

ABSTRACT

Title of Dissertation: THE METAPHYSICS OF MULTI-LEVEL
EXPLANATION

David Michael McElhoses, Doctor of Philosophy, 2013

Dissertation directed by: Associate Professor Mathias Frisch
Department of Philosophy

It is widely presumed that the world exhibits hierarchical structure: that the objects of science can be arranged into *levels* and that some of those levels are *higher than* others. Organisms, on this view, are at a higher level than cells, which are at a higher level than molecules, which are at a higher level than atoms.

Although it is popular, this “layered worldview” faces notable challenges. Some critics contend that it is hopelessly idealized or even ultimately incoherent. Others contend that it makes no difference to our explanatory practices and has no metaphysical or epistemic significance.

I argue that these critics are mistaken. By undertaking a comprehensive analysis of the logical and metaphysical nature of hierarchical structures and their application within contemporary philosophy and ecology, I argue for three claims: that hierarchies exist (insofar as the objects of scientific discourse exist); that the domain of ecological interests is hierarchically structured in a way that is incompatible with *ontological reductionism* – the idea that everything in the scientific domain is, in fact, contained

within the “fundamental level of reality”; and that the hierarchical structure of the world often (but not always) justifies the practice of describing, explaining, and analyzing things using hierarchical terms.

My analysis begins with a review of the existing accounts of hierarchies in ecology, biology, sociology, and economics. Arguing that these accounts are inadequate, I then develop an improved account called *Core Hierarchy Theory* (CHT). CHT, I argue, is an improvement over its predecessors in two respects: generality and simplicity. Other accounts are either too narrow (e.g., failing to count *branching* hierarchies as genuine hierarchies) or make unnecessary theoretical commitments. Using a formalized version of CHT, I then prove four theorems that are relevant to well-known philosophical debates that involve hierarchies. For example, I show that two of the core metaphysical commitments of the ontological reductionist – that the higher level sciences reduce to fundamental science and that all of reality is, in fact, contained within the fundamental level – are in fact inconsistent with the most basic and unrestricted conception of hierarchical structures provided by CHT.

THE METAPHYSICS OF MULTI-LEVEL EXPLANATION

by

David Michael McElhoses

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2013

Advisory Committee:

Professor Mathias Frisch, Chair
Professor Paul Pietroski
Professor Michael Morreau
Professor Lindley Darden

©Copyright by
David Michael McElhoes
2013

To my parents

ACKNOWLEDGEMENTS

This research was supported by the Social Sciences and Humanities Research Council of Canada.

TABLE OF CONTENTS

LIST OF FIGURES.....	ix
----------------------	----

INTRODUCTION

1.0	THE METAPHYSICS OF ECOLOGY: AIMS, SCOPE, AND STRATEGY.....	1
1.1	ISSUE #1: REDUCTIONISM VERSUS HOLISM.....	2
1.2	ISSUE #2: HIERARCHICAL REALISM VERSUS ANTI-REALISM.....	7
1.3	AN EXAMPLE FROM ECOLOGY: THE NITROGEN CYCLE.....	14
2.0	MULTI-GRADE HOLISM: ASSUMPTIONS, THESIS AND ARGUMENT.....	16
3.0	CHAPTER STRUCTURE AND SUMMARY.....	20

CHAPTER 1: THE ECOLOGICAL HIERARCHY

1.0	INTRODUCTION.....	24
2.0	THE ECOLOGICAL HIERARCHY.....	26
2.1	HOW DEEP IS THE ECOLOGICAL HIERARCHY?.....	27
2.2	ORDERING THE HIERARCHY.....	31
2.3	NESTED VS. NON-NESTED HIERARCHIES.....	34
2.4	BEING “AT” A LEVEL IN THE ECOLOGICAL HIERARCHY.....	38
3.0	CHALLENGES TO THE ECOLOGICAL HIERARCHY.....	42
3.1	RICKLEFS CHALLENGE TO THE COMMUNITY.....	43
3.1.1	<i>Disintegrating the Local-Community</i>	43
3.1.2	<i>Re-integrating the Local-Community</i>	47
3.2	POTOCHNIK AND MCGILL VS. SCIENTIFIC HIERARCHIES.....	50
3.2.1	<i>The Coherence Problem</i>	51
3.2.2	<i>The Nesting Defense</i>	53
3.2.3	<i>The Motivation Problem</i>	55
3.3	THE SIGNIFICANCE OF COMPOSITIONAL HIERARCHIES.....	62
3.4	COMPOSITION VS. SCALE.....	65
4.0	CONCLUSION.....	67

CHAPTER 2: THEORIES OF HIERARCHIES

1.0	INTRODUCTION.....	70
2.0	IDIOSYNCRATIC THEORIES OF HIERARCHIES.....	72
2.1	THE APPROACH FROM ECOLOGY.....	72
2.2	THE APPROACH FROM ONTOLOGY.....	74

2.2.1	<i>Levels: Objects or Categories?</i>	75
2.2.2	<i>Dissolving the Dilemma: Levels Can Be Both</i>	76
2.2.3	<i>Are Levels Parts of Hierarchies?</i>	77
2.2.4	<i>Levels of Containment</i>	80
2.2.5	<i>Levels, Elements, and Parts of Classes</i>	82
2.2.6	<i>Being “at” a Level</i>	91
3.0	THE PARTIAL ORDER VIEW	98
3.1	SIMON ON HIERARCHIES	99
3.2	WEBSTER ON HIERARCHIES.....	100
3.3	HIERARCHIES: SET THEORY VS. MEREOLGY	101
3.4	VOORHEES AND BUNGE ON HIERARCHIES	104
3.5	REJECTING THE PARTIAL ORDER ACCOUNT	107
4.0	THE STRICT PARTIAL ORDER VIEW	109
4.1	DAWKINS ON HIERARCHY	109
4.2	HIERARCHIES AS DIRECTED ACYCLIC GRAPHS.....	112
4.3	THE ISSS ON HIERARCHIES.....	113
5.0	A SUPERFICIAL ACCOUNT OF HIERARCHIES	115
6.0	CONCLUSION.....	116

CHAPTER 3: CORE HIERARCHY THEORY

1.0	INTRODUCTION	118
2.0	CORE HIERARCHY THEORY: SUPERFICIAL AND NAÏVE.....	120
2.1	WHAT IS A HIERARCHY?	121
2.1.1	<i>A Collection of Levels</i>	121
2.1.2	<i>... Linked Together by</i>	125
2.2	... HIGHER THAN	126
2.2.1	<i>Transitivity and Complex Non-branching Hierarchies</i>	127
2.2.2	<i>Higher vs. Lower</i>	130
2.2.3	<i>Hierarchical Identity and Distinctness</i>	132
2.2.4	<i>The SPO Analysis of Higher than.</i>	134
3.0	OBJECTIONS TO CORE HIERARCHY THEORY	145
3.1	THE “TOO GENERAL” OBJECTION.....	145
3.2	THE ARBITRARINESS OBJECTION.....	147
3.3	THE EMPTINESS OBJECTION	148
4.0	CONSEQUENCES OF CORE HIERARCHY THEORY	150
4.1	AN ARGUMENT FOR HIERARCHICAL REALISM.....	151
4.2	TRANSITIVITY OF ‘IS A LEVEL OF’	154
4.3	ASYMMETRY, IRREFLEXIVITY, AND ‘IS A LEVEL OF’	155
4.4	COMPOSITIONAL HIERARCHIES ARE LEVELS OF THEMSELVES?.....	157
5.0	CONCLUSION.....	161

APPENDIX A: A FIRST-ORDER RECONSTRUCTION OF CORE HIERARCHY THEORY	164
APPENDIX B: THEOREMS FOR UNRESTRICTED HIERARCHIES	170

CHAPTER 4: THE LAYERED WORLDVIEW AND THE METAPHYSICS OF REDUCTION

1.0	INTRODUCTION	190
2.0	THE LAYERED WORLDVIEW	193
2.1	GLOBAL HIERARCHIES	193
2.2	LOCAL HIERARCHIES	194
2.3	GLOBAL VS. LOCAL HIERARCHIES	195
3.0	HIERARCHIES AND THE METAPHYSICS OF REDUCTION	195
3.1	THE “OLD” REDUCTIONISM	196
3.2	THE “NEW” REDUCTIONISM	198
3.3	REDUCTION, REALIZATION, AND HIERARCHY	200
4.0	OPPENHEIM AND PUTNAM ON THE MICROREDUCTION HIERARCHY	201
4.1	FREE-MOLECULES AND OTHER PROBLEMS	204
4.2	THE METAPHYSICS OF REDUCTION – OLD AND NEW	207
5.0	WIMSATT ON THE MICROREDUCTIONIST HIERARCHY	209
5.1	WIMSATT’S INITIAL ASSUMPTIONS ABOUT LEVELS	209
5.2	COMPOSITION, EXPLANATION AND SCALAR ORDERINGS	210
5.3	WIMSATT VS. OPPENHEIM AND PUTNAM	214
6.0	THE NESTING PROBLEM FOR REDUCTIONISM -- OLD AND NEW	220
6.1	OPTION 1: THERE IS NO HIERARCHY!	225
6.2	OPTION 2: REJECT CORE HIERARCHY THEORY	226
6.3	OPTION 3: REINTERPRET ‘@’ IN CORE HIERARCHY THEORY	228
6.4	OPTION 4: THE REDUCTIONIST’S HIERARCHY AS A ϕ -@ HIERARCHY	231
6.5	OPTION 5: ‘HIERARCHY’ IS AMBIGUOUS	232
6.6	OPTION 6: CHALLENGING THE ASSUMPTIONS	235
7.0	CONCLUSION	237

CHAPTER 5: THE LAYERED WORLDVIEW AND THE METAPHYSICS OF MULTI-LEVEL EXPLANATION

1.0	INTRODUCTION	239
2.0	THE BASIC METAPHYSICAL FRAMEWORK	242
2.1	THE HUMEAN VS. MECHANISTIC APPROACHES	242
2.2	JACKSON AND PETIT ON EFFICACY VS. RELEVANCE	243
3.0	GLOBAL ORDERINGS: SUPERVENIENCE AND NESTING	245

3.1	SUPERVENIENCE AND THE LAYERED WORLD	245
3.2	ARGUING AGAINST SUPERVENIENCE ORDERINGS	247
3.3	FORGING A GLOBAL HIERARCHY.....	249
3.4	ARGUING AGAINST MONOLITHIC GLOBALISM	251
3.5	AN UPWARDS-NESTED GLOBAL HIERARCHY	252
4.0	LOCAL ORDERINGS: THE DUAL-HIERARCHY MODEL	255
4.1	CAUSAL EXCLUSION AND FUNCTIONAL REDUCTION	255
4.2	ORDERS VS. LEVELS: THE DUAL HIERARCHY VIEW	258
4.3	THE DUAL-HIERARCHY VIEW, REDUCTION, AND CORE HIERARCHY THEORY	260
5.0	THE MECHANISTIC HIERARCHY	261
5.1	WHAT IS A MECHANISM?	262
5.2	GLENNAN ON THE DUAL-HIERARCHY VIEW	263
5.3	PRODUCTIVE ACTIVITIES VS. CAUSAL RELEVANCE	265
5.4	CRAVER AND THE LAYERED WORLDVIEW: LEVELS OF MECHANISMS.....	266
5.4.1	<i>Craver, Mechanisms and Globalism</i>	268
5.4.2	<i>Craver on Levels-of-Mechanisms</i>	271
5.5	MECHANISMS, CAUSATION AND MULTI-LEVEL EXPLANATION.....	274
5.5.1	<i>Mechanistic Explanation and Programmatic Explanation</i>	276
5.5.2	<i>Mechanisms and Metaphysical-MLE</i>	277
5.5.3	<i>Nested Hierarchies and Metaphysical-MLE</i>	279
5.6	THE CAUSAL EXCLUSION PROBLEM REVISITED	283
5.6.1	<i>Mechanistic Hierarchies are Dimensioned Realization Hierarchies</i>	284
5.6.2	<i>Causally Excluded Activities and Mechanisms</i>	286
6.0	CONCLUSION.....	288

CHAPTER 6: PROBLEMS IN THE CAUSAL BASEMENT

1.0	INTRODUCTION	291
2.0	THERE IS NO HIGHER LEVEL: THE ELEATIC OBJECTION	292
2.1	THE CAUSAL EXCLUSION ARGUMENT: SIMPLY STATED	292
2.1.1	<i>The Problem of Causal Drainage</i>	295
2.1.2	<i>The Diffusion of Exclusion: Causal Selfishness</i>	298
2.1.3	<i>The Diffusion of Exclusion: The Triviality Response</i>	301
2.1.4	<i>The Diffusion of Exclusion: The Near-Identity Response</i>	303
2.2	THE GENERAL ELEATIC ARGUMENT	306
2.2.1	<i>The Eleatic Principle</i>	307
2.2.2	<i>The Higher level Does Not Make a Difference</i>	308
3.0	THERE IS NO REALIZATION: THE ELEATIC PROBLEM ONCE AGAIN.....	310
3.1	FOUR TESTS FOR DIFFERENCE-MAKING.....	310
3.1.1	<i>The Quinean Test</i>	311
3.1.2	<i>The Modal Juxtaposition Test</i>	312
3.1.3	<i>The Contrastivist Test</i>	314
3.1.4	<i>The Refined Contrastivist Test</i>	317

4.0	RESURRECTING THE REALIZATION HIERARCHY	318
4.1	THE REIFICATION OF REALIZATION: IT MAKES A DIFFERENCE.....	318
4.1.1	<i>The (Causal) Disconnection Between Levels</i>	319
4.1.2	<i>Co-opting the Contrastivist Test</i>	321
4.2	THE REIFICATION OF THE HIGHER LEVEL: IT MAKES A DIFFERENCE	323
4.2.1	<i>The Contrastivist Strategy Revisted</i>	323
4.2.2	<i>Implementing the Contrastivist Strategy</i>	324
4.2.3	<i>The Problem for Property Realization</i>	326
4.3	ELIMINATING THE ELEATIC PRINCIPLE.....	328
5.0	IMPLICATIONS FOR THE CAUSAL EXCLUSION ARGUMENT	331
5.1	WEIGHING THE OPTIONS	332
5.2	INTERPRETING THE RESULT.....	335
6.0	CONCLUSION.....	337

LIST OF FIGURES

Chapter 1:

Figure 1: The Nested Hierarchy.....	34
Figure 2: Ricklefs' Two Interpretations of the Local Community.....	44
Figure 3: A Nested Hierarchy -- The Lower Level is Contained Within the Higher Level	54

Chapter 2:

Figure 4: An Organizational Hierarchy For a Limited Liability Partnership.....	106
Figure 5: A Broken Hierarchy.....	113

Chapter 3:

Figure 6: A Simple Non-Branching Hierarchy.....	126
Figure 7: The Ambiguity of the Missing Higher Than.....	127
Figure 8: Two Distinct Hierarchies H_1 and H_2	132
Figure 9: The Most Basic Higher Than/Hierarchical Structure.....	134
Figure 10: An Upward Branching Hierarchy as Two Basic Hierarchies.....	135
Figure 11: A Downward Branching Hierarchy as Two Basic Hierarchies.....	136
Figure 12: A Zigzagging Strict Partial Order.....	145
Figure 13: A Reflexive Hierarchy.....	158
Figure 14: The Nested Hierarchy.....	159
Figure 15: A Direct Compositional Relationship vs an Indirect Compositional Relationship.....	175

Chapter 4:

Figure 16: Wimsatt's Waveforms -- Regularity and Predictability Across Levels.....	210
Figure 17: Wimsatt's Reductionist (Quasi)Hierarchy.....	214
Figure 18: Kim's Nested Global Hierarchy.....	249

Chapter 5:

Figure 19: The Dual-Hierarchy Approach.....	258
Figure 20: Craver's Depiction of a Mechanism With Three Levels.....	271

Chapter 6:

Figure 21: H_s Realized by L_s	324
Figure 22: Two Causal Situations.....	329
Figure 23: Eight Difference-Making Options.....	331
Figure 24: Inter-level Causation.....	334
Figure 25: Realizing Causation	335

INTRODUCTION

1.0	THE METAPHYSICS OF ECOLOGY: AIMS, SCOPE, AND STRATEGY	1
1.1	ISSUE #1: REDUCTIONISM VERSUS HOLISM.....	2
1.2	ISSUE #2: HIERARCHICAL REALISM VERSUS ANTI-REALISM.....	7
1.3	AN EXAMPLE FROM ECOLOGY: THE NITROGEN CYCLE	14
2.0	MULTI-GRADE HOLISM: ASSUMPTIONS, THESIS AND ARGUMENT	16
3.0	CHAPTER STRUCTURE AND SUMMARY	20

1.0 The Metaphysics of Ecology: Aims, Scope, and Strategy

Philosophical studies in ecology generally come in two varieties: the normative, which studies the moral or aesthetic value of nature; and the epistemological, which studies the methods of investigation and explanation within ecology as a science. Inquiry into the metaphysics of ecology is a far less common endeavor by comparison.

The metaphysical issues that do arise within ecology are often tied-up with epistemological issues, sometimes leaving a murky boundary between them. The more common strategy is to avoid the murkiness by focusing only on the epistemological issues, perhaps with some hope that settling them helps to inform metaphysical inquiry.

Here I take the less-ventured path through the murk. I aim to investigate the *metaphysics* of ecology in a way that informs its epistemology. It is a work directed towards metaphysicians, philosophers of science, philosophers of ecology, and theoretical ecologists. My strategy is to adopt the standard philosophical presuppositions found within the theoretical ecologist's framework and then to reason about its metaphysics, keeping epistemology within the periphery.

Whether or not such a strategy is legitimate depends on the problem being investigated. For one, the problem must have both a metaphysical and an epistemological component, and for another, the murky connection between them must resolve in manner that appears conducive to metaphysical study. Any issue which meets these legitimizing conditions falls within the scope of our inquiry.

There are at least two problems that meet these criteria in ecology. First, there is the debate between reductionists and holists. And second, there is a disagreement among holists (Allen and Starr 1982; c.f. Salthe 1985, 2009) regarding the existence of hierarchies and their relation to the natural world. My tactic here is to examine both of these issues simultaneously, by investigating a notion that is relevant to reductionists, holists, hierarchical realists and anti-realists alike: the notion of a hierarchy.

1.1 Issue #1: Reductionism versus Holism

The debate between reductionists and holists serves as the main backdrop to my investigation. The debate meets both of the aforementioned legitimizing conditions. It has both a metaphysical and epistemological component, and these components are related in a way that is conducive to metaphysical inquiry.

The metaphysical component of reductionism is the idea that the whole and its properties are “nothing above” its parts and their properties (in some sense of ‘nothing above’ that remains mysterious and elusive). The epistemological component involves the claim that knowledge of the part’s properties and relationships is sufficient for knowledge of the whole’s properties (Keller and Golley 2000, p. 172). Here the connection between metaphysics and epistemology appears to be clear. David Keller and

Frank Golley echo common opinion when they urge us to “bear in mind that epistemic reduction presupposes ontological reduction.” (*ibid.*)

Holism is taken to be the polar opposite of reductionism. The metaphysical component of holism involves the claim that the whole and its properties are indeed “something above” its parts and their properties (in some sense of ‘something above’ that remains mysterious and elusive).¹ The epistemological component of holism involves the claim that knowledge of the parts is insufficient for knowledge of the whole. (*ibid.*) Complete ecological understanding, for the holist, requires understanding both parts and wholes in tandem, giving rise to a sort of explanation and a sort of analysis popularly called *multi-level* explanation and analysis.

