levels. My characterisation has a much broader understanding of 'organisation'.

As I discussed at the outset of Chapter Four (4.2) a principle of organisation pertains to the relative positioning of individuals in a system; whatever kind of heuristic is being applied to create those partitions.

So, in my sense, aggregate systems are *also* structures with 'levels of organisation', just as much as component systems are, whereas in the mechanistic account, it is mechanisms that really exhibit *organised* levels structures. Similarly, control systems are not structured into 'levels of organisation' under the mechanistic account. Once again, this is a consequence of the mechanistic view of levels being deeply embedded within the mechanistic framework more generally. Within the framework, the word 'organisation' has a very specific meaning which it retains when its proponents apply it to discussions of levels of organisation. This is a particularly acute problem if we take a *relata-first* approach because the meaning of terms like 'organisation' are fixed when we focus neither on systems generally, nor classes of individuals, but *mechanisms* and all the conceptual paraphernalia that comes attached. Dropping the emphasis on relata means that the mechanistic understanding of 'organisation' can be understood as a particular type of the broader concept of 'organisation' I employ here. It is still true that the interactions that are found at each level of the mechanism are very different to the sort of interactions (or lack thereof) found in aggregative, or control systems. The difference in organisation that the mechanistic analysis wants to pick out is still highlighted on the continuum of research strategies and so the key aspect of what makes a system 'mechanistic' is not lost under my framework.

The third modification is recognise that the ontic-epistemic distinction contributes to a misleading distinction in philosophical discussions of levels of organisation between levels of nature (ontology) and levels of science (epistemic). By separating out the analysis of epistemic products and the processes by which they are produced, we can give a characterisation of levels of organisation that is applicable across a wide range of different systems and in different kinds of inquiry. We can also leave the door open to different normative evaluations of the products (i.e., explanations) that are produced, at least partly, by the use of a levels of organisation concept. Finally, there's no sense in which the concept of levels is 'metaphorical' on my picture. Fortunately, under my account, we don't need to grapple with what it means for a concept to be metaphorical and how a concept of that type relates to scientific inquiry. Instead we swap metaphors for methodology and understand the role that levels of organisation play in representing system structure.

5.3.2 Addressing Scope

Part of the aim of this chapter is to show that the subsumption of the mechanistic account of levels under my pluralistic framework would not only help to clarify the distinctive features of my framework but could also benefit the mechanistic account itself. That is, the cost of taking on board the points in the previous subsection – dropping a relata-first approach, being wary of the ontic-epistemic distinction, and the joint effect of operating under a broader conception of 'organisation' – is outweighed by the benefit of alleviating a major criticism raised against the mechanistic account. To test this claim let's return to the objection according to which the mechanistic account of levels is *too local* in scope and so cannot account for some important uses of levels claims in scientific practice. So, how does the re-interpretation of the mechanistic view under the broader characterisation of levels I have developed help alleviate that objection?

As a first step, it's fair to say that scope is an evergreen issue when thinking about levels concepts, it has certainly been a recurrent theme in this dissertation (see 1.4.1; 2.4.2; 3.5). One of the take-home lessons of the first two chapters was that a scientifically plausible account of levels must at least provide an explicit position on the scope of inter-level relations. I suggested that the appropriate position ought to consist in having the resources to explain why a commitment to global scope is so unconvincing from the perspective of current research in the life sciences *without* needing to invoke any *in principle* arguments against global scope. I argued that the continuum of research strategies has the resources to maintain this particular balance. To recap, this was because plotting strategies on the continuum is an explicitly local affair, not least because there are a plurality of ways to apply heuristics to a system once the boundary for the system has been fixed (recall the examples of the SMT & TOFT) and of course there are multiple ways to fix that boundary in the first place. Nevertheless, it remains possible that systems could be integrated on a global scope. It would, of course, require the integration of all possible system conceptualisations into one overarching model, which in turn would require the fixing of a *global* system boundary and the consistent application of a vertical organisational principle throughout the entire system. Luckily, that seems extremely unlikely from the viewpoint of contemporary scientific inquiry.

This desideratum, I suggest, can also be found in Craver's wider framework. It is evident in his concept of a 'mosaic unity of neuroscience'. Contrasting the mosaic with classical models of reduction, Craver envisages unity – in neuroscience at least – being developed through the piecemeal integration of mechanisms, primarily by identifying constraints that operate on mechanisms at different levels (Craver, 2007, chapter 7). The precise details of the picture are less important here than the overall message, which is precisely that the mechanistic framework provides the conceptual resources for developing mechanisms at a broader scope. However, the details of, and prospects for, any such unifying integration will only be known empirically as inquiries develop.

Given this overlap in desideratum regarding scope, perhaps this aspect of a reinterpretation of the mechanistic account within my framework is not much of a *re*interpretation after all. Instead, the reluctance to commit to 'sameness of levels' claims is, in the end, simply an expression of a view about scope that is shared between the mechanistic account and my own: from the perspective of current research in the life sciences, global scope would be a highly problematic and unjustifiable feature of a levels of organisation concept. At the same time, the scope at which levels claims can be applied is entirely dependent on models actually being developed at that scope.

If the foregoing is correct, then although there's no difference in *aims* regarding scope, there still remains a difference in application as highlighted by the discussion of objections raised by Kaiser and Eronen in 5.2.3 above. Accordingly, we need to see how the situating of the mechanistic view within my own helps in realising this aim. The main resource that my account offers that is salient to this issue is horizontal conditions. Recall that these are:

i. Intra-level partitions
ii. Specification of intra-level interactions
iii. Boundary carving

None of these horizontal principles are direct responses to Craver's 'placement question'. They will not dictate when individuals end up on the same level of organisation and so invoking them does not constitute a straightforward rejection of Craver's reluctance to give the 'sameness of level' issue a robust application. Intra-level partitions differentiate individuals *already* at the same level, for reasons linked to the purpose of the inquiry (recall the intra-level partitions in the HMMT, Fig. 9, 4.3.1). Specifying intra-level interactions picks out differences in methodological techniques used to investigate individuals at that level (recall the differences between techniques used in the multi-level model of glucose homeostasis, Figs. 11A, 11B, 11C., 4.3.2). Finally, boundary carving merely divides individuals in the system from those in the environment.

There's two possible solutions on the table that stem from horizontal considerations. The first is a change of boundary carving condition: to drop the mechanistic requirement for carving a system from its in environment based on only *one* system-level function. Presumably the one-function one-mechanism heuristic is supposed to play a role in keeping the levels structure extremely local. But if we can already ensure that levels claims remain system-specific then there's no need to adopt such a strict horizontal condition. Although there are *statements* to be found expressing a commitment to one-function one-mechanism boundary carving condition (Craver, 2007, p. 123; Darden, 2008, p. 960; Glennan, 2002, p. 344) they are not usually accompanied by detailed argument as to why it would be inconsistent to characterise one mechanism relative to several functions.[9]

Part of Kaiser's solution to the objection raised against the mechanistic account is to define a *biological part* as one that takes place during, and is relevant to, one of the *characteristic* behaviours of the whole (Kaiser, 2015, p. 181). As far as I can tell, there is no conceptual reason why the mechanistic account couldn't adopt at least *this* aspect of a broader notion of 'parthood' and remain perfectly consistent with the rest of the mechanistic framework.

9 For example, a much discussed issue in the mechanistic debate regarding identifying mechanisms (i.e., boundary carving) concerns the *regularity* with which a mechanism produces a phenomenon; how regular must it be and can the mechanistic analysis account for seemingly irregular or even one-off events? (Andersen, 2012, 2018; DesAutels, 2011; Glennan, 2010). But this debate doesn't seem to affect the claim that one mechanism could be productive (regularly or otherwise) of *several* behaviours and be demarcated from its environment relative to those multiple behaviours.

Notice further how this could also alleviate the issue of direct components as pointed out by Eronen. If the mechanism itself is carved using a more permissive boundary condition, then more components will qualify as *direct components* of the same mechanism and thus be considered on the same level of organisation.

More saliently for my purposes than whether a more permissive boundary carving condition directly clashes with the mechanistic account is that there is definitely no reason why a system in the shaded region of Fig. 16 couldn't be carved from its environment relative to several functions. Nothing in my account of levels dictates that this must be so. A broader horizontal condition would interfere neither with plotting the system on the continuum nor analysing the purpose-relative vertical and horizontal conditions with which it was conceptualised. So, if the foregoing interpretation of the mechanistic account as picking out strategies for investigating system conceptualisations within the shaded area of Fig. 16 is correct, then the insights gained regarding those strategies can be retained whilst applying a more permissive boundary carving condition.

For the sake of argument let's suppose that there is a deeper reason for the one-function, one-mechanism horizontal condition that goes beyond consistence within the mechanistic framework. Still, within my characterisation of levels there's another solution available, namely part (b) of the characterisation: the manipulation, integration, and transfer of information, of and between systems of inquiry. We can think of each of the three Paramecium behaviours (or system-level functions) as individual mechanisms; replete with all technicalities of the mechanistic framework. But, we can *integrate* the three mechanisms by utilising part (b) of my characterisation of levels as it pertains to horizontal model transfer (see 4.4). That is, we can treat the development of the Paramecium conceptualised as a single mechanism as a result of meta-systemic integration of information gained about the lower-level processes of different mechanisms. The horizontal integration of these separate mechanisms is what brings the three distinct process onto one level of organisation. This might seem like a convoluted way of conceptualising the single-celled organism but in fact its just an acute instance of what is presumably a wide-spread horizontal transfer. For example, like the Paramecium, cells that line the human trachea (windpipe) also have motile cilia.