Again, the connection between metaphysics and epistemology is clear: epistemic holism presupposes metaphysical holism. If metaphysical holism is false, after all, then metaphysical reductionism is true: the whole is “nothing above” its parts. But this idea does not fit well with epistemic holism. Insofar as the whole is “nothing above” its parts, the reductionist takes it to be nothing *more* than its parts. Thus, if metaphysical holism is false, then there is nothing *more* than the parts for us to understand, *pace* the epistemic holist. Epistemological holism therefore implies its metaphysical counterpart.

Although they remain notoriously unclear, the locutions ‘nothing above’ and ‘something above’ can be understood as relating to a hierarchical worldview: a worldview in which the domain is partitioned into different levels, and some things are

¹ This metaphysical component of holism should not be confused with the claim that the whole is “more than *the sum* of its parts” (note the emphasis on ‘sum’). This is the rallying cry of the emergentist, and not the ecological holist. I have no intention of investigating whether or not one implies the other. Emergence in ecology is an issue that lies beyond the scope of this work.

considered “higher than” others. Such a worldview provides the basis upon which the two views are distinguished. This is one respect in which hierarchies play a role in the reductionism-holism debate. Both sides of the debate have some sort of hierarchical picture of ecology affixed in their minds: the holist places wholes above their parts in that hierarchy, whereas the reductionist does not.

To avoid confusion, note that by ‘reductionism’ I do not mean to refer to any particular *model* of scientific reduction, such as Ernest Nagel’s (1961) or Kemeny and Oppenheim’s (1956) models.² Rather, I am referring directly to the general idea that *inspires* these sorts of models: (roughly) the view that swarms of fundamental particles, their behaviors, and interactions, are all that there *really* is to nature, and are all that is *really* needed to give an account of it.

We should be wary of characterizing reductionism in this way, for the use of ‘really’ is even more mysterious than the notorious ‘nothing above’. If ‘all there *really* is’ means *all there is, simpliciter*, then reductionism is no different from eliminativism (the radical view that nothing physically exists except for the smallest bits of matter and their properties). But the ecological reductionist is not so extreme.³ If ‘really’ means that fundamental particles are, in some sense, *ontologically privileged*⁴ over their composites, then we need both an account of ‘privilege’, and a way to resolve any other mysterious

² Both of which presuppose the deductive-nomological model of explanation advanced by Hempel and Oppenheim (1948).

³ On my understanding, reductionism occupies the middle ground between the extremes of holism and eliminativism. I will not be entertaining the eliminativist view until the final chapter since it runs contrary to standard assumptions found within ecology (see §5.0 below): that organisms and their environments actually exist. Within the philosophy of ecology, the debate is between reductionists and holists; only within metaphysics, as a general field of study, is such a strong eliminativism taken seriously.

⁴ Dan Steel (2004, p. 66), for example, characterizes reductionism as *privileging* a fundamental level.

terms used to formulate *that* account. None of this is necessary on our original characterization. We can avoid the, perhaps infinite, regress-of-mystery by simply adopting the straightforward hierarchical understanding of ‘above’ that both parties have in mind.⁵

That both parties have *a* hierarchical picture in mind is not enough. If the reductionist-holist debate is to make any sense, then the *same* hierarchical picture of ecology must be shared by both parties. Hierarchies are then philosophically useful. They allow us to get away from the notorious locutions ‘nothing above’ and ‘something above’ without introducing expressions of equal notoriety. They allow us to straightforwardly characterize the metaphysical difference between reductionism and holism as follows.

[Reductionism] If the world is hierarchical in the sort of way portrayed within ecology, then wholes are there along with their parts at the bottommost level, whatever that might be.

[Holism] The world *is* hierarchical in the sort of way portrayed within ecology, but the wholes are *not* there along with their parts at the bottommost level, whatever that might be.

By removing the mysterious locution ‘above’ from the metaphysical characterization of reductionism and holism, the distinction between them becomes clear.⁶ Although it

⁵ The pairing of hierarchies with reductionism is nothing new. Dan Steel (2004, p. 60) notes that reductionism is often understood in a way that makes an appeal to hierarchical levels.

⁶ I will henceforth designate the pre-theoretic notion with a lower-case letter, e.g., ‘reduction’, whereas my characterization of them will be designated using uppercase letters, e.g., ‘Reduction’.

remains unclear what the bottommost level is, this is a matter that can be settled empirically.

This distinction might seem suspicious to some. Dan Steel (2004) argues, for example, that there is a weaker sort of reductionism that is compatible with holism.⁷ He calls it *token-token reductionism*. He argues that there are sometimes exceptions to holistic generalizations, and that these exceptions can be accounted for within an explanation involving posits residing at a more “fundamental” level.

Such suspicion is unwarranted. By distinguishing reduction from holism as I do, notice that Steel’s notion of weak-reduction counts as a type of holism. This, I take it, is an advantage of my approach: certainly no *holist* has ever denied that there is a type of reduction that is compatible with holistic explanation; in fact, it is part of the holist view that some explanations involve posits residing at lower levels. As Frank Golley notes, “the holist understands that the results of reductionist research are always relevant at some level in understanding a phenomenon... is tolerant of the reductionist agenda and even supportive of it” (1993, p.28) That there are aspects of reductionism which are compatible with the holist agenda is thus no reason to rouse suspicion over my characterization.

Besides, my manner of articulating the distinction is useful. Not only does it help to clarify the debate,⁸ but it allows for a focused examination of its conceptual

⁷ Steel uses ‘pluralism’ in lieu of ‘holism’. While there might be subtle differences between the meanings of these terms, they do not seem relevant for our purposes.

⁸ I am not the first to take note of this. Stanley Salthe boldly claims that “Reductionism in fact requires hierarchical structure in order to make sense.” (1985 p. ix) While I am sympathetic to his boldness, I am unsure of what he means by “required”. Hierarchical terminology might be useful in characterizing reductionism, but truth of reductionism hardly depends on the world’s being genuinely hierarchically

foundations. At the heart of my distinction is the notion of a hierarchy. Resolving outstanding issues about hierarchies should then, at the very least, help to enrich the debate, and, at the very most, help to settle it.

1.2 Issue #2: Hierarchical Realism versus Anti-Realism

Among the most prominent methodological approaches to holism in ecology one finds the widely cited *Hierarchy Theory* of T. F. H. Allen and Thomas Starr (1982). These authors defend the idea that ecological understanding involves a process of *multi-level* explanation and analysis that emphasizes the role of the observer in the explanatory and analytical process. Their ideas have been influential. Not only is their methodological and analytical approach widely cited within ecology, it also serves as a promising basis for a *unified* ecology (Allen and Hoekstra 1992).⁹

Still, there is disagreement among Holists regarding the nature and existence of hierarchies. Allen, Starr, and Victoria Ahl argue that hierarchies should not “be interpreted as features of the external world, existing independently of an observer’s criteria for delimiting the system.” (Allen and Ahl 1996, p. 33) They take hierarchies to be purely conceptual artifacts.

structured. Notice that, on my characterization, Reductionism does not require anything to be hierarchical. In fact, this is an advantage of my characterization. The reductionist does not quake at the prospect of eliminativism, she simply takes it to be unnecessary and extreme. My conditional characterization captures this: Reductionism is vacuously true if eliminativism is correct. Eliminativists tell us that the world is *not* hierarchical in the sort of way portrayed within ecology, since there are, they say, no ecological entities to be hierarchically arranged.

⁹ Here I am largely unconcerned with the unification movement. Although I will touch on the unification motivation for Hierarchy Theory early on, it is largely tangential to my investigation: it appears to be a purely epistemological issue.

One reason for this approach, given by Allen and Starr, is that there are an infinite number of ways to delineate levels, and none are more ontologically privileged than any of the others; they argue that there is nothing less real about a level consisting in a person with an air-pocket around them than one consisting in the person alone, or one consisting in the person minus a nanometer-thick layer of skin. Since what matters, says Allen and Starr, is the information we gain from our selection of levels, and our selecting levels does not require their reification, there is hence no need to reify them; to reify an infinite number of levels without being required to do so is senseless. Taking hierarchies to be conceptual artifacts, they think, is good enough for science. (1982, p. 10)

Other philosophers of ecology disagree. Stanley Salthe (1985), Richard Levins and Richard Lewontin (2000) claim that hierarchies are real features of the world. Levins and Lewontin argue that a “proper materialism... looks for the actual material relationships among entities at all levels.” (2000, p. 221) Salthe’s metaphysical commitment is made clear when he expresses that he “will not abandon the entitative ontological perspective in favor of the kind of operationalist phenomenalism suggested by, for example, Allen and Starr.” (1985, p. 21)

As with the reductionism-holism debate, this second, far less prominent, issue meets both legitimizing conditions. It has a metaphysical component, an epistemic component, and the components connect up in a way that is favorable to metaphysics. The metaphysical component of the debate is captured in the following distinction between *hierarchical realism* and *hierarchical anti-realism*.

- [Realism] Not all hierarchies are intellectual artifacts or abstracta; i.e. the natural world is a hierarchy or contains hierarchies.
- [Anti-Realism] All hierarchies are either intellectual artifacts or abstracta; i.e. the natural world is not a hierarchy and does not contain hierarchies.

Although they do not characterize their dispute in this way, it should be clear that Salthe, Levins, and Lewontin are committed to hierarchical Realism, whereas Allen, Starr, and Ahl are committed to Anti-Realism.

The epistemological component to their debate involves their respective justification for a multi-level approach to explanation¹⁰ and analysis. Paired with Realism is the metaphysical justification,

- [Metaphysical-MLE] The hierarchical structure of the world justifies our describing, explaining, and analyzing things using hierarchical terms.

Paired with Anti-Realism is the epistemological justification,

- [Epistemic-MLE] The *usefulness* of hierarchical structure justifies our describing, explaining, and analyzing things using hierarchical terms.

¹⁰ Note that at this early stage, I use ‘explanation’ in the *social* sense: as a description that one person gives to another so that the recipient might gain understanding. Later, I will discuss a *metaphysical* sense of explanation that involves the description of a mechanism (Machamer, Darden, and Craver, 2000; Glennan 1996, 2002).

Although my focus is on ecology, multi-level explanation and multi-level analysis are notions that extend beyond their ecological applications. Patrick McGivern, who is primarily concerned with physics, tells us that, “For many systems it is *necessary* to describe them simultaneously on multiple levels... in order to get an accurate representation of their behavior... [it] is extremely useful in modeling systems in a wide variety of disciplines, including fluid mechanics, aerodynamics, chemistry, ecology, geophysics, and neuroscience.” (2008 p. 54, emphasis added; Anderson 1972, pp. 393-396 makes a similar claim)

McGivern’s usage of ‘level’ is clearly the one which relates to hierarchies, rather than that which relates to building tables and hanging picture-frames. His usage of ‘necessary’, on the other hand, is puzzling. He claims that multi-level descriptions are necessary to get an accurate depiction of behavior; but necessary in what sense? And why is it necessary in that sense?

McGivern might mean that multi-level descriptions are sometimes causally-necessary in that some phenomena can be *produced* only via the coordination of things spanning different levels. Or he might mean that multi-level descriptions are epistemologically-necessary in that some phenomena can be *understood* only by representing systemic behavior on many levels. Perhaps both senses are meant: the metaphysical reading might even imply the epistemological reading; it depends on what one means by ‘produced’, ‘coordination’, and ‘spanning’.

Notice that the ambiguity of McGivern’s claim fits nicely within the dispute over Realism, as well as the corresponding Metaphysical-MLE and Epistemic-MLE

distinction. Explanation in ecology, it is widely presumed, involves causal notions.¹¹ The causal reading of ‘necessary’ then seems to support the metaphysical justification: if the production of some phenomena requires multi-level causal coordination – i.e., different causal relations holding between different things residing at different levels – then the hierarchical structure of causation appears to imply (or at least suggest) Metaphysical-MLE. Likewise, the epistemic reading seems to support the epistemic justification: if multi-level descriptions are necessary for our understanding, then Epistemic-MLE becomes hard to deny.

Although there clearly is a connection between Realism, Anti-Realism, and multi-level explanation, it is a weaker connection than what we see in the Reductionist-Holist debate. The epistemic components – Metaphysical and Epistemic-MLE – are not mutually exclusive. Realists can endorse both justificatory approaches: even if the world is, in fact, hierarchical, it might also be that our use of hierarchical terminology is justified instrumentally. Epistemic-MLE, then, does not presuppose Anti-Realism; it is compatible with both Realism and Anti-Realism.

Nevertheless, the connection between the metaphysical and epistemic components of this second issue is sufficient to legitimize our strategy: Metaphysical-MLE appears to presuppose Realism; if the world is not really hierarchical, then the use of hierarchical terminology is not justified by virtue of the world’s hierarchical structure (and the causal disambiguation of McGivern’s claim is then false). The issue over Realism and Anti-Realism, then, is relevant to epistemology, as our strategy requires. If we can show

¹¹ The issue of causation and explanation in ecology is too complicated to be teased apart here. I leave it as a matter for future research.

Realism to be false or problematic, then we have made progress towards an epistemological issue.

Aside from the relationships found within each issue, it is also worth noting the connections *between* the two issues. Reductionism, for example, can be held along with either Realism or Anti-Realism,¹² but is clearly incompatible with both Epistemic and Metaphysical-MLE. Reductionism implies their negation: if accounting for the behavior of the parts at the bottommost level suffices to account for the behavior of wholes, then the account need not mention any other level aside from the bottom. Hierarchies become irrelevant to ecological understanding under Reductionism since it implies that there is no need to coordinate multiple levels to achieve such understanding. Epistemic and Metaphysical-MLE, therefore, presuppose Holism.

The Reductionism-Holism debate also helps to answer an interesting question regarding the connection between Realism and Metaphysical-MLE. Evidently, the latter implies the former, but what of the reverse: does Realism imply Metaphysical-MLE? If Reductionism is true, then Realism has no such implication. Reductionist-Realism is a *prima facie* consistent position that implies the negation of Metaphysical-MLE, for the reasons described in the paragraph above. On the other hand, if Holist-Realism is correct, then it is indeed plausible to suggest that multi-level explanation has a metaphysical justification, and the causal reading of McGivern's claim also gains intuitive credibility.

¹² The existence (or non-existence) of hierarchies does not directly imply Holism, since the hierarchies themselves might be said to be, in reality, "nothing above" the things they contain, or are composed of. In fact, this is a classic way to characterize the reductionist view. Reductionists such as Oppenheim and Putnam (1958) and Jaegwon Kim (2002) pair reductionism with a principle called the *downward-inclusion* principle: all levels, they say, are included within the bottommost level. This principle plays a decisive role in later chapters.

The connections found within, and between, our two main issues are clearly quite complicated. To help simplify matters, let us collect together the most important relationships noted so far, and proceed under the assumption that they are correct:

- (a) Holism is true if and only if Reductionism is not.
- (b) Realism is true if and only if Anti-Realism is not.
- (c) The epistemic component of reductionism presupposes (metaphysical) Reductionism.
- (d) Metaphysical-MLE presupposes Realism.
- (e) Neither Realism nor Anti-Realism implies Holism.
- (f) Reductionism implies that neither Metaphysical-MLE nor Epistemic-MLE is correct.

As with the Reductionism-Holism debate, we do not investigate the Realism vs. Anti-Realism disagreement directly. There is little point. The arguments in favor of either position are hopelessly few. Anti-Realists do nothing to refute Realism; they simply think that it is, for their purposes, unnecessary. Likewise, Realists do nothing to refute Anti-Realism; they think that an avoidance of ontology is unnecessary.

Instead, I investigate this smaller issue indirectly. I use it as a backdrop for a more focused investigation into hierarchies. Resolving outstanding issues about hierarchies, at the very least, helps to clarify and to enrich the disagreement, and, at the very most, helps us to settle it.

1.3 *An Example From Ecology: The Nitrogen Cycle*

Before proceeding to describe my central thesis and argument structure, it is useful to see how these debates play out with respect to an example from actual ecology. Consider nutrient cycling (Gibson 2009, p. 144).

Nutrient cycling tracks the movement of nutrients (e.g., nitrogen and phosphorus) throughout an ecosystem. Understanding how nutrients are cycled within an ecosystem helps to explain changes in landscape and biodiversity. As such, it is important for conservation efforts. Human activities can intervene on these cycles in a significant way.

To explain lake-eutrophication, for example, it is often important to consider changes to the nitrogen cycle. Nitrogen is initially found in the atmosphere. Bacteria convert the nitrogen into organic nitrogen, which is usable by plants to aid in photosynthesis. Plants absorb the nitrogen in the form of ammonium and/or nitrate ions via root hairs. The process is called *assimilation*. Animals eat the plants, defecate, and then die and decompose. The plants themselves die, their leaves fall, and they too decompose. This decomposition is facilitated by *ammonification*: bacteria and fungi convert the organic nitrogen from the decomposing entity into ammonium. Other bacteria (e.g., the *Nitrosomonas* species) then convert the ammonium to nitrites (which are toxic to plants in high levels), at which point other bacteria (e.g., the *Nitrobacte* species) convert the nitrites into nitrates; the process is called *nitrification*. Lastly, *denitrification* begins: bacteria use the nitrates instead of oxygen as an electron acceptor in respiration, releasing dinitrogen gas back into the atmosphere. With dinitrogen gas back in the atmosphere, the nitrogen cycle is complete.

Through human intervention, more nitrogen is added to the ecosystem than what can be effectively cycled. When the excess nutrients enter a lake an overabundance of phytoplankton and cyanobacteria is the result. The overabundance of phytoplankton crowds out other marine plants, causing the lake to become hypoxic, and cyanobacteria release neurotoxins into the lake. Marine vertebrates do not stand a chance: a dead lake results.

The explanation we see here involves plants (growth rates, death rates, decomposition rates), animals, nitrogen, fungi and bacteria of various sorts. All of these entities engage in specialized behaviors that are relevant to the cycling of nitrogen; and, moreover, they are commonly thought to correspond to different hierarchical *levels*. Each of the different organisms is part of the ecosystem, and different parts of the organisms involved are relevant to the cycling of nitrogen.

The Reductionist tells us that to acquire complete understanding of how nitrogen is cycled it is enough to describe the cycle at the finest grain of detail possible: at the level of the ecosystem's tiniest parts using the terminology of fundamental physics. For the Holist, to understand the nitrogen cycle requires understanding that the ecosystem involved is a, say, *grassland* ecosystem, as well as understanding the individual contributions of its parts to the cycling of nitrogen. Fundamental physical descriptions, for the Holist, are not enough: one must understand bacteria as bacteria, not as swarms of fundamental particles, in order to understand the cycling of nutrients.

Hierarchical Anti-Realists tell us that we are making an implicit appeal to hierarchical levels when we described nutrient cycling, not because the world really is

hierarchical, but rather, because that is the way we *perceive* the world to be. Realists express the opposite opinion. They say that the grassland ecosystem really is hierarchical, and this is why the cycling of nutrients can only be understood by describing processes that reside on a variety of different levels.

How, then, should ecology proceed? Should ecologists understand nutrient cycling Reductively or Holistically? If they should understand it Holistically, should they take a Realist position or not? Answering these questions is not a trivial matter. Some initial assumptions are required.

2.0 Multi-Grade Holism: Assumptions, Thesis and Argument

I begin with assumptions that are standard within the various branches of ecology. Keller and Golley (2000) list three basic assumptions: naturalism, scientific realism¹³, and the “comprehensive scope of ecology” (p. 11).

Naturalism, Keller and Golley say, is the idea that “there is but one system of reality... that can be explained using the same methods and terminology.” (*ibid.* p. 12) By ‘scientific realism’, they mean what they call *metaphysical-epistemic realism*, “the position that there are patterns or regularities in nature existing independently of human perceivers, and these patterns, to some extent, are objectively knowable by humans.”

¹³ Note that their usage of ‘scientific realism’, should be distinguished from the special sort of Realism specified in the previous section. By ‘Realism’ I meant *hierarchical* realism: the claim that hierarchies, in fact, exist.

(*ibid.*) By ‘comprehensive scope’, they mean that ecology studies “all of nature” (*ibid.*), by studying the interactions between organisms and their environment.¹⁴

Within these three assumptions, one finds three more: that the “ecological” entities, processes, and phenomena that make up the “pattern of nature” are all *real*; that they interact with one another *causally*,¹⁵ betraying a mild anti-Humean stance; and that they have *parts*.

To these assumptions I add some of my own prejudices. First, my understanding of *parthood* is that of classical mereology (e.g., Simons, 1987) supplemented with the doctrine of mereological monism, the idea that there is one supremely general and fundamental notion of parthood under which all specific sorts of parthood (e.g., material-part, functional-part, undetached-part) and which applies across all ontological categories;¹⁶ second, that the notion of *causation* that is relevant to science is adequately modeled by manipulationist theories of causation (Woodward, 2003); and third, that causal *explanation* proceeds along mechanistic (Machamer, Darden, and Craver 2000;

¹⁴ I am uncertain as to whether Keller and Golley mean that ecology takes scope over all such interactions *at all levels*, including the fundamental level of physics, or whether its scope is restricted to those levels which are, in some sense, *organic*. This uncertainty is no barrier to the present aim. If it were a barrier, we might instead proceed on the organic reading: the idea that ecology takes scope over physics strikes me as being too bizarre without rigorous motivation.

¹⁵ Contrast this with physics, where statistical correlation is commonly thought to replace the need for robustly causal notions.

¹⁶ Mereological monism is the subject of some controversy in analytic metaphysics, as is classical mereology. Nevertheless, I shall make no attempt resolve or settle the controversy. This is work in the metaphysics of ecology, not mereology. I take classical monistic mereology for granted because, for one, it is the standard approach, and for another, it is standard for a good reason. The basic conception of ‘part’ at the heart of classical mereology is quantitative: parts are *some of* the whole and anything which is *some of* the whole is a part of it. Insofar as *some of* is monistic, parthood is too. (Sider 2007) And insofar as *some of* is reflexive, anti-symmetric, and transitive, so is parthood; these three properties together form the core of classical mereology. Classical monistic mereology (supplementation principles notwithstanding), then, is a very short leap from this basic quantitative conception of parthood.

c.f., Glennan 1996, 2002) or Bayesian lines (Woodward, 2003)¹⁷. Although controversy surrounds all three of these assumptions, they are quite common within philosophy. I make no effort to defend them here.

These assumptions and foundational prejudices are used to develop and defend a view called *Multi-Grade Holism*. This view consists in the following three claims, the conjunction of which is my central thesis.

- (T1) Realism is true; Anti-Realism is false.
- (T2) Metaphysical-MLE is false (as a general principle).
- (T3) Reductionism is false; Holism is true.

What might seem initially odd about this view is the conjunction of (T1) and (T2). Realism and Metaphysical-MLE were paired together in our initial characterization. It was also noted that rejecting (T1) was the most obvious way to *affirm* (T2), using claim (d) above as a key premise.

Nevertheless, we will see later on that the sort of Realism to be defended has characteristics which appear to undermine the metaphysical warrant for multi-level explanation. Note, however, that (T2) is not an *essential* characteristic of the view I defend; I am happy to do away with it should someone else is able to solve the problems I raise for Metaphysical-MLE in ecology.

As noted earlier, the tactic I use to achieve my strategic aim has nothing to do with addressing the specific arguments for or against any of the positions mentioned.

¹⁷ Or a combination of the two lines. I use ‘Bayesian’ rather than ‘manipulationist’ since I take it, *pace* Woodward, that explanation is a passively predictive enterprise. Woodward has an active view of explanation: he thinks that to explain a phenomenon is to describe how to manipulate the occurrence of that phenomenon.

Instead I investigate their conceptual foundations: I develop a theory of hierarchies that is suitable to help resolve these issues without begging any questions. That theory of hierarchies is called *Core Hierarchy Theory* and it plays a central role in the argument for Multi-Grade Holism. The overarching structure of the argument is as follows:

- (P1) The ecological domain is a compositional hierarchy. (Chapter 1)
- (P2) Core Hierarchy Theory is our best (*qua* simplest and most general) theory of hierarchies. (Chapter 2-3)
- (P3) Core Hierarchy Theory implies (T1): Realism. (Chapter 3)
- (P4) Core Hierarchy Theory implies (T3): Reductionism is false. (Chapter 4)
- (P5) Core Hierarchy Theory appears to imply (T2): Metaphysical-MLE appears false with respect to some hierarchies. (Chapter 5)
- (C) Thus, our best theory of hierarchies supports (T1)-(T3): Multi-Grade Holism.

If the argument is sound, then my initial aim is achieved: Reductionism has an epistemic component and Metaphysical-MLE is itself partially an epistemic claim. Insofar as Core Hierarchy Theory is the best account of ecological hierarchies, this argument presents a case in which the metaphysics of ecology informs its epistemology, at least insofar as we are compelled to accept the claims that are implied by our best theories.

3.0 Chapter Structure and Summary

The argument Multi-Grade Holism begins with an inquiry into the use of hierarchies within ecology, proceeds to an investigation into the very notion of a hierarchy, and then uses the results of the latter investigation to argue for Multi-Grade Holism. A brief summary of each chapter is found below.

Chapter 1: The first chapter is a case-study in the actual use of hierarchies within science. Ecology, for example, is standardly thought of (by ecologists) as being a hierarchical science. But there are a number of issues surrounding this idea. The most pressing is that, at present, there is no clear and concise account of the ecological hierarchy. It is most difficult to investigate the nature and significance of something when you do not know what that something is. Using the work of ecologist Robert Ricklefs as a foil, I formulate a clear and concise account of the ecological hierarchy, and I defend it from some objections.

Chapter 2: The second chapter is a survey of the various theories of hierarchies that have been offered by scientists and philosophers. I examine the hierarchy theories of Allen and Starr (1982), Roberto Poli (2001; 2004; 2006), Herbert Simon (1962; 1973), Burton Voorhees (1983), Mario Bunge (1969), and Richard Dawkins (1976), dividing them into three general categories: idiosyncratic theories, partial order theories and strict partial order theories. The difference between the latter two is a subtle one, but it is very important. I show that the partial order view faces a class of counterexamples. Given a choice between the two theories, the strict partial order view is favorable, but

nevertheless, it is under-motivated. In the literature, the view is most often stipulated, rather than developed.

Chapter 3: The third chapter then expands on the second. The goal is to develop the strict partial order view of hierarchies from humble beginnings, instead of introducing it as a stipulation. The development culminates into a general account of hierarchies and hierarchical structure called *Core Hierarchy Theory*. I argue that a hierarchy is a mereological collection of things, related by some strict partial order (to which we give the label “higher than”), that meets a special *linkage* condition. In developing this account, I argue both that being a strict partial order is necessary for a relation to count as a *higher than* relation; and that being a strict partial order is sufficient for a relation to count as *higher than*. With the complete account in hand, I then use it to argue for Realism. In the appendices to this chapter, I formalize Core Hierarchy Theory and prove some theorems that are relevant to the debate between Reductionism and Holism.

Chapter 4: Here I review some well-known accounts of a hierarchical worldview from the Reductionist perspective (Gillett 2007; c.f., Wimsatt 1976, 1994, and 2007). These accounts of hierarchical reduction fit within my own characterization of Reductionism found above. Using the results obtained from Core Hierarchy Theory, I then argue against Reductionism: either Core Hierarchy Theory is unacceptable, Reductionism is false, or the whole debate is ill-founded.

Chapter 5: Having argued that we have reason to take the Holist worldview seriously, in this chapter I review Holist (i.e., non-reductive) accounts of the layered worldview in science more generally. I conclude by arguing against Metaphysical-MLE as a general

principle: the ecological hierarchy modeled by Core Hierarchy Theory, despite satisfying Realism, has properties which appear to undermine the metaphysical justification for multi-level explanation in ecology. There are some hierarchies (e.g., mechanistic hierarchies) which seem to provide a metaphysical justification for multi-level explanation, but the ecological hierarchy described in the first chapter is not one of them.

Chapter 6: In the sixth and final chapter, I add the finishing touches to Multi-Grade Holism. I begin by considering objections from a neglected third player, the eliminativist who seeks to challenge the reality of hierarchies. The challenge is this: since higher level things, and the realization relation that holds between them, do not make a causal difference to the world, they (and the hierarchy of them) therefore do not exist within the scientific domain. In response, I argue, first, that higher level events and the realization relation itself, do in fact make a difference to the world, and second, that even if they did not make a difference, the eliminativist argument still fails.

To summarize: What follows is a comprehensive investigation into the metaphysics of ecology. My aim is to engage in a metaphysical study that has consequences for the epistemology of ecology. Falling within the scope of my inquiry is any issue which has both a metaphysical component and an epistemological component, and is such that the two components connect up in a way that makes metaphysics relevant. To accomplish my aim, my strategy is to investigate the debate over Reductionism vs. Holism, as well as the disagreement with respect to hierarchical Realism. My tactic is to offer a deep investigation into a notion that is relevant to both issues: I offer a fully general, non-question-begging, theory of hierarchies and then I explore its consequences, Multi-Grade Holism among them.

Chapter 1

The Ecological Hierarchy

1.0	INTRODUCTION	24
2.0	THE ECOLOGICAL HIERARCHY	26
2.1	<i>HOW DEEP IS THE ECOLOGICAL HIERARCHY?</i>	27
2.2	<i>ORDERING THE HIERARCHY</i>	31
2.3	<i>NESTED VS. NON-NESTED HIERARCHIES</i>	34
2.4	<i>BEING “AT” A LEVEL IN THE ECOLOGICAL HIERARCHY</i>	38
3.0	CHALLENGES TO THE ECOLOGICAL HIERARCHY	42
3.1	<i>RICKLEFS CHALLENGE TO THE COMMUNITY</i>	43
3.1.1	<i>Disintegrating the Local-Community</i>	43
3.1.2	<i>Re-integrating the Local-Community</i>	47
3.2	<i>POTOCHNIK AND MCGILL VS. SCIENTIFIC HIERARCHIES</i>	50
3.2.1	<i>The Coherence Problem</i>	51
3.2.2	<i>The Nesting Defense</i>	53
3.2.3	<i>The Motivation Problem</i>	55
3.3	<i>THE SIGNIFICANCE OF COMPOSITIONAL HIERARCHIES</i>	62
3.4	<i>COMPOSITION VS. SCALE</i>	65
4.0	CONCLUSION.....	67

1.0 Introduction

Hierarchical structures are ubiquitous within ecology. A quick scan of any introductory textbook reveals a variety of different applications. Taxonomic hierarchies and trophic (food-chain) hierarchies are among the most commonly encountered. But there is another, more general, purpose for hierarchies within ecology: they help to identify and organize that which is relevant to the study of ecological phenomena. My main concern in this chapter is to investigate the most well-known hierarchy that serves this purpose: the so-called “ecological hierarchy”.

There are at least three ways in which the ecological hierarchy serves this purpose. First, it helps to differentiate some of the individual *subjects* of study within ecology; for example, populations vs. ecosystems. Second, it helps to differentiate between the different *fields* of study within ecology; for example, population dynamics vs. ecosystems science. And third, it is thought to help to account for the relationships between these subjects and fields; relationships of causal dependence and influence, for example.

Although hierarchies play a significant role in the structure of ecology, the ecological hierarchy remains a rather murky notion that is easily misinterpreted. My primary aim here is to give a clear account of the ecological hierarchy, and to defend its place in ecology from recent objections.

The chapter takes the following structure. The first section is devoted to developing a clear account of the ecological hierarchy, as conceived of and used by ecologists. I begin by offering two competing presentations of the ecological hierarchy: the atomic view, which takes atoms to be part of the ecological hierarchy; and the organism view, which takes organisms to be the fundamental units of the hierarchy. Arguing against the organism view, I proceed with the atomic view in mind.

Building on the arguments of Allen and Starr (1982), I then develop a precise account of what it is to be at a level of the ecological hierarchy, as well as what it is to be higher than something else in the hierarchy. All of this culminates in a clear and concise characterization of the ecological hierarchy itself, at the end of the first section.

In the second section, I describe and respond to two objections to the ecological hierarchy: one from Robert Ricklefs (2008b), and the other from Angela Potochnik and Brian McGill (forthcoming). In brief, Ricklefs argues that the standard conception of the *community* within ecology is incompatible with the standard conception of the ecological hierarchy, and Potochnik and McGill argue that the usage of hierarchies in ecology has been motivated by appealing to certain implications of hierarchies which are either implausible or obviously false. In response, I argue that the account of the ecological hierarchy given in the first section comes away from these challenges unscathed. Having defended the use of hierarchies within ecology, a deeper investigation into their metaphysics is then warranted.

2.0 The Ecological Hierarchy

A hierarchy, roughly put, is an arrangement of some things, where some of those things are *higher than* others. (Craver 2007, Ch. 5) These *things* are generally called “levels”, or are otherwise contained within levels. The most encompassing presentation of the ecological hierarchy places *atoms* at the lowest level (Sadava, 2008; Potochnik and McGill, forthcoming). Atoms come together to form *molecules*, which make up the next lowest level. Molecules come together to form *cells*, which come together to form *tissues*. Tissues come together to form *organs*, which form *organ systems* (e.g., the circulatory system), which in turn come together to form *organisms*: the individual members of a species and the unit of natural selection (Ricklefs, 2008a).

The next level of the hierarchy is the *population* level. Organisms of the same species that are living together come together to form a population (Ricklefs, 2008a): the

unit of evolution and population dynamics. Next, populations of different species come together to form a local *community*. A local community, as standardly conceived (c.f. Sadava, 2008 p. 8, Ricklefs, 2008b p. 742), is a collection of populations of different species found within the same specified area. At the community level, ecologists study interactions between populations (e.g., predation, parasitism, symbiosis, etc.) as well as biodiversity. At the next level, communities come together, along with the abiotic features of the environment, to form ecosystems. At this level, ecologists are concerned with large-scale energy exchanges and nutrient cycles. Lastly, ecosystems come together to form the earth's *bio-sphere*, at the topmost level.¹⁸

2.1 *How Deep is The Ecological Hierarchy?*

The height of the ecological hierarchy is not a major point of disagreement; the main disagreement involves its depth. Some versions of the hierarchy do not include anything below the organism level (Ricklefs 2008).¹⁹ Atoms, for example, might seem to have no place in an ecological hierarchy since there appears to be nothing remotely ecological about them. Such a restriction yields a hierarchy that reaches only from organisms to the biosphere. Ricklefs, for example, remarks that “The organism is the most fundamental unit in ecology.” (*ibid.* p. 3)

The remark is no surprise: ecology, as a field of study, is classically conceived as the study of the relationship between organisms and their environment, broadly

¹⁸ Although some earlier presentations of the ecological hierarchy, influenced by Alfred Tansley (2000), include the universe itself as a level, contemporary ecologists do not tend to do so. Contemporary presentations of the ecological hierarchy are only as high as the biosphere.

¹⁹ Cf. Pickett et al. (1994, p. 20), who include subcellular structures as a level between the cellular and molecular levels.

construed. From the perspective of a *community* ecologist, such as Ricklefs, organisms and groups of organisms are of primary and fundamental concern. Still, *ecosystem* ecologists have a different perspective. They consider neither organisms nor communities to be primary. Within ecosystem science, it is imprudent to ignore the *parts* of the organisms, as well as the environment in which they (and their parts) are situated.

We might then agree with Ricklefs that organisms are the fundamental objects of study within *community* ecology, but this should not be taken to imply that they are metaphysically, or explanatorily fundamental within ecology as a general discipline. Insofar as the ecological hierarchy is offered as a way to structure all that is relevant to ecology as a discipline, leaving atoms and molecules out of the *ecological* hierarchy suggests that they play no role in ecological explanations. But that is a mistake. Ecosystem ecologists study nutrient cycles, and to describe a nutrient cycle, one must consider interactions between molecular compounds – e.g., interactions involving nitrites or nitrates. A similar consideration holds for atoms as well. Atmospheric composition is important to ecologists studying anything from individual organisms to the biosphere, and is largely understood in terms of atomic proportions. Take O₂ and CO₂ for example: one cannot have an O₂ molecule without O atoms. Molecules and atoms, then, have an important explanatory role to play. Thus, if the domain of ecology stretches only from organisms to the biosphere, then important explanatory elements (e.g., atoms and molecules) are left outside of it.

Ricklefs might not see this as much of a problem. One might simply take the ecological hierarchy to bottom out at the organism level, and introduce the missing

explanatorily relevant features as something external to the ecological hierarchy.²⁰ There might even be advantages to such a qualification. For one, it presents a compelling way to distinguish ecology from the general study of biology, yet at the same time brings the two fields together by virtue of their overlapping concern for organisms and populations.²¹ While the global unity of science movement (c.f. Oppenheim and Putnam, 1958; Fodor 1974) is generally seen as a failure, a local unity of the “life sciences” is then made to seem quite plausible under this hierarchical approach. That some ecological explanations require appeals to molecular-biology and chemistry, and as such require an appeal to features outside the primary concern of ecology, only serves to enhance this hierarchical unity: ecology takes as explanans what biology takes as explanandum.

Unification, I take it, is a desirable feature; and so, such an advantage should indeed arouse interest. But nonetheless, it is no reason to think that the ecological hierarchy completely ceases its downward progression at the organism level. Just because the biological and ecological explanandum hierarchies can be distinguished from one another does not imply that they are not in fact integrated into a single hierarchy that contains everything relevant to the production and explanation of ecological phenomenon. To hierarchically organize only the explanandum of ecology is to represent only a small part of the interest within the discipline taken as a whole. Ecologists also

²⁰ Alternatively, one might introduce atomic, molecular, and molecular-biological properties at the levels in which they seem most appropriate in ecological explanations. Atoms, for example, might make an appearance no lower than the ecosystem level, and molecular-biological properties might appear no lower than the organism level. The general idea is that we can introduce atoms and molecular properties at the levels which include the phenomena these entities help to explain. Although this is indeed one way to characterize an ecological hierarchy, it is not the standard view in ecology. The standard view, as we will soon see, is ordered by composition. Atoms are parts of organisms and everything else in the ecological hierarchy; and so, on the standard view, appear lower than organisms, and everything else, in the ecological hierarchy.