Information gained about the transport mechanism of involving cilia on the Paramecium might be transferable to the function of sweeping debris away from the lungs and even *cell migration* thought to be facilitated by cilia that line tracheal cells (Enuka, et al., 2012).

In sum, even though the horizontal integration of three mechanisms into one single-celled organism might seem peculiar, it is just a very specific instance of meta-systemic horizontal model transfer and demonstrates how the machinery of my account of levels can provide a solution to the problems Kaiser and Eronen raise.[10]

How does this solution compare to the ones offered by Bertolaso & Buzzoni (2017), and Kästner (2018). Well, the thrust is certainly similar: the 'levels' in a system are fluid and context-dependent, in my case they are dependent on system conceptualisation which can happen in a plurality of ways via the deployment of different vertical and horizontal organisational principles But I think there's two additional features that my solution offers – that is, two distinctive features of my characterisation that address these issues of scope in a beneficial way for the mechanistic account. The first is to address the ambiguity surrounding the *ontological* status of mechanistic levels. As noted, I agree with Bertolaso & Buzzoni, and Kästner that context is important to levels claims but I think that what drives the sort of objections raised by Eronen and Kaiser is that mechanisms – and the levels therein – are repeatedly described as *ontic* or even *ontological.* This places a conceptual roadblock to the fluidity and context-dependency of levels claims. How can the 'mechanism-component' relationship be subject to many different perspectives when it is supposed to be directly latching on to patterns in the causal structure of the world, for example? The strategy of *re-describing* or *re-conceptualising* systems (mechanisms) and their structure (components & subcomponents) is less obviously amenable to this 'ontic' view of mechanistic levels.

10 I should also note that in addition to the introduction of more permissive boundary condition, Kaiser (2015, p. 183) includes the condition of membership of a shared 'general biological kind', as another way in which two individuals can be placed on the same level of organisation. There's no objection to kind membership that follows from my account of levels but neither is the explicit use of kinds required for the point I want to make here. Accordingly, for the claims I want to make in this chapter, I can stay away from the difficult issue of 'kinds' in the life sciences, without denying that the issues are of course related. Investigating the relationship between 'kinds' and 'levels' within the practice-oriented framework I operate under here may be a fruitful avenue for future inquiry.

However, when subsumed under my framework, we remove this roadblock and embrace local pluralism about levels claims *qua* those claims being about features of system *representations*, not causal features of the world.

Secondly, the distinctive group of ‘horizontal organisational principles’ provided by my characterisation give concrete ways of dealing with the sort of integrative, or perspective-shifting, solutions suggested by Bertolaso & Buzzoni, and Kästner. Of course, this is not an ‘advantage’ over their frameworks per se. Bertolaso & Buzzoni draw on the context-dependence of causal attributions to make their case, and Kästner sketches a new way to think about levels altogether – as perspectives rather levels. I do not claim that my solution is better, rather I’m just drawing attention to the fact that my characterisation of levels of organisation offers a built-in solution to these kinds of worries that can be applied to concrete cases which, in line with the aims of this chapter, highlights another distinctive feature of my account of levels more generally.

5.3.3 Going back to the Start: The LCM, Mechanisms, & Levels of Organisation

As just mentioned, the main of this chapter has been to highlight the distinctive features of my account and the contribution my account can make to the philosophical thinking about levels of organisation in scientific practice. In the foregoing discussion these features were intertwined with the discussion of how to reinterpret the mechanistic picture of levels and so I’ll end this chapter by making those features more explicit. I think the best way to tie together the threads of the work undertaken here and in doing so, make those features more explicit, is to zoom out a little and consider the relationship between the LCM, the mechanistic account, and my own characterisation of levels. We can make a start by bring back Box 2 and Fig. 1, developed all the way back in Chapter One.

(1) Global Scope: the system under consideration is the entire world
(2) Monism about the Levels Concept: there will only be one correct model of levels of organisation
(3) System Type: It is an aggregative system that is simply decomposable
(4) Distinct Levels: no entity 'properly' belongs to more than one level and all its parts are found at the next lowest level
.

Box 2 A summary of the key features of the LCM.

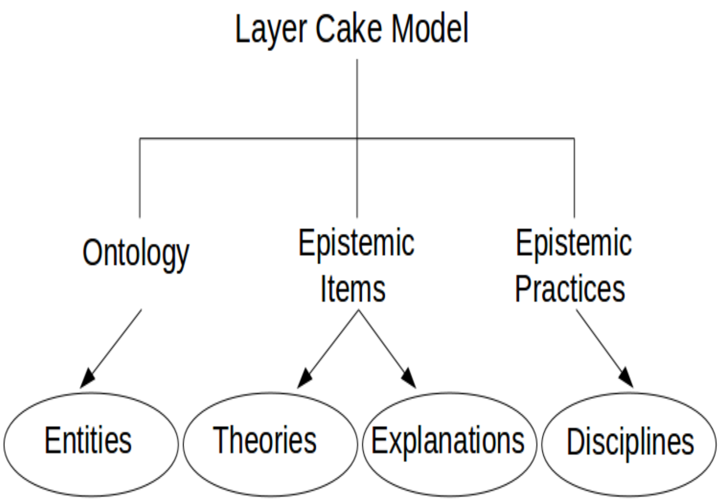

Fig. 1 Three variants of the LCM. Each variant is an interpretation of the LCM, differentiated according to what is being structured into levels in accordance with the six conditions of the LCM.

On the face of it, the mechanistic account seems to reject all of these features of the LCM. It adheres to only a very local application of any given levels concept (contra 1). It is pluralistic in the sense that ‘the levels metaphor’ has multiple applications (contra 2). It certainly rejects the conceptualisation of systems as *only* aggregative (contra 3). Whilst it doesn’t directly reject the distinctiveness of levels (4), it attempts to side-step the issue altogether in virtue of arguing against the application of ‘sameness of level’ claims within the framework. Finally, at least insofar as the mechanistic account is tied to the practice of decomposition strategies (e.g., Bechtel & Richardson, 2010), it will adhere to step-wise decomposition, the second interpretation of aspect 4.

This might seem like a positive result for the mechanistic account, after all the LCM is deeply out of step with contemporary scientific practice owing to the essential tension at its heart – concurrently adhering to stepwise decomposition of aggregate systems at a global scope (see 1.5). However, if anything can be taken away from the work done in Chapters One and Two it’s hopefully that developing a scientifically plausible account of levels is a far more complicated affair than simply rejecting the key features of the LCM.

Let’s start at the top of the list with aspect (1): global scope. Kaiser’s and Eronen’s objections can be understood as arguing that rejecting global scope is a more nuanced matter than the mechanistic account can account for, at least in its initial formulation. In the introduction (5.1) I described the mechanistic account as taking a ‘heavy-handed’ approach to scope: either levels concepts are as global as the LCM or they must be formalised in such a way as to be extremely local. Reinterpreting the mechanistic account under my framework further highlighted how my picture of levels offers a more nuanced approach to scope, one I argued would in fact be amenable to the aims of the mechanistic account.

Given the legacy of the LCM it is perfectly understandable that those who still want to utilise the resources of ‘levels of organisation’ as a tool of analysis, move as far away as possible from the global scope contained within the LCM; a particularly uncomfortable feature of the LCM for philosophical positions that aim to accurately capture aspects of contemporary scientific inquiry. However, in this chapter we’ve seen how adopting the directly opposing position – advocating restrictive views on ‘sameness of levels’ to secure a very local scope of application – comes with its own difficulties as far as applicability to contemporary scientific practice goes. This draws attention to a distinctive feature of my account of levels: it negotiates the balance between explaining why the application of a levels concept at a global scope seems to unconvincing, without relying on in principle arguments that restrict the kinds of models that may be developed in the future.

There are no *conceptual* reasons that fall out of my characterisation of levels that rule out the very possibility of a global scope of application. Rather, the characterisation creates a framework into which empirical content must be plugged. If a conceptualisation can be developed that manages to gather every object into one all-encompassing system, complete with a consistent vertical organisational principle, then the characterisation of levels I offer would have been applied at a global scope. However, and importantly, such a system conceptualisation seems extremely unlikely to be developed from the perspective on contemporary scientific inquiry, not least because it’s not clear at all how working towards such a model would benefit the kinds of research projects I’ve looked at in this dissertation; and presumably many more besides.

The second issue I wanted to address was the seemingly ontological character of the mechanistic account. This, I argued, was a consequence of taking a relata-first approach and the application of the epistemic-ontic distinction to the discussions of levels of organisation. This brings me back to another recurring theme that runs throughout this dissertation: the explicit focus on epistemic practices over both epistemic outputs and ontological interpretations of levels claims. Of course, simply claiming that I’m focusing on epistemic practices doesn’t make issues of metaphysics and ontology disappear. But this is not been my claim. Rather, I’ve been suggesting a switch in where philosophers should *begin* thinking about levels of organisation. Drawing explicit attention to this issue is crucial for discussions of levels owing to another legacy feature of the LCM; the interweaving of the three interpretations of ‘levels of organisation’ such that ontology, epistemology, and scientific inquiry are merged into the same conversation (Fig. 1). This significantly weakens the prospects for a scientifically plausible account of levels by entangling the concept with far more suspect commitments. Instead I’ve been making the case for starting with the way levels are utilised in scientific practice and ensuring that there really is a scientifically plausible characterisation available.