²¹ Population genetics, for example, is obviously concerned with populations.

care very much about how these phenomena of interest come about, which implies that they are interested in things found below the organism level in the integrated hierarchy. To model the structure of ecology as a unified field (or as a collection of disunified sub-fields), one must be sure to include all of the relevant interests, not only those concerns taken to be primary by a particular sub-discipline of ecology (e.g., community ecology, population ecology, etc.). Atomic, chemical, and biological predicates are all found within ecological literature, and not as idle footnotes. I take it, then, that Ricklefs' remark regarding the fundamentality of organisms is not in reference to the hierarchy of ecological interests, but rather only to the top-portion of it: the hierarchy of ecological explananda.²²

Even with the explanandum restriction, there is still a problem. Atomic, chemical, and biological predicates can also appear within the ecological explanandum. Explaining the declining population of a certain species of fish, for example, might involve a description of specific chemical properties found in the area, but explaining how the chemical properties came to be in the area to begin with is something else that needs to be explained. Are chemists and molecular biologists called upon to investigate such an issue? Perhaps they are, but if so, then they are going beyond their normal duties. The issue is plainly one for the ecologist. Chemical properties, then, also count as ecological explanandum. The hierarchy that Ricklefs is describing therefore cannot be an explanandum hierarchy. While there might be another way to characterize Ricklefs' vision, it is tangential to our present concern. Our investigation will proceed with the

²² Or perhaps Ricklefs only means to give the hierarchy of interests relevant to *community* ecologists, such as himself. I can only speculate on his motivation.

complete hierarchy in mind, extending from the biosphere all the way down to the atomic level.²³

2.2 *Ordering the Hierarchy*

Leaving aside the issue of depth, let us now turn to the issue of ordering. In general, there can be no hierarchy without some sort of ordering of its levels and contents. According to the standard conception of the ecological hierarchy, it is this notion of “coming together” – i.e., a compositional relationship – that gives it depth. The main alternative to a compositional ordering is one based on scale (Potochnik and McGill, forthcoming; McGivern 2010).²⁴

Even if we take a compositional ordering for granted, the hierarchy’s contents remain unclear. There is still the task of specifying what is being hierarchically ordered via this “coming together.” There are a number of plausible options. We might take the hierarchy to be one of objects, processes, events, properties, types, kinds or we might think of it as a hierarchy which integrates all of these things into a single hierarchy. Despite the numerous options, some of them share similar features which allow us to form two general categories of contents: I will call them *individuals*, and *instanceables*. With this categorization in mind, the options are narrowed down to three: the individual view, the instanceable view, and the integrated view.

²³ Notice that the sub-atomic level is not part of the ecological hierarchy. As far as I know, sub-atomic particles lay beyond its scope. They appear in no existing explanation of any ecological phenomenon. That, of course, might change. A quantum influence on genetic mutation or on global climate warrants the addition of a sub-atomic level to the ecological hierarchy.

²⁴ Ecologists do not often acknowledge a distinction between these two methods (see, for example, Allen and Starr 1982, Allen and Ahl 1996, and Salthe 1985).

On the first view, the individual view, each level consists only of individuals: singular things that can satisfy predicates (e.g., objects, events, and perhaps processes, behaviors, and phenomena).²⁵ The properties of these individuals are not themselves associated with any level.²⁶

On the instanceable view, one thinks of the ecological hierarchy as an ordering of things which has individual instances (e.g. types, properties, kinds, and perhaps processes, behaviors, and phenomena).²⁷ Types, properties, and kinds, however, are not likely to be composed of anything: composition is generally taken to be a relation between particulars. To produce a compositional hierarchy of types, properties, or kinds it seems that one cannot avoid involving individuals in some regard: presumably, the levels are ordered by virtue of the fact that all particular instances of these instanceables can be decomposed into instances of lower level instanceables. Communities can be decomposed into atoms, for example; so community-relevant properties are higher than atomic properties. Despite the important role of the individual in the hierarchy of instanceables, the individual instances themselves are not found on any level of the type-hierarchy.

²⁵ Our usage of ‘individual’ then differs from David Hull’s (1980). For our purposes, the term ‘individual’ merely picks out something that falls within the range of our singular quantifiers (i.e., “there is an x”), and can satisfy predicates. Unlike Hull, I do not use the term in a substantive metaphysical way. He is asking a different set of question from what we are asking here.

²⁶ To be clear, I should note upfront that by expressions like ‘is included in’, ‘is in’, ‘is contained in’, and ‘is at’ I am talking about the same relationship: the relationship holding between hierarchies, levels, and the things they put into order.

²⁷ I list processes and behaviors as potential examples of both individuals and instanceables since I do not have a clear grasp of their metaphysics. On the one hand, it seems that there are types of processes one can engage in; e.g. making coffee. But on the other hand, it seems that there are token processes as well: my making of the coffee earlier this morning. The same sort of consideration holds for behaviors and phenomena. There is no need to resolve these issues here, however.

Neither of these views adequately represents the sort of hierarchy utilized by ecologists. The ecological hierarchy appears to include individuals and instanceables. The reason: if there are only bare individuals, and no properties, (or vice versa) found within the ecological hierarchy, then the hierarchy is utterly uninteresting to Ecology. Ecologists are not so much interested in things or properties *simpliciter*, but rather, the contribution things and properties make to explanation and understanding.

For example, to explain why a species of bat has a certain physical distribution across a large region, (a population level explanation) requires reference to more than an individual population of bats, but also to its properties and the properties of its parts as well: if an organism *thrives on the faces of cliffs*, then that helps to explain why the population is distributed as it is. To give an explanation at a level of the ecological hierarchy, then, it is more appropriate to associate both individuals (e.g., bat populations) and properties (e.g., thrives on cliffs) with the levels of the hierarchy: explanation involves more than bare individuals and more than bare instanceables.

Insofar as the ecological hierarchy is a hierarchy of ecological interests, the most accurate characterization of the ecological hierarchy is therefore the integrated view. On the integrated view, both individuals and instanceables are associated with levels. We can say that an instanceable is within a level if and only if it has instances within that level. But then, we might ask, what is it for an individual instance to be *in* a level of the ecological hierarchy?

2.3 *Nested vs. Non-Nested Hierarchies*

To answer this question, we must first specify what sort of hierarchy we are dealing with in ecology. Within ecology, there is a well-known distinction between *nested* and *non-nested* hierarchies. Non-nested hierarchies are thought to be those in which one level does *not* contain the lower (Allen and Starr 1982 p. 40; Ahl and Allen 1996 p. 107). Examples of non-nested hierarchies include most food chains, corporate hierarchies, and college rankings. Nested hierarchies, on the other hand, are thought to be those in which a level *does* contain the levels below it (Allen and Starr 1982 p. 38; Ahl and Allen 1996 p. 107-10). Such a relationship can be presented in terms of nested sets or circles. To present the ecological hierarchy as a nested hierarchy, we might use the following:

[Biosphere [Ecosystems [Communities [Populations [Organisms [...]]]]]]²⁸

Or we might use:

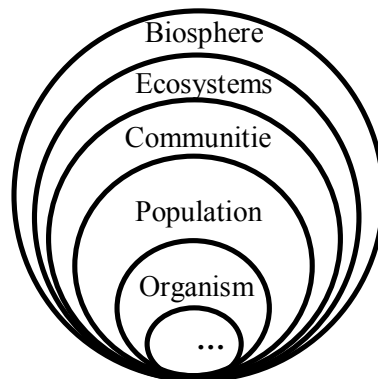


Figure 1: The Nested Hierarchy

Allen and Starr (1982) famously argue that compositional hierarchies involve such nesting. They ask us to consider military hierarchies as an example. Generals, they

²⁸ We should not take this to be an expression of set-theory. The top-level does not describe a two-membered set consisting in a set and the biosphere. Rather, it describes a series of nested subsets. This method of representing hierarchical nesting comes from Salthe 1985.

note, do not contain their subordinates, but armies do contain their battalions and soldiers. The former hierarchy, the command hierarchy, is paradigmatically non-nested; whereas the latter hierarchy, *qua* compositional, is marked as being nested. (1982 p. 39; also, see Salthe 1985, p. 10) Analogously, we are then to conclude, the ecological hierarchy, *qua* compositional, is a nested hierarchy.

Although it is never made explicit by Allen, Starr, or Salthe, the justification for nested compositional hierarchies seems to involve the transitivity of the *containment* relation. If a level x contains individual y, and individual y contains lower level individual z, then x contains z. The contents of lower levels in a compositional hierarchy are thus also contents of higher levels by virtue of being *contained* within the higher level individuals as their parts. Although one might question the transitivity of hierarchical containment, which is needed for their argument, there is, in fact, a growing body of psychological research which suggests that the everyday notion of hierarchical containment is indeed transitive (see Winer 1980; Greene 1989, 1994; Murphy 2002; Deneault and Ricard, 2006; and Slattery et al 2011).²⁹

But, of course, none of this suggests that a non-nested compositional hierarchy is impossible. To generate a non-nested compositional hierarchy, one simply constrains the containment relation in a way that invalidates its transitivity, yielding a refined notion of containment (e.g., x *directly*-contains y iff x contains y and there is no intermediary container z such that x contains z and z contains y). We can then add yet add another distinction to supplement the nested/non-nested dichotomy: a distinction between nested

²⁹ The primary question many of these psychologists seek to answer is at what age are hierarchical notions acquired within children. The answer, it seems, is between five and nine years of age. (Slattery et al 2011)

hierarchies of things, and hierarchies of nested things. The former are represented in a manner similar to Figure 1; whereas the latter would be represented by any hierarchy (perhaps non-nested) whose levels contain nested things. For example, one might reconstruct the compositional military hierarchy using a constrained notion of containment that *does not* imply that soldiers are contained within the army-level. Such a constraint transforms the exemplar nested hierarchy, into a non-nested one.

That the containment relation *can* be constrained is no reason to think that it *is* or *should be* constrained with respect to the ecological hierarchy. Unless the very notion of a nested hierarchy turns out to be incompatible with ecological ambitions, to add constraints to containment would only serve to add unnecessary complications. Assuming that there is no such incompatibility, the simplest and most general presentation of the ecological hierarchy therefore takes it to be nested.

In fact, there is a reason to think that ecological ambitions are *not* incompatible with nesting. The nested hierarchy allows for the inclusion of individuals and instanceables that do not fall within the specified level restrictions. Consider the place of antlers in the ecological hierarchy, for example. Antlers are bony branching protrusion from the skulls of mature males belonging to the Cervidae family.³⁰ Antlers are not ecosystems, communities, populations, organisms, organs, organ-systems, tissues, cells, molecules, or atoms. If level specifications were understood to imply restrictions to the sorts of things that are contained within each level of the ecological hierarchy, then antlers would fall outside the ecological hierarchy altogether. But that conflicts with one of the main purposes of the ecological hierarchy: to organize all that is of interest to

³⁰ Female caribou and reindeer, note, are exceptions. They too have antlers.

ecologists. Antlers are objects of interest to evolutionary ecologists, and are therefore to be included in the hierarchy of ecological interests. To include antlers in the ecological hierarchy one need only remove the restrictions on a levels contents. If the ecological hierarchy is nested, then antlers – *qua* composed of tissues – are appropriately positioned in the organ level, and all higher levels. This is not the only advantage of nesting. If the size of antlers found within a population, community, or ecosystem can be used as a metric for determining the “health” of an ecosystem, community, or population, then including antlers as contents of those higher levels is wholly appropriate. Since antler size is largely influenced by nutrition, their use as a metric is not at all implausible.

Nevertheless, these advantages of nesting might be overruled by conceptual difficulties for nested hierarchies. If the ecological hierarchy is nested, as Allen and Starr argue, then molecules will turn out to be within both the molecular level *and* the ecosystem level, by virtue of ecosystems being composed of molecules. Since the ecosystem level is taken to be higher than the molecular level, the fact that molecules are found within both these levels might be taken to suggest that molecules are higher than themselves. Taking for granted that the higher than relation is asymmetric (as evidenced by the psychological studies mentioned above), nothing can be higher than itself. So, it might be argued, nested hierarchies are deeply incoherent.

Fortunately for Allen and Starr, there is a simple response to the incoherence worry. They merely need to acknowledge a distinction between two relationships: the *higher than* relation, which holds between levels, and the *is at a higher level than* relation, which holds between the contents of levels. Drawing such a distinction allows them to agree to the asymmetry constraint when it comes to the ordering of levels, while

at the same time allowing them to deny it with respect to the ordering of their contents. Nestedness, then, can be retained without falling into incoherence.

2.4 *Being “at” a Level in the Ecological Hierarchy*

Having drawn the nested/non-nested distinction, and having argued for its relevance to the ecological hierarchy, any account of what it is to be contained within, or to be *at*, a level of the ecological hierarchy must take that distinction into account. Previously, we said that hierarchical levels contain both individuals and instanceables. An instanceable is contained within a level whenever it has individual instances at that level. All that remains, then, is an account of what it is for an *individual* to be at a level of the nested ecological hierarchy; an account that respects the nesting constraint, as discussed in the previous section.

Note that the purpose of the account is not to give a semantics of ‘x is at y’. We have already identified it as a sort of containment, inclusion, or (perhaps) parthood relation. Any of these options suffices for the present purpose. Here we only want an account of what it is to be at some given level of the nested ecological hierarchy we have already mentioned: an account which will allow us to determine, for anything found within the domain of ecological interest, which level or levels contain it. It is not enough to simply say that an individual, x, is at the community level if and only if x is a community. Since the ecological hierarchy is nested hierarchy, other things aside from communities are included in the community levels: organisms, for example, are associated with the community level *qua* nested.

The proposal that I have in mind is accomplished in three steps: first, begin with the domain of ecological interests; second, specify a label or name for each level, and rank them as the hierarchical analyst desires (viz., as a higher than ranking); and then, third, offer an algorithm or definition that each individual in the domain is placed into the appropriate level. Given the ordering of ecological levels we have so far – atoms (L_0), molecules (L_1), cells (L_2), tissues (L_3), organs (L_4), organ-systems (L_5), organisms (L_6), populations (L_7), communities (L_8), ecosystems (L_9), and the biosphere (L_{10}) – we can say that:

- (AT-LEVEL) Given any ecological level L_n , for all individuals x , x is at L_n iff either (1) or (2)
- (1) x is atomic³¹ and $n = 0$;
- (2) If x is a non-atomic individual, then for each of its proper parts p , there is some level L_z , (where $0 \leq z < n$) such that p is at L_z .

The account might appear to be circular, but in fact it is not.³² Recall this is not intended to be a reductive semantic or conceptual analysis of *x is at level y*, for any x and y . The ‘is at’ predicate is being taken for granted here as a way to give a more specific analysis, an analysis of what it is to be *at level L_n* of the ecological hierarchy, for any

³¹ Note that by ‘atomic’ I do not mean to refer to an object without proper-parts (i.e., an indivisible whole), but rather, the sorts of things that micro-physicists study.

³² One might also notice that gerrymandered sums of atoms, molecules, cells, etc., seem to be misplaced within this analysis. Note, however, that since our domain is restricted to individuals that are of *ecological interest*, the gerrymandered sums of atoms, molecules, cells, etc., fall outside the scope of the analysis.

chosen L_n from the list above (viz. L_0 - L_{10}). In other words, rather than giving an analysis of 'is at', we're giving an analysis of 'is at L_n ' where $0 \leq n \leq 10$.

What the AT-LEVEL analysis provides us with is a schematic framework to provide a series of analyses, one for each level of the ecological hierarchy L_0 - L_{10} . To arrive at this series, one merely drops the leading quantifier of AT-LEVEL and replaces each instance of the L_n -variable with one of L_0 - L_{10} from the list above. One such analysis in the series, for example: for all individuals x , x is at L_7 (the population level) iff x is composed of and only of individuals found within L_0 - L_6 (the atomic through organism levels). The analyses generated from the AT-LEVEL schema are thus informative and non-circular: on the left hand side of the 'iff' we state our interest in some particular level, and on the right hand side, we satisfy that interest.

The purpose of AT-LEVEL, recall, is to describe where each entity of ecological interest resides within the structure of nested levels. The analysis accomplishes its task: it provides, in principle, a way to associate each entity with the appropriate level(s) in a way that satisfies the nesting constraint. Atoms and atomic properties appear at the bottom level by virtue of (1), and also appear at all higher levels by virtue of (2), satisfying the nesting constraint.³³ Also by virtue of (2), all and only those individuals who are composed of atoms, and nothing else besides atoms, are found at the molecular levels. Keeping with the trend, all and only those individuals composed of things residing on the molecular level or lower, and nothing else besides things found on the molecular level or lower, is found on the cellular level and all levels higher than the cellular level.

³³ Atoms satisfy (2) since they always fail to satisfy its antecedent.

And so on. The process is repeated until one has pinned one's entity of interest to the first level one is able to, starting from the bottom as we do here. To put the point most simply: individuals found at any level L are also found at every level L^* higher than L , since the L -things are composed of things, and only of things, found at levels lower than L^* .

Although every lower level thing is contained within the higher levels, according to AT-LEVEL there are nevertheless some things that are exclusive to the higher levels. Biospheres, for example, are not found at the ecosystem level or lower: since biospheres are composed of ecosystems, they are not composed of and only of things found at levels lower than ecosystems. Similarly, molecules and higher level individuals are not found at the atomic level since they are not composed of anything lower than atoms; there are no such things in the domain of the eco-hierarchy. Lower levels are contained within higher levels, but higher levels are not (completely) contained within the lower levels. There is an imperfect overlap. The nesting constraint has thus been satisfied.

With this understanding of being at a level, we can then go on to account for what it is to be at a higher level than something else in the ecological hierarchy (as well as accounting for being at a lower level than) in a straightforward matter. Note that the following variables are intended to range over both individuals and instanceables.

(HIGHER-LEVEL) for all X and Y , X is higher than Y (i.e., Y is lower than X)
iff the least n_1 such that X is at L_{n_1} is greater than the least
 n_2 such that Y is at L_{n_2} .

There are, of course, other ways to specify what it is for an individual or property to be at a higher level than some other. HIGHER-LEVEL is useful in that it gives us a way to

understand what it is to be at a higher level that is compatible with both nested and non-nested hierarchies. In a nested hierarchy, cells are at a higher level than atoms: the lowest level at which cells can be found (L_2) has a greater rank than the lowest level at which atoms are found (L_0). The same consideration holds under any non-nested conception of the ecological hierarchy as well.

Our first major task is then complete. We have arrived at an account of the ecological hierarchy that is concise and complete. Let us henceforth take this account to be representative of the hierarchical structure of ecological science:

(**ECOHIERARCHY**) The ecological hierarchy =_{df} the hierarchy which (a) consists in all individuals and instanceables relevant to ecological interests; (b) spans from the atomic to the biospheric level in the following order (from bottom to top) – atoms (L_0), molecules (L_1), cells (L_2), tissues (L_3), organs (L_4), organ-systems (L_5), organisms (L_6), populations (L_7), communities (L_8), ecosystems (L_9), and the biosphere (L_{10}) – and (c) satisfies both **AT-LEVEL** and **HIGHER-LEVEL**.

3.0 Challenges to the Ecological Hierarchy

The compositional account of the ecological hierarchy has increasingly been met with scrutiny. In this section, I will present and respond to two challenges that have been raised to the compositional hierarchy. The first comes from Ricklefs (2008b) and the second comes from Potochnik and McGill (forthcoming).

3.1 *Ricklefs Challenge to the Community*

Ricklefs (2008b) argues that the popular conception of community, Gleason's (1926) "local" account, makes for an awkward fit within the ecological hierarchy, and has, perhaps, impeded progress within the discipline. His goal is to produce an alternative characterization of the community that achieves greater fit with the ecological hierarchy, and might even advance the field of ecology itself. I only focus on the hierarchical concern that arises from his challenge.

3.1.1 *Disintegrating the Local-Community*

On the standard view, we noted earlier, communities are thought to be composed of populations living within a specified area – i.e., a location in space and time. But 'specified area' and 'location in space and time' are imprecise. Areas can be specified in many different ways; and 'location' might mean either a region of space, or a single space-time point. The dominant precisification is Gleason's. Call it, the *local-community* concept: a (local) community is a collection of populations with overlapping distribution that co-occur at a point in space-time (Ricklefs 2008b).

In his critique, Ricklefs interprets Gleason's characterization in two different ways. First, he seems to interpret Gleason to mean that a local community consists in *a single space-time point*, surrounded by an assemblage of organisms from different species (2008b, p.744). Call this the *point* interpretation. As I understand the view, the space-time point is identified on a map, and the populations whose distributions overlap that point are then grouped together and identified as the community *relative* to that point.

Who would ever hold the point-view? Ecologists who find it useful to pin handwritten labels to points on maps might see fit to hold it; and so might those among those who regularly use Geographic Information Systems (GISs) to perform similar tasks. It is natural for such individuals to think of communities as mere conceptual tools, rather than as real things in the world.

On the second interpretation, Ricklefs takes Gleason to mean that a local community consists in something real: an *assemblage of species* found within an area surrounding a single point in space-time (*ibid* p. 742). Call this the *assemblage* interpretation. On this interpretation, communities are not relativized to space-time points. They are very real things that fall within a certain range, perhaps determined by the location of a pin on a map, or GIS system.

We can differentiate the two interpretations using the following diagram:

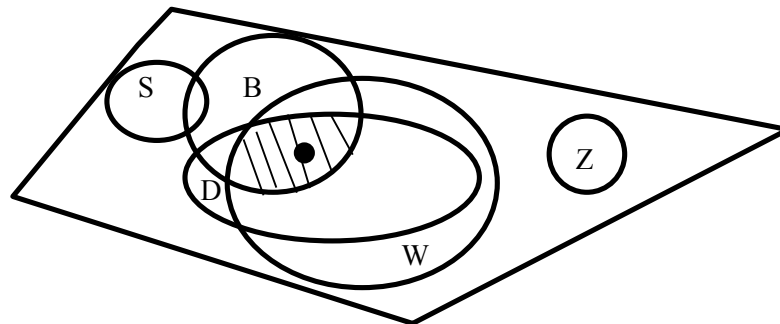


Figure 2: Ricklefs' Two Interpretations of the Local Community

Where the large polygon represents an aerial view of an environmental region with natural boundaries (say, an isolated valley); D represents the spatial distribution of a population of deer; W represents the distribution of a wolf population; B represents the distribution of black bears; S represents the distribution of swans surrounding a local marsh; and Z represents the distribution of zebras found in the local zoo.

On the point interpretation of Gleason, the community lacks spatial extent. It is nothing other than the dot found within the overlap of B, D, and W in Figure 2. On the assemblage interpretation, the community has spatial extent. It consists in the organisms found within the shaded region in Figure 2, or some arbitrarily specified sub-region.

No matter which interpretation ecologists generally have in mind, Ricklefs moves to reject the localized characterization in favor of the following alternative, call it, the *regional-community* concept: a (regional) community is a collection of populations falling within an environmental or geographic gradient; that is, it is a non-arbitrary region of space-time bounded by such things as natural barriers to dispersal and migration, geographic features such as watersheds, or perhaps bounded by features such as temperature, altitude, distance from water, and so on. To account for the difference between his own view and the Gleason standard, Ricklefs says the following:

The local community concept represents what I think of as a vertical perspective on species occurring within an arbitrarily bounded area, vertical because the species present local add together – pile up, so to speak – to form the local assemblage. The occurrence of species elsewhere within a region is of no consequence. In contrast, the regional community concept corresponds to a horizontal perspective on the distribution of populations over ecological and geographic gradients (2008b p. 744)

Figure 2 illuminates this distinction. The Gleason “local” standard takes the community, either as an assemblage of individual organisms found within shaded region (or sub-region thereof), or as some point within that region. Ricklefs preferred “regional” view of

the community can include everything in the valley: bear, deer, wolves, squirrels, and zebras. It depends on the gradient of particular interest.

Ricklefs mentions two reasons for “disintegrating”, as he puts it, the local community from the ecological hierarchy. Each corresponds to one of the two interpretations of community. The problem with both interpretations of the local approach, says Ricklefs, is that they do not fit within the standard compositionally ordered ecological hierarchy. (2008b p. 742) If he is correct, then there is a tension between ECOHIERARCHY and the community-concept: ECOHIERARCHY includes the community-concept, but if Ricklefs is correct, then the community-concept does not fit within ECOHIERARCHY.

With respect to the point interpretation of the community: populations, he argues, are spread out geographically, but space-time points are not. They have no spatial extent; and so, the “point community” (ibid p. 744) ceases to be a composite individual; indeed, he says, it is not an entity at all. (ibid.) Communities are nothing more than a point on a map to which we *assign* some group of organisms. Since the ecological hierarchy takes communities to be individuals, if the first interpretation is correct, then indeed the community has no place within ECOHIERARCHY.

With respect to the second interpretation: populations, he argues, are widely spread out geographically, but local assemblages of species are not. The shaded region (or its sub-regions) in Figure 2 occupy a space that is much smaller than the spatial extent of the associated populations of deer, wolves, and bears; and moreover, and the number of organisms found within the shaded region is much smaller than the number of

organisms in their respective populations (*Ibid.* p. 742). Thus, Ricklefs concludes, “a local community cannot be inclusive of the populations of its components species.” (*ibid*)

On both interpretations, then, it appears that the local community concept cannot fit within the ecological hierarchy. On both interpretations, the whole community is much smaller than the distribution and contents of its component populations. It is not literally composed of those populations. The local community concept, he concludes, must be disintegrated from the eco-hierarchy.

3.1.2 *Re-integrating the Local-Community*

There is, in fact, a third interpretation of the local-community that Ricklefs fails to consider, one that escapes the problems he raises for the eco-hierarchy. To replace the local community concept with a region-based conception is, then, a bit hasty.

The third interpretation can be called the *compositional* interpretation of the local-community. It can also be described using Figure 2 as an example. On the compositional interpretation, the local community does not just consist in those organisms found within the overlapping (i.e., shaded) region; rather, it consists in all of the organisms found within the populations that happen to overlap. With respect to Figure 2, we can interpret the local community as being the assemblage of everything within population D, within population W, and within population B. In other words, instead of taking the community to be a subset of the intersection of D, W, and B, we can instead take it to be the collection of organisms that fall within the *union* of D, W, and B.

On such a view, one might still see fit to add a regional boundary of some sort; but not one that makes a difference to the individuation of the community. It does not

matter whether or not an organism is found within any specified boundary for it to be considered a member of a community whose spatial distribution spatially overlaps that boundary, at least, not on the compositional interpretation of the local community. What really matters for the individuation of community, as with any other community, is that there is an area in which these populations have a chance to *interact*. The compositional interpretation described above satisfies this requirement.

That the interactive requirement has been satisfied is not obvious, by any means. It might seem problematic that those members of the populations who lie outside the overlapping area have no chance to interact with other members of other populations. But insofar as we are defending ECOHIERARCHY here, communities are composed of populations. Although not each and every member of these population interact, it nevertheless remains true that the populations themselves interact. Consider an analogy: one common ways for humans to interact is for their parts to touch – shaking hands is an obvious (albeit boring) example; I see no reason to think that populations are any different in this respect. Their members are their parts under the standard understanding of populations; since some of those parts interact, by our analogy, the populations interact. Of course, the stipulated area might be where the action is from the perspective of the working ecologist, but it by no means exhausts the membership of the community.

The compositional interpretation of the local community concept fits well with the ECOHIERARCHY analysis given earlier. Populations are composed of organisms grouped together by proximity (or the potential for mating, perhaps); and communities are composed of entire populations whose spatial distributions overlap, and who therefore have a chance for interaction and influence (e.g., predation and symbiosis). Ecosystems,

such as the valley ecosystem depicted in Figure 2, we can say, are composed of all the communities (all the overlapping distributions of populations) and abiotic features found within the region.

Although we should oppose Ricklefs' call to revise the community concept, there is no need to *reject* his conception of the regional-community. The compositional approach I propose above fits with Ricklefs' suggestion as well. But rather than help to individuate communities, we can instead use the regional-community to help integrate two levels: the community level and the ecosystem level. The ecosystemic level, recall, includes abiotic features as well as communities of populations as compositional elements. But which abiotic features does the ecosystem include? How do we decide? If Ricklefs' is correct, then the specification of the regional boundaries can be used to determine the abiotic components of the ecosystem. The ecosystem consists in the communities and abiotic elements that are found within some region that is used to pick out some regional community at the community level.

I do not mean to suggest that ecology *needs* a regional community concept, to prevent the community and ecosystem levels from dis-integrating. If we did not introduce regions at the community-level, we could simply introduce them at the ecosystem level; and so, the regional-community is inessential to hierarchical integration. But nevertheless, if Ricklefs is correct and the regional-community concept is useful within ecology, its integration-enhancing features suggest that it does have a place within the ecological hierarchy. I only argue that the regional community should not be thought to *replace* the existing understanding of the community. Ricklefs' contribution should thus be seen as an enrichment and extension of the community-concept. It should not be taken

to be a complete revision. When viewed in this way, Ricklefs' insights do not pose a challenge to ECOHIERARCHY.

3.2 *Potochnik and McGill vs. Scientific Hierarchies*

Whereas Ricklefs raises a very specific problem for the ecological hierarchy – namely, the community concept – Potochnik (a philosopher) and McGill (an ecologist) present a very general problem for the compositional ordering of the hierarchy. They argue that the ecological hierarchy, standardly conceived as a compositional hierarchy, is commonly ascribed problematic implications (§2 p. 9). They write:

In our view, the many overly ambitious conclusions drawn from the simple fact of part-whole composition – and the persistence of those conclusions – demonstrate that hierarchical stratification is not useful as a general conception of ecology or science. (*ibid*)

It appears that their case is this: since the compositionally ordered ecological hierarchy leads scientists and philosophers to draw bad conclusions and hasty generalizations, it should not be used as an organizing principle for the discipline and should be replaced with a different organizing structure: scale.

To evaluate Potochnik and McGill's case against hierarchical ecology, I begin by detailing the conception of hierarchy that they take to be standard in the biological sciences. I argue that the concept of hierarchy that they are arguing against is a departure from the aforementioned ECOHIERARCHY, and is, in fact, incoherent: their case against compositional hierarchies is a straw-man. I conclude the section by providing reasons to prefer ECOHIERARCHY over the scalar alternative they propose.

3.2.1 *The Coherence Problem*

Before attempting to undermine the motivation for the ecological hierarchy, Potochnik and McGill attack the conception of the hierarchy itself. According to ECOHIERARCHY, the organizing structure of ecology is a nested compositional hierarchy that consists of individuals and instanceables, and spans from the atomic level to the biosphere. Potochnik and McGill appear to agree with all of this (the inclusion of instanceables notwithstanding). They write,

The basic idea is that higher-level entities are composed of (and only of) lower-level entities, but the prevalent concept of hierarchical organization involves stronger claims as well. The compositional hierarchy is often taken to involve stratification into discrete and universal levels of organization. It is also often assumed that levels are nested, that is, that an entity at any level is composed of aggregate entities at the next lower level. (§1 p. 2)

Their agreement, however, is illusory. In fact, the characterization quoted above is incoherent. When one clarifies the standard usage of nestedness and discreteness, it becomes clear that levels are nested only if they are *not* discrete.

First, let us get clear on what Potochnik and McGill mean by ‘nested stratification’. They claim that nested stratification depends on “the uniformity of part-whole composition. For strata to emerge, atoms must always compose molecules, populations must always compose communities, etc.” (forthcoming, §2.1) But in fact, this cannot be a requirement of stratification. Two far-flung atoms do not compose a

molecule, and so atoms do not always compose molecules. I think what they mean by nested stratification is better expressed using what might be called *the principle of homogenous decomposition*:

(HOMO-DECOMP) For any entity e on level L , there are some entities “the x s” on level $L-1$, such that each of the x s is a proper part of e , and there is no proper part of e that is not a part of one of the x s.

If HOMO-DECOMP is correct, then everything on every level has an exhaustive decomposition into the next lower level.

On Potochnik and McGill’s understanding, *discreteness* appears to imply that the levels within the hierarchy do not have any entities in common. In other words, the levels of a hierarchy are discrete when and only when there is no overlap between them.³⁴ Atoms, on such a view, are *not* found at the molecular level and molecules are not found at the cellular level.

With such a characterization in mind, Potochnik and McGill argue against the ecological hierarchy by arguing against HOMO-DECOMP and similar principles. They draw on examples of objects at level n that are not composed exclusively of objects at

³⁴ I offer this interpretation out of charity. Although they eventually give a precise statement of what discreteness amounts to: “that an object taken to be at some level n is composed of all and only parts at level $n - 1$.” (§2.1), nobody actually holds this view. It implies something plainly absurd: that if I am composed of organ-parts, then I am not composed of atoms, molecules, cells, or tissues. I know of no ecologist who holds that organisms are *not* composed of molecules. If this is the understanding of discreteness they mean to argue against, then they are wasting their time. Their argument is a straw-man. HOMO-DECOMP is a far more accurate way to present the actual compositional intuitions employed within ecology. See Kim (2002) for a discussion of this, and similar principles.

level n-1. Kim calls these sorts of problems *free-molecule* problems (Kim 2002; also see Guttman 1976). For example,

ecosystems are said to be composed of... communities, but individual molecules, such as molecules of food waste, are also an important component. Tissues are only partly composed of cells; also crucial are the macromolecules that hold the cells together. (Potochnik and McGill, forthcoming, §2.1)

These molecules, Potochnik and McGill presume, are parts of the ecosystem (or tissue), but are not part of any entity at the community level. Since (their vision of) the ecological hierarchy implies HOMO-DECOMP, which succumbs to these sorts or counterexamples, they suggest that we have a strong reason to reject the ecological hierarchy.

3.2.2 *The Nesting Defense*

Within all of their criticisms Potochnik and McGill are using a characterization of the nested ecological hierarchy that is clearly at odds with the standard conception of nestedness discussed in the first section. Nestedness, recall, involves containment; specifically, the containment of the lower levels by the higher. Insofar as the army-level contains soldiers, the molecular level contains atoms, *pace* Potochnik and McGill.³⁵ The following diagram should help to illustrate nestedness:

³⁵ Note that Kim (2002) suggests a different solution to this problem, although Potochnik and McGill fail to consider it. He suggests that HOMO-DECOMP should be replaced with a weaker principle: one that requires each entity to decompose into entities, each found at the L-1 level *or lower*. Adding “or lower” means that entities need not exhaustively decompose into the L-1 level, and so the free-floating molecules no longer count as a counterexample. But this solution has its own implications some might find to be problematic. It entails that every higher level thing is also a lower level thing: e.g. ecosystems, qua having an exhaustive decomposition into the population level or lower, are thus at the population level. And qua

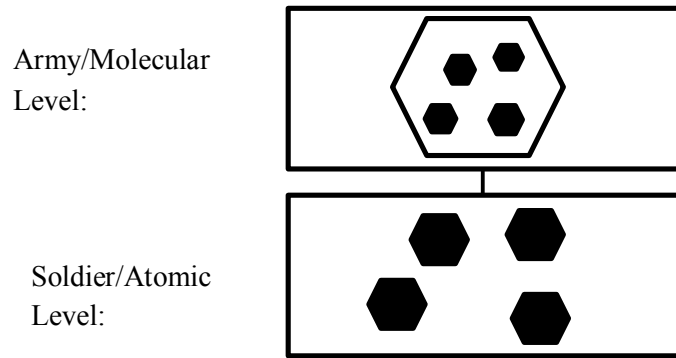


Figure 3: A Nested Hierarchy – the lower level is contained within the higher-level (c.f. Piero 2009, p. 23).

In fact, given the notion of nestedness accounted for within the AT-LEVEL analysis given in the previous section, HOMO-DECOMP is trivialized: since each level contains all of the lower levels, all of the proper parts of an entity, *e*, reside at the next level down; and so, *e* trivially possesses an exhaustive decomposition into entities at that level.³⁶

Potochnik and McGill’s mistake lies in their presumption that each level is discrete, that it consists only of entities characteristic of that level. For example, the ecosystem level contains only ecosystems, and the cellular level contains only cells. If Figure 3 is to serve as an illustration of the nested hierarchy, then that presumption must be rejected. And since Figure 3 appears to illustrate the notion of nestedness introduced by Allen and Starr (1982), and it is their work on hierarchies that introduces the nesting

having an exhaustive decomposition into the molecular level or lower, ecosystems are therefore at the molecular level. Reductionists should have no problem with this implication; but many ecologists are not reductionists. One of the purposes of hierarchical ecology, as presented by Allen and Starr (1982) and Allen and Ahl (1996) is to provide an alternative to reductionism.

³⁶ Does this mean that the laws of atomic physics must also be considered a part of ecology, if atoms are included? No. Such laws should be included only insofar as they are found within the domain of ecological interest to begin with. As far as I know, the laws of atomic physics do not play an active role in ecological explanations. They appear to reside beyond the scope of ecological interest, despite the fact that atomic predicates, and the individual atoms that satisfy them, do generally fall within the scope of ecology. Just because physics and ecology have overlapping concerns does not imply that all of physics gets subsumed within ecology.

concept to begin with, it appears that Potochnik and McGill are working with a conception of the ecological hierarchy that is not representative of what one finds within ecology.

ECOHIERARCHY, therefore, is not as hopeless as Potochnik and McGill purport; here we have shown them to be attacking an incoherent straw man. In the next section, we will examine the motivational facet of their argument against the use of compositional hierarchies within ecology.

3.2.3 The Motivation Problem

Potochnik and McGill's attack on the coherence of ECOHIERARCHY is only one of the problematic implications they investigate. They also consider five other implications of compositional hierarchies, which, they seem to think, are needed to motivate the use of hierarchies within ecology. They cite James Feibleman (1954) as the source of many of these implications.

The implications they cite include metaphysical claims and epistemic claims, providing two different dimensions along which hierarchical ecology can be motivated. Potochnik and McGill argue that these motivational implications are false. Without a reason to think that the ecological hierarchy is metaphysically or epistemically significant, they conclude that we have no reason to retain the hierarchical approach to ecology. Although they make a number of good points, I ultimately think that their challenges fail to undermine the motivation for hierarchical ecology. Let us address each motivating implication in turn.

The first purported implication involves a relationship between properties; specifically, an adherence to mereological supervenience.³⁷ Within a compositional hierarchy, higher level entities are thought to have properties not found on lower levels, since they are particular organizations of their lower level parts. But if that is all the entity is, an organization of its parts, then its properties are the properties of the organization of parts. This, it is often suggested, implies mereological supervenience. Call this *the supervenience implication*.

The problem with the supervenience implication, Potochnik and McGill argue, is that it is not universal. They cite examples of higher level properties that do not supervene on lower level properties. Camouflage, they point out, is a property of organisms, but whether or not a particular skin-color pattern counts as camouflage depends on more than the properties of the organism's parts. It depends on environmental factors as well, such as the sensory powers of predators. Moreover, they point out, phenotypic properties of organisms do not always supervene on genetic properties, even though phenotypic properties appear on a higher level than genotypic properties. Thus, the use of compositional hierarchies does not imply the supervenience of all properties on the properties of parts. And so, mereological supervenience cannot be used to motivate the metaphysical significance of hierarchies: mereological supervenience is not a general property of compositional hierarchies.

I take it that this first challenge is a good one. Whoever offers the supervenience implication as a motivating the metaphysical significance of hierarchies seems to be conflating compositional hierarchies and supervenience hierarchies. ECOHIERARCHY,

³⁷ I.e., the idea that the properties of the whole supervene on the properties of its parts.

recall from above, makes no mention of supervenience. Although an asymmetric supervenience relation might be considered hierarchical (see Kim 2002), there is no reason to think ECOHIERARCHY (or any other compositional hierarchy) is such a hierarchy. Just because some properties of the whole supervene on their parts does not directly imply that every property found on every level of a compositional hierarchy supervenes on the properties found within the next lower level. Potochnik and McGill are correct: since the ecological hierarchy is not mereologically supervenient, its being mereologically supervenient cannot be a reason to think the ecological hierarchy is metaphysically significant.