After that a further, very substantial question, that remains concerns the metaphysical implications of, and perhaps assumptions within, the systems that have been developed. In return for this temporary suspension of further metaphysical and ontological issues, I've been able to explore the prospects of a more 'deflated' sense of levels and have tried to show that the notion is in fact a rich conceptual tool for understanding scientific practice.

The final issue pertains to both aspects (2) and (3) of the LCM: namely its pluralism. It's true that the mechanistic account leaves room for a limited pluralism about levels of organisation: a continuum between aggregrative systems and mechanisms. So it clearly goes beyond aspect (3). However, as evidenced by Craver's taxonomy of levels (Fig. 15) many more vertical organisational principles are ruled out as belonging to the concept. For example, those that sit towards the right-hand side of the continuum of research strategies, i.e., interaction dominant near decomposition, and control. This, I argued, stemmed from the embeddedness of the mechanistic account of levels in the broader framework of a mechanistic analysis of explanations. Specifically, it means that the use of 'organisation' has a much narrower meaning than is found in my account. In my view, 'organisation' is a much broader concept and pertains to the relative positioning of individuals in a system, *whatever* vertical partitioning heuristic is applied to that system in order to achieve that arrangement of individuals.

I don't think there exists a philosophical consensus on what 'organisation' means so I take it that this broader interpretation is fine so long as I shoulder the burden of proof for showing what work my interpretation can do. I hope to have evidenced this in the previous two chapters by demonstrating the utility of the more inclusive concept of levels of organisation that it affords overall. If that has been convincing then so long as levels of organisation are epistemic tools for system building, then there's no *prima facie* reason to exclude many vertical organisational principles from qualifying under the concept. As long as they play a role in representing system structure, they belong to the toolbox of levels of organisation concepts.

5.5 Conclusion

The primary goal of this final chapter has been to situate the characterisation of levels I have been developing within a broader philosophical discussion of levels of organisation in scientific practice, and, in doing so, to draw attention to the distinct features of my account and the specific work they allow the characterisation to do. My strategy to achieve this has been to demonstrate how the most prominent account of levels of organisation in the current philosophical literature – levels of mechanisms – can be integrated into my broader framework. This involved a reinterpretation of some aspects of the mechanistic view, specifically those that intersected with the distinctive features of my account that I wanted to draw attention to: scope, pluralism, organisation, and clarity regarding the interpretation of levels-claims as pertaining only to their use in scientific practice. The result is positive for all concerned. I have illustrated the distinctive features of my view and showed how they help clarify and alleviate problems associated with accounts of levels. If they felt like it, proponents of the mechanistic account could harness the reinterpretation I've offered to avoid the criticisms of their account without losing the core of what the mechanistic picture seeks to clarify, namely levels structures in component dominant systems. Along the way, I hope to have drawn attention to other key contributors to the levels discussions, namely Kaiser, Brooks, and Eronen, respectively, chiefly by way of discussing their points of divergence with the mechanistic view.

General Conclusion

In this general conclusion I'll retrace my steps through the dissertation and bring out the main conclusions of the work. I'll end with a short discussion of what I think are the future directions of research opened up by the work undertaken here. My goal in this dissertation has been to develop a scientifically plausible account of levels of organisation. The work towards that goal began by getting to the root of a juxtaposition that seems to have emerged regarding philosophical discussions of levels. The Layer Cake Model (LCM) had become the most prevalent and influential account of levels of organisation in philosophy, whilst being a scientifically implausible view of levels of organisation. This juxtaposition has put the utility of the concept in jeopardy, as evidenced by the recent arguments of 'levels sceptics' (Eronen, 2015; Potochnik & McGill, 2012; Thalos, 2013).

The first question that I addressed was simply: what is the layer cake model? In Chapter One I answered this question through a careful explication of Oppenheim & Putnam's elaboration on the LCM. There were two main results. The first was a collection of general features of the LCM (Box 2): global scope of application; a monism about levels of organisation; a system type that is aggregative and subject to simple decomposition; and distinctness of levels (including step-wise decomposition). The second was the introduction of three possible interpretations of the LCM demarcated according to the unit of analysis that was organised by the LCM structure (Fig. 1). These were an ontological interpretation (entities & their properties); an interpretation regarding epistemic outputs (theories, laws, and explanations); and an interpretation regarding epistemic practices (disciplines).

Finally, I argued that the LCM's scientific implausibility emerges from an essential tension at it's core: the combination of several of its general features, namely global scope, step-wise decomposition, and an understanding of the system as fully aggregative and subject to simple decomposition. I argued that there's nothing pernicious about these features taken separately, but put in combination they explained why scientific counterexamples to the LCM are easy to generate.

The work in Chapter One took care of one half of the juxtaposition: the scientific implausibility of the LCM, but it's pervasiveness in philosophical discussions of levels remained to be addressed.

In Chapter Two I argued that the key to understanding the LCM's impact on philosophy lies in its deep conceptual and historical connections to discussions of reduction and emergence. I argued that via a further overlap with the unity of science project, the LCM had become intertwined with Nagelian reduction. Neither the classical unity of science project nor Nagelian reduction are popular views in contemporary philosophy of science. A reasonable hypothesis might suggest that the LCM suffered the same fate in virtue of belonging to this particular family of concepts. On the contrary, I argued that the multiple interpretations of the LCM allowed it to out manoeuvre critiques of Nagelian reduction and maintain its influence on philosophical discussions of levels.

To demonstrate this claim I argued that two directly contrasting positions on reduction remain fully compatible with the LCM: ontological emergence and epistemological anti-reductionism. I sought to develop key lessons from their compatibility that illuminated what an account of reduction and emergence must make explicit in order to shed any compatibility with the LCM. These lessons were as follows. Firstly, that the discussion opened up the way to move beyond a dichotomy between reduction and emergence. Secondly, that the subsequent account should provide a clear understanding of the role that aggregation plays in inter-level relations and a clear sense of how and when aggregation fails. I argued that the over-reliance on a 'property-first' approach had become an obstacle in addressing this issue for traditional versions of reduction and emergence and proposed to adopt a 'system-first' approach instead. Thirdly, that a pluralism about levels might be best demonstrated by interrogating the distinctness of levels in the LCM. Finally, that the issue of scope of application would have to be explicitly accounted for in the new account.

In Chapter Three I put these lessons to work in developing a new account of inter-level relations. The first move I made was to switch my focus from ontology and epistemic outputs to solely focus on epistemic practices – that is, reduction as an activity rather than as a relation between objects in the world or epistemic outputs of inquiry, such as theories or explanations. This allowed me to develop an account of research strategies according to which no strategy is reductive or non-reductive *simpliciter*, but only more, or less, reductive compared to another and within a frame of reference. That frame of reference was to be a continuum of research strategies that conceptualised their target system of a sliding scales of (a) susceptibility of the system to decomposition, and (b) prevalence of key system dynamics.

I tested this new framework by applying it to analyse recent developments in cancer research. This chapter was largely concerned with laying conceptual and methodological groundwork for the account of levels to come. Particularly the introduction of the practice-oriented framework within which the account is developed; the range of inter-level relations captured by the continuum itself; the interplay between strategies of decomposition and dynamics; and the methodology of seeking to illustrate my claims through detailed case studies in the life sciences.

Another important aspect of the continuum is that it allowed me to put the LCM behind me in what I hope was a conclusive way. This is because the framework of inter-level relations it captures has the resources to address those four key lessons from Chapter Three and in doing so to make a clean break from LCM-style thinking about levels, reduction, and emergence.

Chapter Four saw the introduction of my characterisation of levels, which is as follows:

> Levels of organisation are collections of vertical and horizontal organisational principles deployed as representational tools for (a) conceptualising a target system of inquiry and (b) the manipulation, integration, and transfer of information, of and between systems of inquiry.

The two parts of the characterisation were designed to explicitly address the dimensions of evaluation spelled out in the General Introduction. These were that the characterisation be *scientifically plausible* – understood as being able to capture a wide-range of deployments of the concept in scientific practice. Secondly, that the characterisation would have *analytic utility*, meaning that it can illuminate certain aspects of scientific practice and ultimately contribute to a better understanding of the complex processes of knowledge production found in scientific inquiry. I introduced the key aspects of the characterisation and then proceeded to demonstrate the two parts – (a) and (b) – in three examples drawn from three different points on the continuum of research strategies.

The general idea behind the account is that the concept of levels of organisation captures the process by which systems are *partitioned* in various ways and for various reasons.

These partitions can be vertical in the sense that they produce stratification in the system and differentiate individuals from one another by placing them relatively higher and lower than one another. The continuum of research strategies captures a range of key vertical partitioning relations. Partitions can also be horizontal and I introduced three senses in which this might apply to representing system structure: intra-level partitions, specification of intra-level interactions, and boundary carving.

In the final chapter my aim was to situate my characterisation in relation to contemporary philosophical discussions of levels. I argued that a much discussed contemporary account of levels – the mechanistic account of levels of organisation – could be reinterpreted and situated within my characterisation. This process, I argued, is of benefit to all those concerned. I was able demonstrate the strong pluralism and generality of my account, as well as drawing attention to its distinctive features. Additionally, I was able to offer a solution to a major objection to the mechanistic account, regarding its own scientific plausibility. The basis of the reinterpretation was that the mechanistic account of levels captures a specific class of system representations on the continuum – roughly component dominant and subject to near decomposability.