Potochnik and McGill's other challenges are not so convincing. Their second challenge pertains to some purported implications for complexity. According to Potochnik and McGill, compositional hierarchies are thought to be significant since they are thought to imply that complexity always increases as levels are ascended. They cite Jagers op Akkerhuis (2008) who says "the organization of nature is profoundly hierarchical, because from its beginning, interactions between simple elements have continuously created more complex systems, that themselves served as the basis for still more complex systems." (p. 2)

Potochnik and McGill argue against the complexity-implication by pointing out that ecosystems sometimes exhibit remarkably simple behaviors: the simplicity of certain predator-prey cycles serves as their main example. Since compositional hierarchies do not genuinely imply that complexity always increases along with hierarchical ascension, the increasing complexity of higher levels cannot be used as a reason to think that compositional hierarchies are metaphysically or epistemologically significant.

Given the nesting constraint on the ecological hierarchy, Potochnik and McGill's complaint is quite clearly incorrect. Although the ecosystemic level does contain some simple behaviors, the complexity found at lower levels is also found at higher levels, by virtue of our nesting constraint.³⁸ So, complexity does increase as levels are ascended. While this does not warrant the claim that *all* higher level systems are more complex than all lower level systems, this is not a claim a defender of hierarchical ecology needs to maintain.

The third purported implication involves rates of change. Potochnik and McGill claim that compositional hierarchies imply that rates of change always decrease as levels are ascended. To support this claim, they cite O'Neill et al. (1986), who argue that higher levels will change more slowly than lower levels since higher level changes are the result of a great many lower level interactions taken together.

The problem with this implication, Potochnik and McGill argue, is that there are many examples of lower level changes that are much slower than higher level changes. Changes to a genotype, they point out, occur very gradually, whereas changes to an individual organism's phenotypic traits, say, the amount of body hair, can change over the course of a few minutes of reckless shaving. Since compositional hierarchies do not imply that rates of change always slow along with hierarchical ascension, the slowing of rates of change as levels increase lends no credence to the claim that the ecological hierarchy is metaphysically or epistemologically significant.

³⁸ It is worth noting that Feibleman (1954 p. 59), the progenitor of the "implications" (he calls them *Laws*) to which Potochnik and McGill are objecting, explicitly includes the nesting constraint within his *first* law governing levels and complexity: on more than one occasion he remarks that the "levels are cumulative upward." (*ibid.* p. 60)

This argument is not convincing. Of course, it must be acknowledged that rates of change do not *always* slow as levels increase: the nesting constraint, recall, implies that rapidly changing lower level individuals are also found on higher levels as well. But this does not imply that a heuristic generalization cannot be made here. That is, Potochnik and McGill have not ruled out that, in a statistically significant number of cases, rates of change tend to become slower with an increase in level. Until they do, their challenge to the epistemic significance of the ecological hierarchy is not likely to gain much traction.

Indeed, Potochnik and McGill's phenotype/genotype counterexample might well be an isolated case. Upon discovering a genotypic change within a sub-population which give rise to fitness enhancing phenotypic traits (under the right developmental circumstances), the ecologist appears warranted in expecting a gradual change to the phenotypic traits of the general population. Although the generalization is not without exception, changes to the whole population do appear to come slower than changes to the individual organisms that compose the population. If we did not understand ecology as a compositional hierarchy, then it is unclear what would underwrite these expectations. Changes to the properties of the whole depend on changes to the properties of the part. And that there is such dependence strikes me as an evidential reason for thinking that the compositional ecological hierarchy is indeed metaphysically and epistemologically significant. The inferences we make about, for example, rates of change, appear to be informed by this metaphysical relationship of part-whole dependence.

The fourth purported implication involves the relationship between hierarchies, causation, and causal mechanisms. Compositional hierarchies, they suggest, imply that what goes on at a given level is accomplished by a mechanism at the level below. In

support of this, they cite Craver and Bechtel (2007) who argue that mechanisms play a central role in making sense of inter-level causal claims. According to Craver and Bechtel, levels of mechanisms are a species of compositional, or part-whole relationships, and causal relationships only occur intra-level (i.e., within a single mechanistic level).

To argue against this implication, Potochnik and McGill give examples of cases of inter-level causal processes; they argue that causation can be either upward or downward. To support the existence of so-called “downward” causation, they cite a growing consensus for a move away from “the traditional idea that [the population size of an organism] is determined locally through interactions with competitors and predators... and toward the idea that global abundance (abundance of a species across its entire range), shaped by processes such as the evolution of specialists vs. generalists, is a crucial determiner.” (forthcoming, §2.2) They also claim that one cannot explain holistic medical problems such as “metabolic syndrome,” (linking heart disease, diabetes, obesity, among other conditions) by looking at isolated molecules. These considerations, they claim, undermine the idea that change is always mediated by causal mechanisms at lower levels, and so, mechanistic mediation cannot provide a reason for thinking the ecological hierarchy is metaphysically or epistemologically significant.

I doubt that many will find these causal considerations compelling, mostly because the central example is rather unclear: upward and downward causation is an extremely contentious issue, and I am quite skeptical that anything can be inferred until the contention is resolved. Potochnik and McGill only hint at some possible examples of inter-level causation. And while it might turn out that they are correct, it is not at all clear that they are. It is possible that these are not genuine examples of inter-level causation.

And so, their examples are insufficient to conclude that this is indeed a problem for the significance of ECOHIERARCHY and compositional hierarchies in general. For now, how causation works within hierarchical contexts must remain an open question, and as such, cannot tell us anything about the metaphysical or epistemological significance of hierarchies.

Aside from these counterexamples, Potochnik and McGill also complain that compositional hierarchies are responsible for the widespread idea that lower level *theories* are more epistemically secure than higher levels, are better supported, and have a wider application than higher level theories. These purely epistemic implications, they argue, are also not genuine; they do not provide a reason to think that hierarchical ecology is epistemically significant.

I can understand why someone might think that epistemic security of lower levels is implied by hierarchies. Quite often we explain the behavior of higher level things by analyzing and describing the behaviors of their parts. Nothing Potochnik and McGill say gives us reason to doubt this practice. Since the parts reside on different compositional levels, compositional levels are then explanatorily significant.

Indeed, Potochnik and McGill's survey of the significance of hierarchies is hardly exhaustive. Aside from the complexity implication, the rate-of-change heuristic, and the prolificacy of downward-looking explanations, there might be other ways in which the metaphysical or epistemological significance might be maintained.

One major obstacle for Potochnik and McGill's general ambition is that they give no good account of what they mean by a hierarchy, compositional or otherwise. Without

some account of hierarchies in general, it is extremely difficult to figure out what sorts of claims actually are, and are not, genuine implications of compositionally ordered hierarchical structures.

3.3 *The Significance of Compositional Hierarchies*

Even without a clear account of hierarchies, there still are a few points to be made in favor of hierarchical ecology, to add to those listed in the preceding sub-section. The first is the simplest: understanding the structure of the world, or some part thereof, is often thought to be a common scientific and philosophical ambition.

If there is a reason to think the whole of ecology (or some part thereof) is structured hierarchically, then articulating that reason counts as an interesting contribution to ecological science in its own right; for lack of a better word, its significance is intrinsic. And there are such reasons. At the outset we characterized ecology as the study of the interactions between the individual and its environment, broadly conceived. Since individuals are parts of ecosystems, and ecosystems are an environment for the species that reside within them, ecology thus appears to carry a compositional burden by its very nature.

Another point in favor of ECOHIERARCHY is one we have already noted. A compositional ordering helps to unify ecology from a patchwork of disconnected sub-fields which ask different questions and employ different methods for answering them, to a singular inter-related picture of systemic interactions which integrates well with other hierarchical approaches to science. Even if the global unity of science movement has been abandoned, a local approach to ecological (and biological) unity might nevertheless

be facilitated by the hierarchical approach to ecology. In this regard, hierarchies are significant to ecological science.

Moreover, the compositional approach to hierarchical ecology is significant in that it captures and makes clear an important relationship between the individuals being ordered within the hierarchy: that, *ceteris paribus*, when something, say an ecosystem, undergoes change, it does not do so independently of changes to the communities and abiotic features that help to compose it.

I cannot help but to agree with Potochnik and McGill that the ecological hierarchy is not a supervenience hierarchy, where “being higher than” is interpreted as implying “supervenes on”. Nevertheless, changes to wholes generally imply some change to some of its parts. To use Potochnik and McGill’s example: the phenotypic properties of the platypus do not supervene on its genotypic properties. But still, they do supervene on the relative spatial locations of its parts. The platypus cannot have a duck-bill, unless it has parts configured into a duck-bill shape. That the compositional hierarchy captures this chain of change-dependencies strongly suggests that the compositional hierarchical approach to ecology is indeed metaphysically significant, at least insofar as composition is metaphysically significant.

Those resistant to ECOHIERARCHY will no doubt continue to discount both the metaphysical and epistemic significance of hierarchies. But I urge them to reconsider. The significance of change-dependencies is not only some idle metaphysical feature; it is an epistemically significant feature. Within community ecology there is a metric known as *beta-diversity*. Exactly how the metric should be understood is something of a

contentious matter (Tuomisto, 2010); but, roughly put, it is a measure of the difference in species diversity between ecosystems (or communities) or change of diversity within a single ecosystem (or community) over time. When concerned with the beta-diversity of a single ecosystem over time, the ecologist pays close attention to the breeding and dying (c.f. Ricklefs, 2008b; Whittaker, 2001), and perhaps migration patterns, of organisms within the ecosystem of interest.³⁹

The ecological hierarchy, qua compositional, allows for a generalization of such diversity at all levels: generally, we can say that the beta-diversity of something increases when there is an increase in the number of different kinds (or “species”) of its parts. For example, an ecosystem’s beta-diversity increases when the number of species in its component communities increases; whereas the beta-diversity of a community increases when the number of species composing the community increases; whereas the beta-diversity of a population, we can say, increases when the number of different types of individual organism increases, given some way of categorizing organisms of a single species (e.g., tall-neck giraffes vs. short-neck giraffes) that is relevant to the interests of ecology. On such a generalization, we could even talk about the relative beta-diversity of different molecules, if it was fitting to do so. In short, ECOHIERARCHY suggests that the familiar concept of beta-diversity can be made into a general metric that applies at all levels of composition, should we see fit to apply it. That strikes me as a significant implication of the standard compositional ordering of the ecological domain. One Potochnik and McGill do not consider.

³⁹ Although the beta-diversity metric is typically thought to apply to a region, this is simply a matter of prejudice. Beta-diversity can be described relative to any area of interest. The fact that the area of ecological interest is often regional is inessential to beta-diversity as a general metric.

3.4 *Composition vs. Scale*

Potochnik and McGill might retort that this beta-generalization feature is not unique to the ECOHIERARCHY. The beta-diversity metric is also commonly associated with a scalar (i.e., size-based) approach to ecology; and so, they might see this trend as supporting a scalar hierarchy for ecology insofar as it supports a compositional ordering. Indeed, they ultimately suggest that scalar orderings provide the right framework for ecology (c.f. Potochnik and McGill, forthcoming, and Rueger and McGivern 2010).

Although sympathy for scalar orderings is widespread throughout ecology, and such orderings are thought to be largely coextensive with compositional orderings, (even the composition-oriented Allen and Starr 1982, p. 271 claim that hierarchical levels can be understood as being associated with some scale), it is not at all clear to me that the scalar approach supports a generalization of beta-diversity.

Moreover, Potochnik and McGill do not address the main problem with pure scalar hierarchies: scales of interest can change depending on what is being studied. A population of bacteria, for example, can be much smaller than a single elephant. Similarly, individuals belonging to a certain species of fungi can grow to be kilometers in size. That is larger than some communities, and perhaps even some quaint ecosystems.

These vast discrepancies in scale are only one reason to complain about using scalar orderings as an organizing structure for ecology (and perhaps science in general). There are other problems as well. In general, scale is not useful in coming to a clear understanding of the different objects of study in ecology, such as ecosystems and communities. How does one specify scale such that it captures all and only the

ecosystems? What about the communities? That ecosystems are larger than communities does not help us understand what an ecosystem is, or what a community is. But notice, by understanding ecosystems as being composed of communities plus abiotic features, we have both a way of understanding what an ecosystem is, and a way of understanding how it relates to communities. Part-of is informative in ways that larger-than is not. By understanding ecological organization in terms of parthood, we gain some grasp of each thing being ordered. Moreover, such an understanding delivers a stable (i.e., an unchanging) hierarchical structure that ecologists can exploit when organizing and categorizing elements relevant to the discipline.

Besides, if ecologists find scale orderings useful, there is nothing stopping them from using them. There is no need to see scale and composition as being mutually exclusive approaches: parts, after all, are quite typically smaller than their wholes.

Before we continue, let us summarize what we have done in this section. We began by specifying what Potochik and McGill take to be the standard view of the ecological hierarchy, and we saw that, on the surface, it matches my own ECOHIERARCHY analysis. Upon deeper investigation, however, it turns out that their conception is incoherent. They claim that the levels in the hierarchy are both nested and discrete, which is impossible. Given the understanding of nesting introduced by ecologists such as Allen and Starr (1982), ECOHIERARCHY, qua nested, explicitly denies the discreteness of levels.

Next, we saw a list of five purported “implications” of hierarchies: hierarchies imply mereological supervenience; hierarchies imply that complexity increases as levels

are ascended; hierarchies imply that rates of change slow as levels are ascended; hierarchies imply that causation at a level is the result of causal mechanisms at lower levels; and hierarchies imply epistemic security for lower level theories and explanations. Potochnik and McGill's counterexamples show that some of these are not genuine implications. Nevertheless, we have other reasons to think that compositional hierarchies are significant: complexity does increase as ECOHIERARCHY's levels are ascended, rates of change typically do slow as levels are ascended, the structure of ecology is of intrinsic significance, compositional hierarchies adequately characterize an important set of change-dependencies. Not only is this characterization metaphysically significant, but it is epistemologically significant as well, since it permits a generalization of the beta-diversity metric to apply to all levels.

Lastly, I listed a number of drawbacks to the scalar approach Potochnik and McGill favor instead of compositional hierarchies. Scalar metrics are arbitrary, there are vast discrepancies of scale among organisms, and scalar orderings are of limited use when it comes to characterizing and individuating the subjects of interest within ecology.

4.0 Conclusion

The purpose of this chapter was to introduce the notion of an ecological hierarchy, to make the notion precise, and to defend the notion from some recent objections.

Hierarchies, we have seen, are used, and are useful, in theoretical ecology.

Hierarchies serve two main purposes within ecology. First, hierarchies provide an organizing structure for the discipline itself. By organizing the subjects of ecological study into levels, we have a way of understanding how various research programs within

ecology relate to one another. Each level is associated with its own subjects, theories, and research methods, and the hierarchical ordering of these levels corresponds to the way human beings perceive nature to be structured.⁴⁰ Second, the ecological hierarchy gives us a way to individuate the subjects of ecological study, making clearer what sorts of things are relevant to ecological phenomena, and what sorts of things we can expect to figure into ecological explanations.

This second purpose plays an important role in the ECOHIERARCHY account of the ecological hierarchy, as well as our analysis of what it is to be at a level of this hierarchy. I argued that the ecological hierarchy is a nested hierarchy ordered by composition: nested in that each level contains everything found in the levels below; compositionally ordered in that that everything found on any level except the bottommost level is composed of things found within the lower levels.

With this account of the ecological hierarchy in hand, I then sought to defend it from recent objections. First, we examined a criticism from Robert Ricklefs which suggests that the common “local” understanding of the community level fails to fit into the ECOHIERARCHY account featured above. His arguments are not convincing.

Second, we examined a series of criticisms from Angela Potochnik and Brian McGill. They argue that the ecological hierarchy is a “vexed” concept, and is riddled with difficulties. For one, they claim that the compositional structure of the ecological hierarchy, as standardly conceived by ecologists, exhibits an inner-conflict: they claim that there is no nested hierarchy with discrete levels. Although I agree with this claim, I argued that it is no slight against ECOHIERARCHY. The discreteness presumption is

⁴⁰ And perhaps to the way nature really is structured.

simply false; discreteness is not part of the conception of the ecological hierarchy as a nested. Some hierarchies are discrete; but ECOHIERARCHY is not one of them. Their concern is therefore not a reason to reject hierarchical ecology.

Although the hierarchical approach to ecology is able to withstand the aforementioned challenges raised by Ricklefs, Potochnik, and McGill, there is still much left to accomplish in order to meet the ambitions laid out in the introductory chapter. We have not yet answered the question of hierarchical realism: is hierarchical structure is part of the structure of the world, or is it a human invention? We have not yet investigated the justification for the multi-level approach to explanation and analysis, nor have we taken any steps to investigate the Reductionism vs. Holism debate. And above all, we have not yet said exactly what a hierarchy is.

The remainder of this dissertation involves a close examination of these outstanding issues. We will begin with an investigation of the existing general theories of hierarchies.

Chapter 2

Theories of Hierarchies

1.0	INTRODUCTION.....	70
2.0	IDIOSYNCRATIC THEORIES OF HIERARCHIES	72
2.1	THE APPROACH FROM ECOLOGY.....	72
2.2	THE APPROACH FROM ONTOLOGY	74
2.2.1	<i>Levels: Objects or Categories?</i>	75
2.2.2	<i>Dissolving the Dilemma: Levels Can Be Both</i>	76
2.2.3	<i>Are Levels Parts of Hierarchies?</i>	77
2.2.4	<i>Levels of Containment</i>	80
2.2.5	<i>Levels, Elements, and Parts of Classes</i>	82
2.2.6	<i>Being “at” a Level</i>	91
3.0	THE PARTIAL ORDER VIEW	98
3.1	SIMON ON HIERARCHIES	99
3.2	WEBSTER ON HIERARCHIES	100
3.3	HIERARCHIES: SET THEORY VS. MEREOLGY	101
3.4	VOORHEES AND BUNGE ON HIERARCHIES	104
3.5	REJECTING THE PARTIAL ORDER ACCOUNT	107
4.0	THE STRICT PARTIAL ORDER VIEW	109
4.1	DAWKINS ON HIERARCHY.....	109
4.2	HIERARCHIES AS DIRECTED ACYCLIC GRAPHS.....	112
4.3	THE ISSS ON HIERARCHIES.....	113
5.0	A SUPERFICIAL ACCOUNT OF HIERARCHIES	115
6.0	CONCLUSION	116

1.0 Introduction

If anything is evident from the preceding chapter, it is that there are serious confusions when it comes to a general understanding of hierarchies. This is not at all surprising.

Fully developed theories of hierarchies qua hierarchies are hard to come by.

In this chapter, I review some attempts to formulate a general theory of hierarchies, grouping them into three categories: idiosyncratic accounts, partial order accounts, and *strict* partial order accounts. The idiosyncratic views come from by Allen and Starr, and Roberto Poli. Although I argue against the idiosyncratic views, Poli's approach is useful. I use it to argue that there is a tight connection between hierarchies and mereology. The partial order view is attributed to Herbert Simon, and upheld by Jackson Webster, Burton Voorhees, and Mario Bunge. I argue that the partial order view is too narrow to capture the sense of 'hierarchy' actively in use within science and philosophy.

The strict partial order views are attributed to Richard Dawkins, and the International Society for Systems Science (ISSS). This *strict* approach has intuitive merit, but it is not without its problems. Aside from being undermotivated, underdeveloped, and underexplored, it faces conceptual difficulties. Where the partial-approach is too narrow, the strict approach errs by being too general: it admits hierarchies that are clearly not hierarchies.