There were a couple of sticking points that had to be negotiated in order to give a smooth reinterpretation. These issues stemmed from the fact that the mechanistic account of levels is deeply embedded in the mechanistic approach to philosophy of science. I argued that the key insights of the account could remain without some of these more conceptually loaded features of the view. I identified these issues as follows: from the application of the ontic-epistemic distinction in philosophy of explanation to issues of levels of organisation; generating a pluralism about levels through a relata-first approach, rather than through a pluralism about structure; and a more esoteric understanding of the concept of ‘organisation’. I concluded by taking a step back and considering the LCM, the mechanistic account, and my own in relation to the general features of Box 2 in order to make clear the specific position my characterisation provides.

I began this dissertation with two contrasting quotes regarding levels of organisation. The first, from Wimsatt, claimed that they are a deep non-arbitrary features of the world. The second, from Eronen, claimed that there’s little point in further developing an account of levels of organisation from the perspective of contemporary scientific inquiry.

This latter position Eronen labels as ‘deflationary’ as it dispenses with ‘levels’ in favour of relations of composition and scale. So, how does my characterisation sit in relation to these views? Well, it can be described as deflationary in relation to Wimsatt’s view and certainly in relation to the LCM. It makes no claims about the structure of the world or even about the structuring of classical philosophical units of analysis for science – theories, laws, and explanations. Given the role we’ve seen the LCM play in the unity of science project and in discussions concerning, for example, the metaphysical status of causal interaction between properties, this is a pretty deflationary position. I should stress once again, however, that nowhere have I argued *against* levels playing these sorts of role, nor would I argue that my own characterisation of levels *cannot* contribute to such debates – quite the opposite, I would hope. The ‘deflationism’ in my account is a matter of priority and focus. I have argued that if we’re interested in the scientific plausibility and analytic utility of levels of organisation, we should start by looking at the role they play in the conceptualisation of target systems of inquiry in scientific practice.

Indeed, I’ve argued that when we do start from this perspective, we can be far more optimistic about levels that Eronen seems to be. I have sought to demonstrate the use of levels concepts in a heterogeneous group of examples: cutting-edge research into the human Microbiome, the integration of well-established clinical knowledge and in vitro modelling techniques concerning a self-regulating system (glucose homeostasis), and highly abstract mathematically modelling of protein signalling pathways. If I am correct in my analysis of these projects then it looks as though there’s plenty of life in the levels concept yet. If any of the foregoing work in this project has been illuminating or perhaps even convincing, then the concept of levels itself can do a lot of work in analysis the conceptualisation of target systems of inquiry and the development of those inquires over time.

Of course, none of this is to suggest that the case file on levels of organisation can now be closed. Instead, I hope to have staked a claim for embracing its utility in the analysis of scientific practice and to continue to develop the features of the account in relation to other aspects of scientific practice. To this end, I’ll conclude by suggesting several lines of inquiry that flow from that acceptance of the characterisation of levels I’ve offered here. Firstly, the focus of my project here was narrowed to building a characterisation of levels from the ground up. A complementary project would be to examine the features of this characterisation in relation to other established features of system structure.

Specifically I am thinking of features such as *modularity*, *robustness*, and *plasticity*. Modularity seems to have an important role to play in component dominant systems as well as considerations of manipulating specific component functions. The issue of robustness must have a close relation to the sorts of system dynamics consider in Chapter Three given that robustness can be an a feature of a system that has a lot of redundancy with respect to its components. Similarly, plasticity concerns the dynamic development of a system in response to both endogenous and exogenous changes and this be an important feature of system conceptualisations from across the range of the continuum. Bringing these features of system into the discussion was beyond the scope of my purpose here but going forward it seems like a natural progression to consider how the decomposition and dynamics of system conceptualisation relate to these issues.

A second broad theme of development from the end point of this dissertation is to consider the relationship between levels of organisation and the other aspects of a research strategy that contribute towards the conceptualisation of a target system of inquiry. For example, I position the horizontal and vertical organisational principles that constitute a levels concept as 'theoretical assumptions', with the exception of the horizontal principle of specification of intra-level interactions, which connected with methodological assumptions. It must surely be the case that more considerations that fall under my label of 'methodological assumptions' will have a dynamic relationship with these purpose-relative representational tools. What resources are available to researchers, the skills and interests of a team's membership, and both the benefits and limitations of adopting particular experimental techniques to investigate the system.

Another aspect of research strategies that would be of great benefit to investigate further is the process by which boundaries are developed for a system conceptualisation in different contexts. In this dissertation, boundary conditions played a role in fixing the level of abstraction for comparing strategies on the continuum (Chapter Three) as well as comprising a horizontal organisational principle (Chapter Four). I was careful to note that a more developed account of boundary carving would have to be context sensitive and very well informed from an empirical point of view. Developing such an account and relating it to the work done here on system conceptualisation would constitute a substantial and fruitful project.

My hope is that these future lines of research will be made more approachable and have a well-developed base on which to develop thanks to the work completed in this dissertation on the concept of levels of organisation.

Bibliography

Alexander, S. (1920). *Space, Time, and Deity. 2 vols.* London: Macmillan.

Andersen, H. (2012). The case for regularity in mechanistic causal explanation. *Synthese*, *189*(3), 415–432.

Andersen, H. (2017). Reduction in the Biomedical Sciences. In M. Solomon, J. Simon, & H. Kincaid (Eds.), *Routledge Companion to Philosophy of Medicine*. New York: Routledge.

Andersen, H. (2018). What would Hume say? Regularities, laws, and mechanisms. In S. Glennan & P. Illari (Eds.), *The Routledge Handbook of Mechanisms and Mechanical Philosophy* (pp. 157–168). New York: Routledge.

Ayala, F. J. (1974). Introduction. In F. J. Ayala & T. Dobzhansky (Eds.), *Studies in the philosophy of biology: reduction and related problems* (pp. vii–xvi). Berkeley: University of California Press.

Baker, L. R. (2003). Metaphysics and Mental Causation. In J. Heil & A. R. Mele (Eds.), *Mental Causation* (pp. 75–96). Oxford: Oxford University Press.

Baker, S. G. (2009). Improving the Biomarker Pipeline to Develop and Evaluate Cancer Screening Tests. *Journal of the National Cancer Institute*.

Baker, S. G. (2015). A Cancer Theory Kerfuffle Can Lead to New Lines of Research. *Journal of the National Cancer Institute*, *107*(2), 1–8.

Barnes, B., Bloor, D., & Henry, J. (1996). *Scientific knowledge: A sociological analysis.* Chicago: Chicago University Press.

Barnes, E. (2012). Emergence and Fundamentality. *Mind*, *121*(484), 873–901.

Baumgartner, M. (2009). Interventionist Causal Exclusion and Non reductive Physicalism. *International Studies in the Philosophy of Science*, *23*(2), 161–178.

Baumgartner, M. (2013). Rendering Interventionism and Non-Reductive Physicalism Compatible. *Dialectica*, *67*(1), 1–27.

Baxendale, M. (2018). Mapping the continuum of research strategies. *Synthese*, *In Press.*

Bechtel, W. (2008). *Mental mechanisms: philosophical perspectives on cognitive neuroscience*. New York: Routledge.

Bechtel, W. (2010). The Downs and Ups of Mechanistic Research: Circadian Rhythm Research as an Exemplar. *Erkenntnis*, *73*(3), 313–328.

Bechtel, W. (2017). Using the hierarchy of biological ontologies to identify mechanisms in flat networks. *Biology & Philosophy*, *32*(5), 627–649.

Bechtel, W., & Abrahamsen, A. (2005). Explanation: a mechanist alternative. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*(2), 421–441.

Bechtel, W., & Richardson, R. C. (2010). *Discovering complexity: decomposition and localization as strategies in scientific research*. Cambridge, Mass.: MIT Press.

Bedau, M. (2008). Downward Causation and Autonomy in Weak Emergence. In M. Bedau & P. Humphreys (Eds.), *Emergence: contemporary readings in philosophy and science* (pp. 155–188). Cambridge, Mass.: MIT Press.

Bedau, M., & Humphreys, P. (Eds.). (2008). *Emergence: contemporary readings in philosophy and science*. Cambridge, Mass.: MIT Press.

Bedessem, B., & Ruphy, S. (2015). SMT or TOFT? How the Two Main Theories of Carcinogenesis are Made (Artificially) Incompatible. *Acta Biotheoretica*, *63*(3), 257–267.

Bennett, K. (2007). Mental Causation. *Philosophy Compass*, *2*(2), 316–337.

Bertolaso, M. (2009). Towards an Integrated View of the Neoplastic Phenomena in Cancer Research. *History and Philosophy of the Life Sciences*, *31*(1), 79–97.

Bertolaso, M. (2011). Hierarchies and causal relationships in interpretative models of the neoplastic process. *History and Philosophy of the Life Sciences, 33*(4), 515–535.

Bertolaso, M., & Buzzoni, M. (2017). Causality and Levels of Explanation in Biology. In M. Paolini Paoletti & F. Orilia (Eds.), *Philosophical and scientific perspectives on downward causation* (pp. 164–180). New York: Routledge.

Bickle, J. (2006). Reducing mind to molecular pathways: explicating the reductionism implicit in current cellular and molecular neuroscience. *Synthese*, *151*(3), 411–434.

Bissell, M. J., & Hines, W. C. (2011). Why don't we get more cancer? A proposed role of the microenvironment in restraining cancer progression. *Nature Medicine*, *17*(3), 320–329.

Bokulich, A. (2003). Horizontal Models: From Bakers to Cats. *Philosophy of Science*, *70*(3), 609–627.

Bontly, T. D. (2002). The Supervenience Argument Generalizes. *Philosophical Studies*, *109*(1), 75–96.

Brigandt, I. (2013). Systems biology and the integration of mechanistic explanation and mathematical explanation. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *44*(4, Part A), 477–492.

Brigandt, I., & Love, A. (2017). Reductionism in Biology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017). Metaphysics Research Lab, Stanford University. Retrieved from https://plato.stanford.edu/archives/spr2017/entries/reduction-biology/

Broad, C. D. (1925). *The mind and its place in nature.* London: Paul, Trench, Trubner.

Brooks, D. S. (2016). *The Levels of Organisation Concept in Biology*. Bielefield: Universität Bielefeld.

Brooks, D. S. (2017). In Defense of Levels: Layer Cakes and Guilt by Association. *Biological Theory*, *12*(3), 142–156.

Bruno, W., Martinuzzi, C., Andreotti, V., Pastorino, L., Spagnolo, F., Dalmasso, B., … Ghiorzo, P. (2016). Heterogeneity and frequency of BRAF mutations in primary melanoma: Comparison between molecular methods and immunohistochemistry. *Oncotarget*, *8*(5), 8069–8082.

Burge, T. (2003). Mind-Body Causation and Explanatory Practice. In J. Heil & A. R. Mele (Eds.), *Mental Causation* (pp. 97–120). Oxford: Oxford University Press.

Carnap, R. (1934). *The Unity of Science*. London: Kegan Paul, Trench, Trubner, and Co.

Cartwright, N. (2004). Causation: One Word, Many Things. *Philosophy of Science*, *71*(5), 805–819.

Cartwright, N., Cat, J., Fleck, L., & Uebel, T. E. (1996). *Otto Neurath Philosophy Between Science and Politics.* Cambridge, Mass.: Cambridge University Press.

Casati, R., & Varzi, A. C. (1999). *Parts and places: the structures of spatial representation*. Cambridge (MA): MIT Press.

Chang, H. (2012). *Is water H2O?: Evidence, realism and pluralism.* London: Springer.

Chang, H. (2014). Epistemic Activities and Systems of Practice: Units of Analysis in Philosophy of Science After the Practice Turn. In L. Soler, S. Zwart, M. Lynch, & V. Israel-Jost (Eds.), *Science After the Practice Turn in the Philosophy, History, and Social Studies of Science* (pp. 67–79). Taylor & Francis.

Chang, H. (2015). Reductionism and the Relation Between Chemistry and Physics. In T. Arabatzis, J. Renn, & A. Simões (Eds.), *Relocating the history of science: essays in honor of Kostas Gavroglu* (pp. 193–209). Cham: Springer.

Churchland, P. M. (1981). Eliminative Materialism and the Propositional Attitudes. *The Journal of Philosophy*, *78*(2), 67–90.

Churchland, Patrica. (1986). *Neurophilosophy: Towards a Unified Science of the Mind-Brain*. Cambridge, Mass.: MIT Press.

Churchland, Patricia, & Sejnowski, T. J. (1992). *The computational brain*. Cambridge, Mass.: MIT Press.

Clayton, P., & Davies, P. (Eds.). (2006). *The Re-emergence of Emergence: The Emergentist Hypothesis from Science to Religion*. Oxford: Oxford University Press.

Conlan, S., Kong, H. H., & Segre, J. A. (2012). Species-Level Analysis of DNA Sequence Data from the NIH Human Microbiome Project. *PLOS ONE*, *7*(10), e47075.

Couch, M. (2011). Mechanisms and constitutive relevance. *Synthese*, *182*(3), 119–145.

Craver, C. F. (2007). *Explaining the Brain*. Oxford: Oxford University Press.

Craver, C. F. (2015). Levels. *Open MIND*, *8*, 1–26.

Craver, C. F., & Darden, L. (2013). *In search of mechanisms: discoveries across the life sciences*. Chicago: The University of Chicago Press.

Craver, C. F., & Tabery, J. (2017). Mechanisms in Science. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017). Metaphysics Research Lab, Stanford University. Retrieved from https://plato.stanford.edu/archives/spr2017/entries/science-mechanisms/

Cummins, R. (1975). Functional Analysis. *Journal of Philosophy*, *72*, 741–765.

Darden, L. (1991). *Theory Change in Science: Strategies from Mendelian Genetics*. New York: Oxford University Press.

Darden, L. (2006). *Reasoning in Biological Discoveries: Essays on Mechanisms, Interfield Relations, and Anomaly Resolution*. New York: Cambridge University Press.

Darden, L. (2008). Thinking Again about Biological Mechanisms. *Philosophy of Science*, *75*(5), 958–969.

Darden, L., & Maull, N. (1977). Interfield Theories. *Philosophy of Science*, *44*(1), 43–64.

Davidson, D. (1970). Mental Events. In L. Foster & J. W. Swanson (Eds.), *Experience and Theory*. London: Duckworth.

DesAutels, L. (2011). Against Regular and Irregular Characterizations of Mechanisms. *Philosophy of Science*, *78*(5), 914–925.

Dickson, W. M. (1998). *Quantum Chance and Non-Locality*. Cambridge: Cambridge University Press.

DiFrisco, J. (2017). Time Scales and Levels of Organization. *Erkenntnis*, 1–24.

Dobzhansky, T. (1937). *Genetics and the origin of species*. New York: Columbia University Press.

Druker, B. J., Talpaz, M., Resta, D. J., Peng, B., Buchdunger, E., Ford, J. M., … Sawyers, C. L. (2001). Efficacy and safety of a specific inhibitor of the BCR-ABL tyrosine kinase in chronic myeloid leukemia. *The New England Journal of Medicine*, *344*(14), 1031–1037.

Druker, Brian J., Guilhot, F., O'Brien, S. G., Gathmann, I., Kantarjian, H., Gattermann, N., … IRIS Investigators. (2006). Five-year follow-up of patients receiving imatinib for chronic myeloid leukemia. *The New England Journal of Medicine*, *355*(23), 2408–2417.

Dupré, J. (1993). *The Disorder of Things*. Cambridge, Mass.: Harvard University Press.

Enuka, Y., Hanukoglu, I., Edelheit, O., Vaknine, H., & Hanukoglu, A. (2012). Epithelial sodium channels (ENaC) are uniformly distributed on motile cilia in the oviduct and the respiratory airways. *Histochemistry and Cell Biology*, *137*(3), 339–353.

Eronen, M. I. (2013). No Levels, No Problems: Downward Causation in Neuroscience. *Philosophy of Science*, *80*(5), 1042–1052.

Eronen, M. I. (2015). Levels of Organization: A Deflationary Account. *Biology and Philosophy*, *30*, 39–58.

Eronen, M. I., & Brooks, D. S. (2018). Levels of Organization in Biology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2018). Metaphysics Research Lab, Stanford University. Retrieved from https://plato.stanford.edu/archives/spr2018/entries/levels-org-biology/

Favela, L., H. (2015). *Understanding Cognition via Complexity Science*. University of Cincinnati: Electronic PhD Dissertation. Retrieved from https://etd.ohiolink.edu/

Favela, L., H., & Martin, J. (2017). 'Cognition' and Dynamical Cognitive Science. *Minds and Machines*, *27*(2), 331–355.

Feibleman, J. K. (1954). Theory of Integrative Levels. *The British Journal for the Philosophy of Science*, *5*(17), 59–66.

Fettweis, J. M., Serrano, M. G., Girerd, P. H., Jefferson, K. K., & Buck, G. A. (2012). A new era of the vaginal microbiome: advances using next-generation sequencing. *Chemistry & Biodiversity*, *9*(5), 965–976.

Floridi, L. (2011). Modern Trends in Philosophy of Information. In P. Adriaans & J. van Bentham (Eds.), *Philosophy of information* (Vol. 8, pp. 117–135). Amsterdam: Elsevier, NH.

Fodor, J. (1968). *Psychological Explanation*. New York: Random House.

Fodor, J. (1974). Special Sciences (or: The Disunity of Science as a Working Hypothesis). *Synthese*, *48*(2), 97–115.

Frantz, C., Steward, K., & Weaver, C. (2010). The Extracellular Matrix at a Glance. *Journal of Cell Science*, *123*, 4195–4200.

Frey, S., Wolkenhauer, O., & Millat, T. (2010). Quantifying Properties of Cell Signaling Cascades. In P. Iglesias & B. Ingalls (Eds.), *Control Theory and Systems Biology* (pp. 69–84). Cambridge, Mass.: MIT Press.

Gay, L., Baker, A.-M., & Graham, T. A. (2016). Tumour Cell Heterogeneity. *F1000Research*, *5*.

Gebharter, A. (2015). Causal Exclusion and Causal Bayes Nets. *Philosophy and Phenomenological Research*, 1–23.

Gillett, C. (2016). *Reduction and emergence in science and philosophy*. Cambridge: Cambridge University Press.

Glennan, S. (1996). Mechanisms and the Nature of Causation. *Erkenntnis (1975-)*, *44*(1), 49–71.

Glennan, S. (2002). Rethinking Mechanistic Explanation. *Philosophy of Science*, *69*(S3), S342–S353.