What we are left with, in the end, is only a superficial notion of what a hierarchy is. It is progress nonetheless. In the next chapter, I show how we can use the superficial analysis to develop a robust and well-supported general theory of hierarchies qua hierarchies: a theory that rises builds on the failings of the strict, partial, and idiosyncratic views.

2.0 Idiosyncratic Theories of Hierarchies

To arrive at an answer to the question of hierarchical Realism mentioned in the introduction, we first need an acceptable account of hierarchies and at least some of their properties. As we will see, the idiosyncratic approaches are unsuccessful. We will look at two idiosyncratic approaches: the Anti-realist “ecological” account from Allen and Starr, as well as the ontological approach of Roberto Poli.

2.1 *The Approach from Ecology*

Ecologists are not shy when it comes to discussions of hierarchies. Recall from the introduction that Timothy Allen, writing with Starr (1982) and later with Victoria Ahl (1996), is champion of the Anti-realist approach to hierarchies, whereas the biologist and ecologist Stanley Salthe (1985) advances a Realist approach.

Although Allen and Starr’s (1982) work on hierarchies is one of the most well-known in theoretical ecology, the analysis that they end up with is unclear. The only explicit definition they give comes early on:

We have defined a hierarchy to be a system... where holons with slow behavior are at the top while successively faster behaving holons occur lower in the hierarchy. (p. 37)

There are a number of reasons to reject this definition.⁴¹ First, not all hierarchies have anything to do with behaviors (consider the hierarchy of pure sets, for example). Second,

⁴¹ The term ‘holon’ comes from Koestler (1967). A *holon* is something that exhibits a mereological duality: it is something that is simultaneously an individual, and is also a part of another individual.

there is an explicit reference to process-rates (i.e., speed – faster and slower), when not all hierarchies are diachronic.⁴² Third, Allen and Starr do not take into account that hierarchies can be inverted – i.e., turned “upside down” – to produce a second hierarchy that ranks its levels in reverse order. For example, a hierarchy in which faster is higher and slower is lower. Call this, the *inversion principle*.

(Inversion Principle) For every hierarchy H, there is a hierarchy H* such that it ranks the levels of H in the reverse order.

The inversion principle is not unique to hierarchies. Given any relation R, (e.g., higher than) a converse relation R* can be defined. If R is a sort of ranking (e.g., a hierarchical ranking), then R* is simply that ranking, reversed.

Perhaps the most striking problem with Allen and Starr’s account comes from within ecology itself: it contains straightforward counterexamples. Consider the hierarchical relationship between communities and their constituent populations. Communities are at higher levels than populations, but community-holons (i.e., individual communities) do not always exhibit slower behavior than population-holons. Suppose, for whatever reason (say, climate change), that the number of invasive insect species in a locale sharply increases. Changes in the alpha-diversity of the local community (i.e., changes to the number of different species that compose the community) then occurs more quickly than the alpha-diversity of any given population (i.e., changes to the number of different sub-types within a population) within that community. We thus

⁴² These appear to be the sorts of claims that lead Potochnik and McGill to raise their concerns regarding the appeal to hierarchies within ecology (as discussed in the previous chapter). If their conception of hierarchies is informed by the Allen and Starr analysis, it is no wonder that they are skeptical.

have a tidy example in which the higher level is changing more rapidly than the lower, *pace* Allen and Starr's characterization.

Their characterization, therefore, suffers from a lack of sufficient generality. Not only does it conflict with the inversion principle, which follows from the standard understanding of ranking relations, it is also subject to straightforward counterexamples from within ecology.

2.2 *The Approach from Ontology*

Roberto Poli (2001, Gnoli and Poli 2004, and 2006) also takes seriously the idea of a metaphysical hierarchy. He sees a study of hierarchies as part of a study in ontology, on par with mereology. His approach differs from the ecologist's in that it shifts the central notion from hierarchy, to *level*.

Levels, for Poli, have something to do with relationships between categories. He tells us that a level is simply a collection (or "complex") of categories that is bound and directed in the appropriate way to other groups of categories. (2001; 2006; Gnoli and Poli 2004, p. 156)

Poli's work on levels raises an important issue: what is the relationship between hierarchies, levels, and mereology? Are levels parts of hierarchies? Are those things found within each level parts of the levels? Poli rejects this mereological approach, but I do not find Poli's case convincing. I offer a positive argument for thinking that levels are indeed parts of hierarchies.

2.2.1 Levels: Objects or Categories?

Although Poli is one of the few philosophers who voices complaint that there is no good definition of ‘level’ or ‘hierarchy’ on offer (2004), he does not offer a precise definition of these terms. He is more interested in drawing a distinction between the *categorical approach* to levels and the *object-approach*, and then defending the former.

The distinction is easily stated: Poli’s categorical approach takes levels to be categories, which he compares to the metaphysics of universals, rather than the strict mathematical sense of Category Theory. The object-approach, on the other hand, takes levels to be particular objects, constituted by a collection of things residing at that level. Some levels are of a higher ‘granularity’ than others on Poli’s characterization of the object view.

He argues against the object-approach in two ways: First, he claims it is an ideological legacy from positivism with no substantial support,⁴³ and then argues that since the categorical approach is more general, it can subsume the positivist picture and do much more besides: it can be used to model what he calls ‘tangled’ hierarchies,⁴⁴ as opposed to mere linear hierarchies (2001 p. 266).

Second, he argues that the object-approach is tied up with mereology – the theory of parts and wholes – and this “makes it impossible to discover whether a properly

⁴³ I must confess that I do not understand his reference to positivism.

⁴⁴ In mathematics, ‘tangled hierarchy’ is an idiom that refers to a sort of loop. The idiom is a substantial departure from the standard understanding of hierarchies. Higher than is generally taken to be either asymmetric or anti-symmetric; either way, a hierarchical ordering of levels is in no way loopy. What might be loopy, however, is the hierarchical ordering of the levels contents.

developed theory of levels, as distinct from the theory of wholes and their parts, has something to add to our understanding of reality.” (2004, p. 156)

Poli’s dismissal of mereology is an important aspect of his view. In his defense of the categorical approach, Poli claims that levels are not parts of objects. He writes,

In order to mark as clearly as possible the difference between the theory of wholes and the theory of levels, let us boldly claim that levels *are internal to items but not as their parts*... [this is] the main principle of the theory of levels (as different from the theory of wholes and their parts). Claiming that levels are not parts means that levels are not elements of entities.
(2004, p. 156)

Although the connection is not obvious, Poli concludes from this that the categorical approach is the best way to go. Presumably, he thinks that since levels are not parts of entities, but are, in some other sense, related to those entities, the only other alternative is to say that they are categories, which we are to understand as universals or classes.

2.2.2 Dissolving the Dilemma: Levels Can Be Both

Although Poli presents an interesting picture of the metaphysics of levels, his arguments are not very convincing. The dilemma between the categorical and object-based approaches is a false one: they are not mutually exclusive.

Consider the hierarchy of natural numbers. In such a hierarchy, we are clearly ranking individual numbers, rather than classes or universals. Although Poli might want

to say that we are really ranking unit-classes of numbers, there is no argument (that I know of) for that claim.

On the other hand, we should be compelled to agree with Poli that there are hierarchies in which the levels *are* categories. Let S be the category of ecosystems and let M be the category of molecules. Is there any sense denying that there can be a hierarchy in which M is higher than S? I think not. Levels can be categories. Levels can be objects. There is no dilemma for us to confront.

Indeed, it is evident that in a hierarchical analysis (metaphysical or otherwise) levels can be whatever you want. Whether levels are particulars, objects, categories, groups, sets, universals, or classes (including equivalence classes) depends on the intent and purpose of the hierarchical analyst, and what she sees as being appropriate to quantify over and to put into hierarchical order. When developing a *general* theory of hierarchies qua hierarchies, then, we must be sure to avoid favoring one horn of Poli's dilemma over the other.

2.2.3 *Are Levels Parts of Hierarchies?*

The relationship between hierarchies and mereology is less than clear. Poli's view on this matter, which we quoted in the previous section, is in direct opposition to that of Herbert Simon, whose ideas serve as a foundation for much of the subsequent discussion on hierarchies in the philosophy of ecology. Simon notes that by 'hierarchy', he means a certain sort of "partitioning in conjunction with the relations that hold among its *parts*." (1962, p. 2 n.6, emphasis added). In other words, levels, for Simon, are parts of

hierarchies that have a special *higher than* relationship holding between them. Poli, as we have seen, disagrees. He claims that levels are *not* parts of the things that have them.

One drawback to Poli's view is that he does not offer a positive account of the relation holding between levels and the hierarchies that have them; the relationship is left mysterious. This drawback is not as innocent as it might seem. It, in fact, leads to the following argument in favor of Simon's view.

Without claiming outright that levels are *parts* of hierarchies, let us take for granted that hierarchies consist (in some sense of 'consist' that does not imply parthood) only of ordered levels, the contents of those levels, and combinations thereof.⁴⁵ An account of hierarchies which takes hierarchies to consist in things other than what is listed here faces a more parsimonious competing account without these extra commitments.

With that in mind, let us now consider Poli's claim. If levels are not parts of hierarchies, what, then, *are* its parts? If hierarchies only have *improper* parts, then hierarchies are mereologically simple (i.e., atomic). But hierarchies are not mereologically simple. An organizational hierarchy for an international conglomerate, for example, is clearly quite complex: it is not uncommon for corporate and government mandates to break large organizations into smaller pieces. Upon doing so, the organizational hierarchy breaks into pieces as well. Insofar as pieces are parts, and I take it for granted that they are, hierarchies are then not mereologically simple. Hierarchies, therefore, have proper parts.

⁴⁵ That is, combinations of levels, combinations of contents, and perhaps combinations of levels and contents.

Working under the assumption that levels are not parts of hierarchies, and that hierarchies consist only in levels and their contents (or combinations thereof), it follows that hierarchies either have the contents of levels as parts, or combinations thereof. If it were to have any other parts aside from these, then hierarchies would consist in more than what we would expect. Parsimony does not favor such a suggestion.

Under the present assumption, note that combinations of levels cannot be considered parts of hierarchies: if combinations of levels are parts of hierarchies, then so are levels: since combinations of levels are parts of hierarchies, and individual levels are parts of combinations of levels, it follows from the transitivity of parthood that levels are parts of hierarchies, which runs contrary to our assumption. So, the proper parts of hierarchies, under our assumption, must either be the contents of levels or combinations of those contents (or both).

Notice that the same sort of reasoning holds with respect to combinations of contents; if combinations of contents are parts of hierarchies, and the contents of levels are part of the combinations of those contents, then by the transitivity of parthood, we can straightforwardly infer that the contents of levels are parts of hierarchies. Therefore (closing our open assumption), if levels are not parts of hierarchies, then the contents of those levels are parts of hierarchies.

But here we arrive at a reason to reject Poli's claim in favor of Simon's. If those contents, $x_1 \dots x_n$ of a hierarchy's levels are appropriately combined such that *those combinations* can be arranged as being higher than each other, then there is no point in positing Poli's metaphysically mysterious "levels" to begin with: all reference to levels

can be replaced with reference to their combined contents. In other words, to say that levels are *not* parts of hierarchies makes the appeal to levels in hierarchies superfluous: if levels are not parts of hierarchies, then an ordering of their collected contents, which *are* proper parts of the hierarchy, suffices to produce a hierarchical arrangement without positing any of what Poli calls “levels.”

Thus, insofar as one posits levels at all, one should take them to be nothing more than combinations of their contents. To suggest otherwise is to forsake parsimony; it is to add extra metaphysical or theoretical content when it is not called for. An adequate theory of levels and hierarchies, therefore, takes levels to be parts of hierarchies, insofar as one mentions *levels* at all.⁴⁶

2.2.4 *Levels of Containment*

The argument from parsimony might raise worries about one of the central arguments from the preceding chapter. ECOHIERARCHY, recall, was formulated and defended using the “nesting” argument from Allen and Starr. That argument appears to presume that the correct way of thinking about the relationship between hierarchies, their levels, and the contents of those levels involves the notion of *containment*: hierarchies contain levels which contain other things. The argument from parsimony implies otherwise: it implies that the relationship is one of parthood, and containment is not parthood. A toolbox, for example, contains tools, but is not composed of them.

⁴⁶ Note, however, that we do not want to say that levels must be collections (or combinations) of their contents. We have been understanding collections as those singular things with proper parts: those things that are being collected. To suggest that levels *must* be collection thus rules out the possibility of a hierarchy of mereological atoms, ordered by velocity, acceleration, distance-from-my-left-foot, or some other property. The possibility of such hierarchies should not be ruled out a priori: levels can be collections or they can be atoms. An acceptable account of hierarchies allows them to be both.

The intuition that hierarchies contain their levels, which, in turn, contain their contents is easily retained under the mereological approach. Although containment is not generally a sort of parthood, parthood is nevertheless a sort of containment. Allen and Starr's argument, recall, relies on the coherence of 'contains-as-part', which is not something we have good reason to question. The worry, then, is not serious.⁴⁷

That there is an implicative difference between containment and parthood might be used to raise a different problem for the parsimony argument. If there is a difference between containment and parthood, and the two approaches to understanding the relationship between hierarchies, levels, and their contents, are equally parsimonious, then it might appear that we have no reason to favor the mereological approach over the containment approach. Poli's opposition to Simon might then be seen as a disagreement over whether levels are containers or whether levels are parts. Parsimony, one might argue, favors neither, for it is a virtue possessed by *both* approaches.

Here I protest. The mereological "levels-as-parts" approach is more parsimonious than the pure containment view. If hierarchies contain levels which contain their contents, but those levels and those contents are not parts of hierarchies or levels, there must then be a boundary which separates the contained from the uncontained. If the levels-as-containers approach were to provide a foundation for a general account of hierarchies, then a general account of hierarchical boundaries must be provided. Without such an

⁴⁷ Allen and Starr's argument is easily restated in mereological terms. Given any compositional hierarchy (e.g., the ecological hierarchy), if L is a level of that hierarchy, and x is part of that level, then any part of x is also a part of L. In a compositional hierarchy where wholes are ranked higher than their parts, the parts of lower levels are thus also parts of higher levels, assuming parthood is transitive. Nestedness, then, is a feature of compositional hierarchies under both the levels-as-parts view and the levels-as-containers approaches.

account, the levels-as-containers approach would remain incomplete. Try as I might, I cannot make good sense of such a boundary. But even if I could, the account of these hierarchical boundaries would be a needless complication. Insofar as we are committed to parthood, that is all we need for a robust theory of the relationship between hierarchies, levels, and their contents. There is no need to posit hierarchical boundaries on the mereological approach; what separates the contained from the uncontained is that the former are parts, whereas the latter are not.

2.2.5 *Levels, Elements, and Parts of Classes*

Another way to defend Poli from the parsimony-based argument is to suggest that there is another sense of ‘consist’ and ‘combine’ that has nothing to do with parthood. In many philosophical circles, it has become standard to think that there are two entirely different ways to make one thing out of many: the one can be mereologically *composed* of the many, in the sense that it has those many things as its *parts*, or, the one can be a *set* (or class) of the many, in the sense that it has the many as its *members*. Membership, according to the initiates of these circles, is not parthood. David Lewis (1991), for example, uses the distinction as a basis from which to argue that the parts of a set are all and only its subsets.

Perhaps, then, the relationship between hierarchies and their levels is one of set-membership rather than parthood. If that is indeed what Poli has in mind,⁴⁸ then he has a reply to the above parsimony-based argument: organizational hierarchies are indeed

⁴⁸ I do not think this is what Poli has in mind; for he also claims that levels are not elements. Elementhood, as it is commonly understood, is nothing other than the set-membership relation. Even so, the set-theoretic suggestion is one worth considering for independent reasons.

complex, and they might indeed have parts, but that does not mean that *levels* are their parts. If hierarchies are sets or classes, then the charge of mereological atomicity can be avoided by pointing out that the corporate hierarchy does have proper-parts: the subsets of its levels. Yet, those subsets are not themselves (necessarily) levels of the hierarchy. Levels are organizational departments; they are not (necessarily) sets of departments.

Of course, such a response is convincing only insofar as we can agree that the members of a set are, in general, *not* to be considered parts of it. Outside of the aforementioned philosophical circles, it is common for the uninitiated to intuit that members are indeed parts of sets. It is doubtful that many speakers of English would object to the claim that the set of all American presidents has Richard Nixon as part, or that fifty-two playing cards strewn about the floor are all parts of the set of those playing cards.

Of course, these are only the intuitions of the uninitiated. Intuitions are not arguments. In fact, it is a well-known line of argument that separates the initiated from the uninitiated to begin with. Its most elegant presentation comes from David Lewis:

A member of a member of something is not, in general, a member of it; whereas a part of a part of something is always a part of it. Therefore, we learn not to identify membership with the relation of part to whole. (1991, p. 3)

Such simple arguments are a treat for philosophers. But in this case the simplicity is deceptive. The argument is invalid.

The problem is in the definite description “*the* relationship of part to whole”; Lewis neglects to acknowledge that the part to whole relationship comes in different varieties: direct-parts, functional-parts, and spatial-parts are all sorts of part to whole relations, and all of them are easily defined using the same primitive mereological notion of parthood. Not all of these varieties possess the same formal properties of primitive-parthood, however. Consider direct-parthood: x is a direct-part of y iff x is a part of y and there is no z such that x is a proper part of z and z is a proper part of y . Clearly, direct-parthood is intransitive: if x is a direct-part of z , and z is a direct-part of y , then it follows from the definition of direct-parthood that x is *not* a direct-part of y . Nevertheless, direct-parthood is a sort of part to whole relation, despite its intransitivity.

Let us then consider an analogue of Lewis’ argument, with ‘direct-part’ taking the place of set-membership: a direct-part of a direct-part of something is not, in general, a direct-part of it; whereas a part of a part of something is always part of it. Therefore, we learn not to identify direct-parthood with *the* relation of part to whole.

Clearly *this* argument is invalid. The premises are as evident as in Lewis’ original argument, but the conclusion – there is exactly one relation of part to whole and it is not identical to direct-parthood – is plainly false. Direct-parthood is surely *a* relation of part to whole; one such relation among others.⁴⁹ Since Lewis’ argument shares the same logical form as the invalid argument (membership simply takes the place of direct-parthood), it too is therefore invalid.

⁴⁹ Some might see this claim as being inconsistent with our endorsement of mereological monism in the introductory chapter. Rest assured, for there is no inconsistency here. Mereological monism does not forbid other, more restricted, varieties of parthood, i.e., ϕ -parthood; it only requires that these ϕ -parts obtain their part-ness by virtue of being parts in the unrestricted sense captured by classical mereology (or whatever mereology one endorses) as well.

Where does this leave Poli and Lewis? It leaves them without a good reason to think that members of sets are not parts of sets.⁵⁰ In Lewis' case, this undermines the motivation for his main conclusion (viz., that the parts of a class are all and only its subclasses) by undermining his justification for his first premise:

First Thesis: One class is a part of another iff the first is a subclass of the second.
(1991, p. 4)

Although Lewis is able to convincingly argue for one direction of this biconditional – that if S is a subclass of C , then S is a part of C – he does little to motivate the other direction. This is unfortunate. The argumentative gap, as slight as it is, opens a way to undermine Lewis' project before it begins. Where C and D are (non-null) classes, and where $D = \{C, 1\}$, and $C \neq \{1\}$, if it turns out that members are parts of sets, then we have a counterexample to the First Thesis: we have a case in which one class, C , is a member-part of D , but is not a subclass of D . The only argument Lewis (1991) offers against this sort of counterexample is his transitivity-argument against the idea that members are parts of sets. That argument is invalid. If there are positive reasons for thinking that members are parts of sets, then we have reason to reject his main thesis, undermining much of his mereologically grounded set-theory.