Glennan, S. (2005). Modeling mechanisms. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*(2), 443–464.

Glennan, S. (2010). Ephemeral Mechanisms and Historical Explanation. *Erkenntnis (1975-)*, *72*(2), 251–266.

Glennan, S., & Illari, P. (Eds.). (2018). *The Routledge Handbook of Mechanisms and Mechanical Philosophy*. New York: Routledge.

Godfrey-Smith, P. (2008). Reduction in Real Life. In J. Hohwy & J. Kallestrup (Eds.), *Being reduced: new essays on reduction, explanation, and causation* (pp. 52–74). Oxford: Oxford University Press.

Green, S. (2015). Revisiting generality in biology: systems biology and the quest for design principles. *Biology & Philosophy*, *30*(5), 629–652.

Greenman, C., Stephens, P., Smith, R., Dalgliesh, G. L., Hunter, C., Bignell, G., … Stratton, M. R. (2007). Patterns of somatic mutation in human cancer genomes. *Nature*, *446*(7132), 153–158.

Grice, E. A., Kong, H. H., Conlan, S., Deming, C. B., Davis, J., Young, A. C., … Segre, J. A. (2009). Topographical and Temporal Diversity of the Human Skin Microbiome. *Science*, *324*(5931), 1190–1192.

Guttman, B. (1976). Is Levels of Organisation a Useful Biological Concept? *BioScience*, *26*(2), 112–113.

Hacking, I. (1983). *Representing and intervening: Introductory topics in the philosophy of natural science*. Cambridge: Cambridge University Press.

Hacking, I. (1996). The Disunities of Science. In P. Galison & D. Stump (Eds.), *The Disunity of Science: Boundaries, Contexts, and Power* (pp. 37–74). Stanford: Stanford University Press.

Hall, N. (2004). Two Concepts of Causation. In J. D. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 225–276). Cambridge, Mass.: MIT Press.

Hanahan, D., & Weinberg, R. A. (2000). The Hallmarks of Cancer. *Cell*, *100*(1), 57–70.

Hanahan, D., & Weinberg, R. A. (2011). Hallmarks of Cancer: The Next Generation. *Cell*, *144*(5), 646–674.

Hare, R. M. (1952). *The Language of Morals*. Oxford: Oxford University Press.

Haugeland, J. (1998). *Having Thought: Essays in the Metaphysics of Mind*. Cambridge, Mass.: Harvard University Press.

Healey, R. A. (1989). *The Philosophy of Quantum Mechanics: An Interactive Interpretation*. Cambridge: Cambridge University Press.

Hempel, C. G., & Oppenheim, P. (1948). Studies in the Logic of Explanation. *Philosophy of Science*, *15*(2), 135–175.

Hesse, M. (1966). *Models and Analogies in Science*. Notre Dame: Notre Dame University Press.

Hodis, E., Watson, I. R., Kryukov, G. V., Arold, S. T., Imielinski, M., Theurillat, J.-P., … Chin, L. (2012). A landscape of driver mutations in melanoma. *Cell*, *150*(2), 251–263.

Holden, J., G., Van Orden, G., C., & Turvey, M., T. (2009). Dispersion of response times reveals cognitive dynamics. *Psychological Review*, *116*(2), 318.

Hooker, C. A. (1981). Towards a General Theory of Reduction, Part I: Historical and Scientific Setting, Part II: Identity in Reduction, Part III: Cross-categorical Reduction. *Dialogue*, *20*, 38–59, 201–236, 496–529.

Hu, S., Xie, Z., Onishi, A., Yu, X., Jiang, L., Lin, J., … Zhu, H. (2009). Profiling the Human Protein-DNA Interactome Reveals ERK2 as a Transcriptional Repressor of Interferon Signaling. *Cell*, *139*(3), 610–622.

Huang, B., Fettweis, J. M., Brooks, J. P., Jefferson, K. K., & Buck, G. A. (2014). The changing landscape of the vaginal microbiome. *Clinics in Laboratory Medicine*, *34*(4), 747–761.

Hull, D. (1972). Reduction in Genetics: Biology or Philosophy? *Philosophy of Science*, *39*(4), 491–499.

Hull, D. (1974). Informal Aspects of Theory Reduction. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, *1974*, 653–670.

Humphreys, P. (1997). Emergence, Not Supervenience. *Philosophy of Science*, *64*, S337–S345.

Humphreys, P. (2016). *Emergence*. New York: Oxford University Press.

Huneman, P., & Walsh, D. M. (Eds.). (2017). *Challenging the Modern Synthesis: Adaptation, Development, and Inheritance*. Oxford: Oxford University Press.

Iglesias, P., & Ingalls, B. (Eds.). (2010). *Control Theory and Systems Biology*. Cambridge, Mass.: MIT Press.

Illari, P. (2013). Mechanistic Explanation: Integrating the Ontic and Epistemic. *Erkenntnis*, *78*(2), 237–255.

Illari, P., & Williamson, J. (2012). What is a mechanism? Thinking about mechanisms across the sciences. *European Journal for Philosophy of Science*, *2*(1), 119–135.

Issad, T., & Malaterre, C. (2015). Are Dynamic Mechanistic Explanations Still Explanations? In P.-A. Braillard & C. Malaterre (Eds.), *Explanation in biology: an enquiry into the diversity of explanatory patterns in the life sciences* (pp. 265–292). Dordrecht: Springer.

Jackson, F., & Pettit, P. (1992). In defense of explanatory ecumenism. *Economics and Philosophy*, *8*, 1–21.

Janda, J. M., & Abbott, S. L. (2007). 16S rRNA Gene Sequencing for Bacterial Identification in the Diagnostic Laboratory: Pluses, Perils, and Pitfalls. *Journal of Clinical Microbiology*, *45*(9), 2761–2764.

Kahn, C. R. (1985). The Molecular Mechanism of Insulin Action. *Annual Review of Medicine*, *36*(1), 429–451.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty : heuristics and biases*. Cambridge: Cambridge University Press.

Kahneman, D., & Tversky, A. (1972). Subjective Probability: a judgment of representativeness. *Cognitive Psychology*, *3*(3), 430–454.

Kahneman, D., & Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica*, *47*, 263–291.

Kaiser, M. I. (2011). The Limits of Reductionism in the Life Sciences. *History and Philosophy of the Life Sciences*, *33*(4), 453–476.

Kaiser, M. I. (2012). Why it is time to move beyond Nagelian reduction. In D. Dieks, W. J. Gonzalez, S. Hartmann, M. Stöltzner, & M. Weber (Eds.), *Probabilities, Laws, and Structures* (pp. 255–272). Heidelberg: Springer.

Kaiser, M. I. (2015). *Reductive explanation in the biological sciences*. Cham: Springer.

Kaiser, M. I., & Krickel, B. (2017). The Metaphysics of Constitutive Mechanistic Phenomena. *British Journal for the Philosophy of Science*, *68*(3), 745–779.

Kästner, L. (2018). Integrating mechanistic explanations through epistemic perspectives. *Studies in History and Philosophy of Science Part A, In Press*.

Kauffman, S. A. (1970). Articulation of Parts Explanation in Biology and the Rational Search for Them. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 257–272.

Kellert, S. H., Longino, H. E., & Waters, C. K. (2006). Introduction: The Pluralist Stance. In S. H. Kellert, H. E. Longino, & C. K. Waters (Eds.), *Scientific pluralism* (Vol. 19, pp. vii–xxviii). Minneapolis, MN: University of Minnesota Press.

Kemeny, J. G., & Oppenheim, P. (1956). On Reduction. *Philosophical Studies*, *7*(1), 6–19.

Kim, J. (1998). *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge, Mass.: MIT Press.

Kim, J. (2002). The Layered Model: Metaphysical Considerations. *Philosophical Explorations: An International Journal for the Philosophy of Mind and Action*, *5*(1), 2–20.

Kim, J. (2005). *Physicalism, or something near enough*. Princeton, N.J.; Woodstock: Princeton University Press.

Kim, J. (2006). Emergence: Core ideas and issues. *Synthese*, *151*(3), 547–559.

Kincaid, H. (1990). Molecular Biology and the Unity of Science. *Philosophy of Science*, *75*, 575–593.

Kitcher, P. (1984). 1953 and all That. A Tale of Two Sciences. *The Philosophical Review*, *93*(3), 335–373.

Kitcher, P. (1993). Function and Design. *Midwest Studies In Philosophy*, *18*(1), 379–397.

Kostić, D. (2016). The topological realization. *Synthese*, 1–20.

Kuhn, T. S. (1970). *The structure of scientific revolutions.* (2nd ed.). Chicago: Chicago University Press.

Kuorikoski, J., & Ylikoski, P. (2013). How Organization Explains. In V. Karakostas & D. Dieks (Eds.), *EPSA11 Perspectives and Foundational Problems in Philosophy of Science* (pp. 69–80). Springer International Publishing.

Ladyman, J., Lambert, J., & Wiesner, K. (2013). What is a complex system? *European Journal for Philosophy of Science*, *3*(1), 33–67.

Ladyman, J., Ross, D., Collier, J., & Spurrett, D. (2007). *Everything must go: metaphysics naturalized*. Oxford: Oxford University Press.

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge;* (pp. 91–196). Cambridge: Cambridge University Press.

Laland, K. N., & Brown, G. R. (2011). *Sense and Nonsense: Evolutionary Perspectives on Human Behaviour*. Oxford: Oxford University Press.

Latour, B., & Woolgar, S. (1979). *Laboratory life: The construction of scientific facts.* Princeton, N.J.: Princeton University Press.

Laudan, L. (1977). *Progress and its problems: toward a theory of scientific growth*. Berkeley: University of California Press.

Lausen, F. (2014). Reductionism as a Research Directive. *Journal for General Philosophy of Science*, *45*(2), 263–279.

Leonard, H. S., & Goodman, N. (1940). The Calculus of Individuals and its Uses. *Journal of Symbolic Logic*, *5*, 45–55.

Leśniewski, S. (1992). Foundations of the General Theory of Sets. In S. J. Surma, J. Srzednicki, D. I. Barnett, & V. J. Rickey (Eds.), D. I. Barnett (Trans.), *Stanislaw Leśniewski: Collected Works* (Vol. 1 & 2, p. 390). Heidelberg: Springer.

Levy, A., & Bechtel, W. (2013). Abstraction and the Organization of Mechanisms. *Philosophy of Science*, *80*(2), 241–261.

Loewer, B. (2001). From Physics to Physicalism. In C. Gillett & B. Loewer (Eds.), *Physicalism and its Discontents* (pp. 37–56). Cambridge: Cambridge University Press.

Love, A. C. (2008). Explaining Evolutionary Innovations and Novelties: Criteria of Explanatory Adequacy and Epistemological Prerequisites. *Philosophy of Science*, *75*(5), 874–886.

Love, A. C. (2011). Hierarchy, causation and explanation: ubiquity, locality and pluralism. *Interface Focus*, 1–11.

Machamer, P. (2004). Activities and Causation: The Metaphysics and Epistemology of Mechanisms. *International Studies in the Philosophy of Science*, *18*(1), 27–39.

Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 1–25.

Mackie, J. L. (1974). *The Cement of the universe: a study of causation.* Oxford: Clarendon Press.

MacLeod, M., & Nersessian, N. J. (2015). Modeling systems-level dynamics: Understanding without mechanistic explanation in integrative systems biology. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *49*, 1–11.

Maffini, M. V., Soto, A. M., Calabro, J. M., Ucci, A. A., & Sonnenschein, C. (2004). The stroma as a crucial target in rat mammary gland carcinogenesis. *Journal of Cell Science*, *117*(8), 1495–1502.

Mally, A., & Chipman, J. K. (2002). Non-genotoxic carcinogens: early effects on gap junctions, cell proliferation and apoptosis in the rat. *Toxicology*, *180*(3), 233–248.

Marras, A. (2000). Critical Notice on Kim's Mind in a Physical World. *Canadian Journal of Philosophy*, *30*(1), 137–159.

Marusyk, A., Almendro, V., & Polyak, K. (2012). Intra-tumour heterogeneity: a looking glass for cancer? *Nature Reviews Cancer*, *12*, 323–334.

Mayr, E. (1942). *Systematics and the Origin of Species from the Viewpoint of a Zoologist*. New York: Columbia University Press.

McGivern, P. (2015). *Emergence*. New York: Taylor & Francis.

McLaughlin, B. P. (1995). Varieties of Supervenience. In E. E. Savellos & Ü. D. Yalçin (Eds.), *Supervenience: new essays* (pp. 16–59). Cambridge: Cambridge University Press.

McLaughlin, B. P. (2008). The Rise and Fall of British Emergentism. In M. Bedau & P. Humphreys (Eds.), *Emergence: contemporary readings in philosophy and science* (pp. 19–59). Cambridge, Mass.: MIT Press.

Menzies, P. (2012). The Causal Structure of Mechanisms. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *43*, 796–805.

Mill, J. S. (1882). *A System of Logic*. New York: Harper and Brothers.

Mitchell, S. D. (2009). *Unsimple truths science, complexity, and policy*. Chicago: University of Chicago Press.

Moore, G. E. (1922). *Philosophical Studies*. London: K. Paul, Trench, Trubner & Co,.

Morange, M. (2007). The field of cancer research: an indicator of present transformations in biology. *Oncogene*, *26*(55), 7607–7610.

Morgan, C. L. (1923). *Emergent Evolution*. London: Williams and Norgate.

Mumford, S., & Anjum, R. (2011). *Getting causes from powers*. Oxford; New York: Oxford University Press.

Nagel, E. (1979). *The Structure of Science: Problems in the Logic of Explanation* (2nd ed.). Indianapolis: Hackett Publishing.

Needham, J. (1931). *Chemical Embryology*. Cambridge: Cambridge University Press.

Needham, J. (1932). Thoughts on the Problem of Biological Organisation. *Scientia*, *52,* 84–92.

Neurath, O. (1937). Unified Science and Its Encyclopaedia. *Philosophy of Science, 4*(2), 265–277.

Neurath, O., Carnap, R., & Morris, C. (Eds.). (1938). *International Encyclopedia of Unified Science* (Vol. 1). Chicago: University of Chicago Press.

Novikoff, A., B. (1945). The Concept of Integrative Levels and Biology. *Science*, *101*(2618), 209–215.

Nyman, E., Brännmark, C., Palmér, R., Brugård, J., Nyström, F. H., Strålfors, P., & Cedersund, G. (2011). A Hierarchical Whole-body Modeling Approach Elucidates the Link between in Vitro Insulin Signaling and in Vivo Glucose Homeostasis. *Journal of Biological Chemistry*, *286*(29), 26028–26041.

O'Connor, T., & Wong, H. Y. (2005). The Metaphysics of Emergence. *Noûs*, *39*(4), 658–678.

Oh, E.-Y., Christensen, S. M., Ghanta, S., Jeong, J. C., Bucur, O., Glass, B., … Beck, A. H. (2015). Extensive rewiring of epithelial-stromal co-expression networks in breast cancer. *Genome Biology*, *16*, 128.

O'Malley, M. A., Brigandt, I., Love, A. C., Crawford, J. W., Gilbert, J. A., Knight, R., … Rohwer, F. (2014). Multilevel Research Strategies and Biological Systems. *Philosophy of Science*, *81*(5), 811–828.

Oppenheim, P., & Putnam, H. (1958). The Unity of Science as a Working Hypothesis. In H. Feigl, M. Scriven, & G. Maxwell (Eds.), *Minnesota Studies in the Philosophy of Science* (Vol. 2, pp. 3–36). Minneapolis: Minnesota University Press.

Papineau, D. (2001). The Rise of Physicalism. In C. Gillett & B. Loewer (Eds.), *Physicalism and its Discontents* (pp. 3–36). Cambridge: Cambridge University Press.

Pearl, J. (2000). *Causality: models, reasoning, and inference*. Cambridge: Cambridge University Press.

Pigliucci, M., & Müller, G. B. (Eds.). (2010). *Evolution, the Extended Synthesis*. Cambridge, Mass.: MIT Press.

Plotnikov, A., Zehorai, E., Procaccia, S., & Seger, R. (2011). The MAPK cascades: Signaling components, nuclear roles and mechanisms of nuclear translocation. *Biochimica et Biophysica Acta (BBA) - Molecular Cell Research*, *1813*(9), 1619–1633.

Plutynski, A. (2013). Cancer and the goals of integration. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *44*(4), 466–476.

Popper, K. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.

Potochnik, A. (2010). Levels of Explanation Reconceived. *Philosophy of Science*, *77*(1), 59–72.

Potochnik, A. (2017). *Idealization and the Aims of Science*. Chicago: Chicago University Press.

Potochnik, A., & McGill, B. (2012). The Limits of Hierarchical Organization. *Philosophy of Science*, *79*, 120–140.

Povich, M., & Craver, C. F. (2018). Mechanistic levels, reduction, and emergence. In S. Glennan & P. Illari (Eds.), *The Routledge Handbook of Mechanisms and Mechanical Philosophy* (pp. 185–197). New York: Routledge.

Pray, L. A. (2008). Gleevec: the Breakthrough in Cancer Treatment. *Nature Education*, *1*(1), 37.

Primas, H. (1991). Reductionism: Palaver Without Precedent. In E. Agazzi (Ed.), *The Problem of Reductionism in Science* (pp. 161–172). Dordrecht: Kluwer.

Putnam, H. (1975). The Nature of Mental States. In *Mind, Language, and Reality: Philosophical Papers, Volume 2* (pp. 429–440). Cambridge: Cambridge University Press.

Raatikainen, P. (2010). Causation, Exclusion, and the Special Sciences. *Erkenntnis*, *73*(3), 349–363.

Rescher, N. (1955). Axioms for the Part Relation. *Philosophical Studies*, *166*, 8–11.

Rescher, N., & Oppenheim, P. (1955). Logical Analysis of Gestalt Concepts. *The British Journal for the Philosophy of Science*, *6*(22), 89–106.

Reutlinger, A., & Saatsi, J. (Eds.). (2018). *Explanation Beyond Causation: Philosophical Perspectives on Non-Causal Explanations*. Oxford: Oxford University Press.

Rosenberg, A. (1985). *The structure of biological science*. Cambridge: Cambridge University Press.

Rosenberg, A. (2006). *Darwinian Reductionism: Or, How to Stop Worrying and Love Molecular Biology*. Chicago: University of Chicago Press.

Rosenfeld, S. (2013). Are the Somatic Mutation and Tissue Organization Field Theories of Carcinogenesis Incompatible? *Cancer Informatics*, *2013*(12), 221–229.

Ruben, D. H. (1983). Social Wholes and Parts. *Mind*, *92*, 219–238.

Rückert, F., Grützmann, R., & Pilarsky, C. (2012). Feedback within the Inter-Cellular Communication and Tumorigenesis in Carcinomas. *PLOS ONE*, *7*(5), e36719.

Rueger, A., & McGivern, P. (2010). Hierarchies and levels of reality. *Synthese*, *176*(3), 379–397.

Salmon, W. C. (1984). Scientific Explanation: Three Basic Conceptions. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, *1984*, 293–305.

Sarkar, S. (1992). Models of reduction and categories of reductionism. *Synthese*, *91*(3), 167–194.

Schaffner, K. F. (1967). Approaches to Reduction. *Philosophy of Science*, *34*(2), 137–147.

Schaffner, K. F. (1993). *Discovery and Explanation in Biology and Medicine*. Chicago: University of Chicago Press.

Schaffner, K. F. (2006). Reduction: the Cheshire cat problem and a return to roots. *Synthese*, *151*(3), 377–402.

Seeley, T. D. (1995). *The wisdom of the hive the social physiology of honey bee colonies*. Cambridge, Mass.: Harvard University Press.

Seeley, T. D., & Tovey, C. A. (1994). Why search time to find a food-storer bee accurately indicates the relative rates of nectar collecting and nectar processing in honey bee colonies. *Animal Behaviour*, *47*(2), 311–316.

Sibley, F. (1959). Aesthetic Concepts. *Philosophical Review*, *68*, 421–450.

Silberstein, M. (2006). In defence of ontological emergence and mental causation. In P. Clayton & P. Davies (Eds.), *The Re-emergence of Emergence: The Emergentist Hypothesis from Science to Religion* (pp. 203–226). Oxford: Oxford University Press.

Silberstein, M., & McGeever, J. (1999). The Search for Ontological Emergence. *The Philosophical Quarterly*, *49*(195), 201–214.

Simon, H. A. (1996). *The sciences of the artificial*. Cambridge, Mass.: MIT Press.

Simons, P. (1987). *Parts: A Study in Ontology*. Oxford: Oxford University Press.

Sober, E. (1999). The Multiple Realizability Argument against Reduction. *Philosophy of Science*, *66*, 524–564.

Soler, L., Zwart, S., Lynch, M., & Israel-Jost, V. (Eds.). (2014). *Science After the Practice Turn in the Philosophy, History, and Social Studies of Science*. Taylor & Francis.

Sonnenschein, C., & Soto, A. M. (2008). Theories of Carcinogenesis: An Emerging Perspective. *Seminar Cancer Biology*, *18*(5), 372–377.

Sonnenschein, C., & Soto, A. M. (2011). The Death of the Cancer Cell. *Cancer Research*, *17*(13), 4334–4337.

Soto, A. M., & Sonnenschein, C. (1987). Cell Proliferation of Estrogen-sensitive Cells: The Case for Negative Control. *Endocrine Reviews*, *8*(1), 44–52.

Soto, A. M., & Sonnenschein, C. (2005a). Emergentism as a default: Cancer as a problem of tissue organization. *Journal of Biosciences*, *30*(1), 103–118.

Soto, A. M., & Sonnenschein, C. (2005b). Response to Coffman. *BioEssays*, *27*(4), 460–461.

Soto, A. M., & Sonnenschein, C. (2011). The Tissue Organizational Field Theory of Cancer: A Testable Replacement for the Somatic Mutation Theory. *Bioessays*, *33*(5), 332–340.

Spirtes, P., Glymour, C. N., & Scheines, R. (2000). *Causation, prediction, and search.* (2nd ed.). Cambridge, Mass.: MIT Press.

Steel, D. (2004). Can a Reductionist be a Pluralist? *Biology and Philosophy*, *19*(1), 55–73.

Stephens, P. J., Greenman, C. D., Fu, B., Yang, F., Bignell, G. R., Mudie, L. J., … Campbell, P. J. (2011). Massive Genomic Rearrangement Acquired in a Single Catastrophic Event during Cancer Development. *Cell*, *144*(1), 27–40.

Stoljar, D. (2010). *Physicalism*. Oxon: Routledge.

Stratton, M. R., Campbell, P. J., & Futreal, P. A. (2009). The cancer genome. *Nature*, *458*(7239), 719–724.

Strevens, M. (2009). *Depth: An Account of Scientific Explanation*. Cambridge, Mass.: Harvard University Press.

Taylor, E. (2015). An explication of emergence. *Philosophical Studies*, *172*(3), 653–669.

Thalos, M. (2013). *Without Hierarchy: The Scale Freedom of the Universe*. Oxford: Oxford University Press.

The Human Microbiome Project Consortium. (2012). Structure, function and diversity of the healthy human microbiome. *Nature*, *486*(7402), 207–214.

Tomczak, K., Czerwińska, P., & Wiznerowicz, M. (2015). The Cancer Genome Atlas (TCGA): an immeasurable source of knowledge. *Contemporary Oncology*, *19*(1A), A68–A77.

Van Gulick, R. (2001). Reduction, emergence and other recent options on the mind/body problem. A philosophic overview. *Journal of Consciousness Studies*, *8*(9–10), 1–34.

Van Riel, R. (2014). *The Concept of Reduction*. Cham: Springer.

Varzi, A. C. (2006). A Note on the Transitivity of Parthood. *Applied Ontology*, *1*(2), 141–146.

Vaux, D. L. (2011). Response to “The tissue organization field theory of cancer: A testable replacement for the somatic mutation theory”. *BioEssays*, *33*(9), 660–661.

Waters, C. K. (1990). Why the Anti-Reductionist Consensus Won’t Survive: The Case of Classical Mendelian Genetics. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, *1990*, 125–139.

Waters, C. K. (2004). What Concept Analysis in Philosophy of Science Should Be (and Why Competing Philosophical Analyses of Gene Concepts Cannot Be Tested by Polling Scientists). *History and Philosophy of the Life Sciences*, *26*(1), 29–58.

Waters, C. K. (2006). A Pluralist Interpretation of Gene -Centered Biology. In S. H. Kellert, H. E. Longino, & C. K. Waters (Eds.), *Scientific pluralism* (Vol. 19, pp. 190–214). Minneapolis, MN: University of Minnesota Press.

Waters, C. K. (2008). Beyond theoretical reduction and layer cake anti-reduction. In M. Ruse (Ed.), *The Oxford handbook of philosophy of biology* (pp. 238–262). Oxford: Oxford University Press.

Weinberg, R., A. (1998). *One Renegade Cell: How Cancer Begins*. New York: Basic Books.

Weinberg, R. A. (2014). Coming Full Circle—From Endless Complexity to Simplicity and Back Again. *Cell*, *157*(1), 267–271.

Wilson, J. (2016). Metaphysical Emergence: Weak and Strong. In T. Bigaj & C. Wüthrich (Eds.), *Metaphysics in Contemporary Physics* (Vol. 104, pp. 345–402). Leiden: Brill Rodopi.

Wimsatt, W. C. (1976). Reductive Explanation: A Functional Account. In R. S. Cohen, C. A. Hooker, A. C. Michalos, & J. W. V. Evra (Eds.), *PSA 1974* (pp. 671–710). Springer Netherlands.

Wimsatt, W. C. (1994). The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets. *Canadian Journal of Philosophy*, *24*(sup1), 207–274.

Wimsatt, W. C. (2006). Reductionism and its heuristics: Making methodological reductionism honest. *Synthese*, *151*(3), 445–475.

Wimsatt, W. C. (2007). *Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, Mass.: Harvard University Press.

Wolkenhauer, O., & Green, S. (2013). The search for organizing principles as a cure against reductionism in systems medicine. *FEBS Journal*, *280*(23), 5938–5948.

Woo, P. C. Y., Lau, S. K. P., Teng, J. L. L., Tse, H., & Yuen, K.-Y. (2008). Then and now: use of 16S rDNA gene sequencing for bacterial identification and discovery of novel bacteria in clinical microbiology laboratories. *Clinical Microbiology and Infection: The Official Publication of the European Society of Clinical Microbiology and Infectious Diseases*, *14*(10), 908–934.

Woodger, J. (1930). The 'Concept of Organism' and the Relation Between Embryology and Genetics, Part I. *The Quarterly Review of Biology*, *5*(1), 1–22.

Woodward, J. (2003). *Making things happen a theory of causal explanation*. Oxford; New York: Oxford University Press.

Woodward, J. (2015). Interventionism and Causal Exclusion. *Philosophy and Phenomenological Research*, *91*(2), 303–347.

Wright, C. (2012). Mechanistic explanation without the ontic conception. *European Journal for Philosophy of Science*, *2*(3), 375–394.

Wright, C. (2015). The ontic conception of scientific explanation. *Studies in History and Philosophy of Science Part A*, *54*, 20–30.

Yang, E. (2013). Eliminativism, interventionism and the Overdetermination Argument. *Philosophical Studies*, *164*(2), 321–340.

Yusuf, I., & Fruman, D. A. (2003). Regulation of Quiescence in Lymphocytes. *Trends in Immunology*, *24*(7), 380–386.

Zuchowski, L. (2017). Modelling and knowledge transfer in complexity science. *Studies in History and Philosophy of Science Part A, In Press.*

Zuchowski, L. C. (2017). Complexity as a contrast between dynamics and phenomenology. *Studies in the History and Philosophy of Modern Physics, In Press.*