⁵⁰ Lewis' second reason for thinking that a set's members are not its parts is undermined by the same counterexample as his transitivity argument. The second argument (1991, p. 5) is that a whole can be exhaustively divided into parts in many different ways (e.g., the sum of all cats can be exhaustively divided into cats, cat-parts, or microphysical particles), whereas a set does not divide exhaustively into its members in the same number of ways; so, members, in general, are not parts of sets (whereas subsets are parts, Lewis suggests, since their exhaustive division matches that of the sum). But notice, there is a mismatch between the number of ways a whole can be exhaustively divided into its parts and the number of ways it can be exhaustively divided into its *direct* parts. There are far less of the latter than there are the former. Since we cannot validly conclude from this mismatch that direct-parts are not parts, we cannot validly conclude that members are not parts using Lewis' analogous argument.

There are at least two reasons for thinking that members are, in fact, parts of sets; the first is, ironically, an analogue of Lewis' argument for taking subsets to be parts of sets: it conforms to common speech. As Lewis writes, it seems

natural to say that the class of women is part of the class of human beings, the class of even numbers is part of the class of natural numbers, and so on. Likewise it seems natural to say that a hyperbola has two separate parts – and not to take that back when we go on to say that the hyperbola is a class of x-y pairs... a standard word for 'subset' is "*Teilmenge*", literally, 'part-set'. (1991, pp. 4-5)

He notes that there are two explanations for this conformity to common speech. First, the straightforward explanation: we speak as if subsets are parts of sets because they *really are* parts of sets. Second, the "devious" explanation: we speak in this way because there is an *analogy* of formal character between the subset relation and the primitive-parthood relation; that is, they have the same formal properties. Although Lewis does not say it, we can rule out the devious explanation. The idea that subsets are parts of sets is intuitive *without* considering the respective formal characters of parthood and membership. The intuition arises straightforwardly upon realizing that combining (or taking the union of) the set $\{1,2\}$ with the set $\{3\}$ gives us the set $\{1,2,3\}$, and that $\{1,2\}$ and $\{3\}$ are both subsets of $\{1,2,3\}$. Learning these two basic facts about set-theory is enough for anyone to intuit that subsets are parts of sets, independent of any considerations of formal character.

In fact, Lewis' argument can be used against him: an analogous consideration holds for membership. It is common to hear people say that a member of a committee is part of the committee, or a member of a species is part of that species, despite that the usage of 'part' in this case is intransitive. Hands are parts of committee members and species members, but hands are not often taken to be parts of committees or species. Such usage suggests that membership implies parthood, even if it is not the same as primitive-parthood. Membership, then, is simply a special type of parthood: a restricted type of ϕ -parthood.

That these cases of membership do not make explicit reference to sets is no reason to worry. In general, sets are defined to be (sorts of) collections of objects; and, in general, we commonly say that collections have parts. Picasso's paintings, for example, are parts of art collections. Moreover, notice that the usage of 'part' in these sorts of cases is not transitive. Although Picasso's paintings are parts of art collections, we do not say that the oddly placed eye and nose in any given Picasso painting is a part of one of those collections. There is, then, at least one respect in which a sort of parthood relation matches the formal character of the membership relation: intransitivity.⁵¹

Also consider Lewis' example of the hyperbola as a class of x, y pairs. Not only do we consider the left and right halves to be parts of the hyperbola, there is also the intuition that each point on the hyperbola is a part of it. Since these points are simply x, y pairs, which are *members* of the hyperbola (viz., the hyperbola-class), if we agree with

⁵¹ On a related note, if one thinks that a member of a set is *some of* that set, then the argument given in favor of mereological monism can be reapplied here: if 'parthood' is simply the name we give to the *some of* relation as it applies to wholes, and we agree that a set is a kind of whole, and that the elements or members of that set are *some of* that whole, then it straightforwardly follows that elements are parts of sets.

Lewis that common usage supports the idea that subsets are parts of sets, then it also supports the idea that sets have their members as parts.

Intuition and common usage, then, suggests that members are parts of sets. Why is that? There are two competing explanations: the devious explanation according to which the intuition arises because of an analogy of formal character, and the straightforward explanation that it arises because members *really are* a sort of part. In this case, the devious explanation can be ruled out entirely. Recall that the intuition that members are parts of sets, according to Lewis, is to be *rejected* on the grounds of formal character. Members of sets, then, should be taken to be parts of those sets; even if membership is not identical to primitive-parthood.

To summarize: insofar as we are convinced that subsets are parts of sets on the basis of intuition and common speech, we should then be equally convinced that members of a set are parts of it as well. What convinces Lewis, and others, that set-membership is *not* a sort of parthood is the difference in the formal properties of membership and *primitive*-parthood. But this difference is irrelevant. There are different varieties of parthood, all of which are defined on primitive-parthood.⁵² And so, Lewis is correct that membership is not *identical* to primitive-parthood, but our intuitions are nevertheless retained under the notion of membership-parthood, or \in -parthood, which, like direct-parthood, is generally intransitive.

Those still attached to Lewis' idea might respond that their intuitions are not as strong in the case of membership. Let us then buttress these intuitions with second

⁵² That is, assuming the doctrine of mereological monism: if there is only one sort of parthood that applies across all ontological categories, then all other varieties of parts, i.e., φ -parts, are only parts *by virtue of* being primitive parts.

argument: the argument from explanatory economy. There is clearly a difference between the following two sets $A = \{1,2\}$ and $B = \{2\}$. Why is there a difference? The answer is straightforward: 1 is a member of A but not B. But now consider the two sets $A^* = \{\{1,2\}\}$ and $B^* = \{\{2\}\}$. Our previous answer does not apply in this case since 1 is not a member of A^* . Although we can introduce the *member of a member* of relation to explain the difference, there is a simpler solution: we can say the difference is a difference in parts: 1 is a part of A^* but not B^* . Notice that the simpler solution maintains its simplicity, regardless of how many set-braces we add. Consider $A^{**} = \{\{\{\{1,2\}\}\}\}$ and $B^{**} = \{\{\{\{2\}\}\}\}$. To explain the difference in terms of membership requires an additional ‘... a member of’ for each additional pair of braces. If members are parts, nothing whatsoever needs to be added; the same answer holds no matter how deep the sets are embedded.

There are, of course, other possible explanations. To explain the difference between A^{**} and B^{**} in terms of non-mereological containment might seem to provide a solution as simple as the suggested appeal to parthood, while avoiding a commitment to the idea that members are parts of sets. Note, however, that taking sets to be containers suffers from the same explanatory drawback noted earlier: if sets contain their members, but not as parts, then we must ask: what is the boundary that separates the contained from the uncontained? It is certainly not the set-braces themselves: they are merely bits of

useful syntax, designating no set-theoretic entity in particular.⁵³ Non-mereological containment, then, does not appear to be a viable option.⁵⁴

Sets, we should then conclude, appear to have at least two *different* sorts of parts: their subsets and their members.⁵⁵ There is no reason to be alarmed or concerned by this. The conceptual foundations of mathematics remain unchanged: membership is still primitive, and subset is defined as it normally is. Working mathematicians can continue to go about their business in whichever manner they see fit. Only philosophers who hope for a Lewisian foundation for set-theory have any reason to be concerned: subset continues to imply part, but *pace* Lewis, the converse does not hold.

Although the main thesis of this section, that members are parts of sets *pace* Lewis, might be cause for philosophic concern, this is not the forum to explore them. They must be left aside as a direction for future research. At this point, it becomes prudent to return to the main issue at hand: hierarchies and their levels.

⁵³ Moreover, we cannot say that the empty-set provides the boundary; for if we say that the empty set, $\{\}$, provides the containing boundary, then the empty set loses its distinctive character. It is no longer empty.

⁵⁴ One might think to avoid the appeal to parthood by appealing instead to the *transitive-closure* of elementhood on each set. But the explanation obtained by doing so is hardly informative: to explain the difference between $\{\{\{1,2\}\}\}$ and $\{\{\{2\}\}\}$ as a difference in their respective transitive-closures, R and R' , leaves us to wonder what the difference between R and R' is; R and R' , after all, are both sets. Notice that the mereological explanation is straightforward: R and R' are different for the same reason $\{\{\{1,2\}\}\}$ and $\{\{\{2\}\}\}$ are; they have different parts.

⁵⁵ I say 'at least' since sets seem to have parts other than their members or subsets. Consider the set $\{\{\{1,2\},3\},4\}$. Intuitively, 2 is part of this set, but it is neither a member nor a subset of it. To explain this intuition by pointing out that 2 is a member of a member of a member of the set only serves to reinforce the intuition that members are parts of sets.

2.2.6 Being “at” a Level

Even if we take for granted that levels are parts of hierarchies, this tells us nothing of the relationship levels bear to their contents.⁵⁶ The use of ‘contents’, however, is suggestive. It suggests that to be in, or “at”, a level is to be *contained* within that level. But what sort of containment? There are two possibilities: mereological containment (i.e., contains-as-part) or non-mereological containment.

There are two reasons that favor mereological containment as general approach to being at a level. One is negative; the other, positive. The negative argument is familiar. If levels contain their contents, but not as parts, then what is the boundary which separates the contained from the uncontained? It is implausible that there is an answer to this question that holds generally across all particular instances of hierarchies. In some instances we might specify spatial boundaries. But such a specification fails to provide a general analysis of hierarchies: some level-boundaries might not be spatial at all. Consider, for example, the hierarchy of pure sets, or Maslow’s hierarchy of human needs. Neither of these well-known hierarchies have anything to do with spatial containment.

Other varieties of non-mereological containment face exactly the same concern. Since we are investigating the notion of a hierarchy as a univocal notion, a general answer to the separation-question is needed. If there is no general answer, then the notion of a hierarchy is fractured into many different notions. Such fracturing is a detriment to hierarchical science. If we must clarify what we mean by ‘hierarchy’ each and every time we use the term, then the label is useless: hierarchies become platitudes without

⁵⁶ Although the claim ‘contents are parts of levels’ was used to argue that levels are parts of hierarchies, that claim only appears within the scope of an assumption, and as the consequent of a conditional. In this section, my goal is to convince the reader of its truth, *simpliciter*.

substance. But that is a counterintuitive prospect. It certainly seems that our various references to hierarchies involve the application of a single general concept. Without a good reason to think otherwise, casting suspicion on the univocality of ‘hierarchy’ is not warranted.

That is the negative argument. Let us now take up the positive argument that being “at” a level, in the most general sense, means *part*. The argument, once again, involves considerations of parsimony. The argument is this: if those things “at” levels were not parts of those levels, then levels would be a needless complication. It would be more straightforward for us to give a hierarchy without “levels” at all, simply by ranking the collections of those things that are “at” each level as being higher or lower, leaving their mereologically detached levels behind.

That is, even if the relationship between levels and their contents is non-mereological, a prior commitment to mereology (revealed by one’s use of ‘part’ in regular or scientific discourse) commits us to a collection of those contents. Insofar as the levels are ranked as higher or lower, the mereological collections of those contents are also ranked as higher or lower than one another.⁵⁷ Committing oneself to the non-mereological interpretation does not release one’s commitment to a competing mereological interpretation. But the extra commitment is unnecessary. Since we are committed to a mereological interpretation on both the mereological and non-

⁵⁷ Since the contents of each level are “at” that level, there is a mapping from those contents to those ordered levels. Since there is such a mapping, and since the *collections* of things “at” each level are just those things taken as one whole, there is, for each level, a mapping from that level to exactly one of those collections such that the mapping preserves the original ranking. Therefore, insofar as the levels are ranked as higher or lower, the collections themselves have the same ranking. They too are ranked as higher or lower. Insofar as there is *any* account of levels whatsoever, there is also a mereological account.

mereological interpretations of being “at” a level, the non-mereological interpretation is the one to avoid.

In simpler terms, if those things “at” each level are not parts of it, then there is good reason to drop levels altogether. It is simpler for us to take all the things that we say are “at” each level, collect them together mereologically, and rank those collections using *higher* and *lower*. Since these odd and detached things we have so-far been calling *levels* have become superfluous, they can be dropped from the hierarchical framework altogether, releasing the term ‘level’ from its hitherto semantic baggage. This leaves us with ordered collections of things that were formerly “at” these, now non-existent, levels. What shall we call these ordered collections? They are ordered as being higher or lower, so why not call them *levels*? After all, there is a vacancy in the denotation of the term; the previous tenant has, just moments ago, been evicted. The term seems fitting, so let us put it to use: insofar as one mentions levels at all, we can thus conclude that “x’s being at a level y” (i.e., x’s being a part of collection y) implies that x is part of y. If we do not accept the implication, then we retreat to a sense of ‘level’ that is inessential, unimportant and uninteresting.

To summarize, there are two reasons to think that those things “at” each level are part of that level. One is negative, the other positive. The negative: the non-mereological approach to being “at” at level is fundamentally mysterious; the mereological approach is not. The positive: we are committed to a mereological approach to *being at a level* in any case, whereas we are not committed to a non-mereological approach at all; to respect parsimony, we drop the excess commitment, leaving us with the mereological approach. Something’s being “at” a level thus implies that it is a part of that level.

To resist this line of argument, one might suggest that it rests on the observation that the term ‘contents’ is commonly used to describe the relationship between levels and those things that are “at” those levels. It might be objected that this observation does not expose a deep understanding of the relationship itself, but only exposes the most common methodological choice. There are other approaches to understanding the relationship between levels and their “contents” aside from literal containment.

One might, for example, simply take “at” to designate the general relationship of association; that is, being “at” a level is to be associated with that level, and nothing more. A univocal account of being “at” a level can be given using the relationship of general association, which does not imply parthood. Being at a certain mailing address, for example, does not imply that one is part of that address.

Note, however, that it is not my ambition to give a semantic analysis of “is at” in the most general sense; what it is to be “at” a physical address is beyond the present concern. The goal in this section is to provide a substantive understanding of the relationship between a level and those things “at” that level. To understand this relationship as one of general association is not substantive. It is trivial. One purpose of giving a substantive account of hierarchies is to resolve the mysterious sense of “association” that levels bear to their contents.

The aforementioned argument from parsimony delivers such a resolution. Regardless of the account of “is at” we favor, any hierarchy that falls under that account, by virtue of our prior commitment to parthood, implies a mereological account anyway. Those with a preference for minimizing one’s commitments should therefore favor a

mereological account of *being at a level* over any others. To fly in the face in parsimony here is ill advised: since a mereological account of “is at” is unavoidable, as was argued above, insisting on a second, non-mereological, understanding of being “at” a level is ruinous to a univocal account of hierarchies. It will not do to leave “is at” mysterious; and since every non-mereological way of resolving ‘is at’ implies a competing mereological resolution, to insist on a non-mereological understanding is to introduce another sense of ‘hierarchy’ altogether.

One might also object that even if we grant that the contents of levels are parts of those levels, there are clearly contexts in which we do *not* want to say that a part of a level counts (in that context) as being “at” that level. For example, consider a three-level hierarchy of Batman’s bat-mobile, where the top level is the *car* level, the middle level is the *engine* level, and the lowest level is the *piston* level. Since pistons are not engines, it appears to follow that the pistons are not “at” the engine level; yet, pistons are parts of that level, qua parts of the engine, which, we have argued, is a part of the engine level. Since parthood is transitive, it appears that parthood does not imply “is at” in all contexts.

In making such an objection, one must be careful not to presume that the piston level is restricted to pistons. Level specifications are guides to what can be found “at” that level, but they are not generally restrictive. That something is not a piston does not, on its own, imply that it is not at the piston level. Consider the place of teeth in the ecological hierarchy, for example. Teeth are not ecosystems, communities, populations, organisms, organ-systems, tissues, cells, molecules, or atoms. If level-specifications were generally restrictive, in the sense that they imply restrictions to the sorts of things that are “at” levels, then teeth would reside outside the hierarchy entirely. Yet, teeth are objects of

interest to evolutionary ecologists, and are therefore to be included in the hierarchy of ecological interests. Level-specifications, therefore, are not restrictive as one might presume.

Nevertheless, it must be acknowledged that there are contexts in which some parts of levels do not count as being “at” that level in that context. Such an acknowledgement can fortunately be reconciled with the mereological account of being at a level. The reconciliation comes with a simple qualification: being part of a level suffices to be “at” that level in the most general *unrestricted* sense. Although there are contexts in which we want a narrower reading of ‘is at’ that filters-out some of the uninteresting parts of the level in question, all this tells us is that ‘is at’ is subject to contextual manipulation. When one removes the contextual filter from our three-level automotive hierarchy, a more general hierarchy appears: one that includes pistons at the engine level, regardless of the interests of the hierarchical analyst. In short: if the analyst is not interested in the pistons at the engine level, she is free to ignore them. By doing so, she has introduced a new sense of ‘is at’, one that narrows down the inventory of a level’s parts by simple fiat. It has already been argued that the contents of levels are its parts, the choice of which part counts, in the context of interest, as being “at” that level is up to the analyst.

To help us understand what the hierarchical analyst is doing with these parts-by-fiat, consider an analogous example. Suppose Bob is part of a hockey team but is, in games, relegated to the bench. Bob never actually plays. But that does not prevent us from (correctly) claiming that that *the team* is playing, while the game is in progress. During the game, our interest is only in the parts of the team that are actually on the ice. We do not deny that Bob is part of the team; he is merely an *uninteresting* part: a part that

we do not care about in the context of the game, either because of an injury, streak of bad luck, or a general lack of talent. In other words, we are, by contextual-fiat, ignoring Bob. We do not deny that Bob is part of (or “is on”) the team, in general; it is only something we deny in the context of the game, since Bob does not fall within the scope of our present interest. Being at a level, in this regard, is no different than being on the team. What *counts* as being at a level, or being on the team, in any given context is restricted by fiat, as determined by our interests. In short, when specifying an account of what it is to be “at” a level *in the most general and unrestricted sense*, it suffices to think of “is at” as “is part of”. Restrictions as to what *counts* as being at a level are later issued by contextual fiat.

The theory of hierarchies found within ontology, then, is a helpful foil in coming to an understanding of hierarchies qua hierarchies. By challenging Poli’s claim that hierarchy theories should be detached from mereology, we have narrowed our candidates for ‘is at’ to only a few options: to be at a level might mean to be part of it, a member of it, a subset of it, to be non-mereologically contained within it, or being “at” a level might be taken as primitive. Taking “at” to be primitive is unhelpful in coming to a substantive account of hierarchies qua hierarchies, since the relation between levels and their contents remains mysterious. To take it to designate a sort of non-mereological containment implies a container – a boundary. Without a general account of these boundaries, the containment view does not provide an adequate understanding.

Following David Lewis (1991), it was noted that inclusion is a sort of parthood. This leaves us with two plausible options: to be “at” a level is either to be part of it, or to

be a member of it. Arguing, *pace* Lewis, that membership is a sort of parthood, we are left with only one plausible options: parthood.

Even if one discovers a good reason to reject the idea that members are parts of sets, there is an independent reason to think that being “at” a level is not a sort of membership. Although ‘is at’ was introduced in the first chapter as a synonym for ‘is in’, and it is often taken for granted that members of sets are, in some sense, *in* those sets, the formal properties of set-membership are not favorable to an account of “is in” with respect to hierarchies. Knowing that an individual is “in” a level of hierarchy H suffices for us to know that the individual is “in” H. That is, ‘is in’ is a transitive relation, whereas set-membership is not. Set-membership, therefore, does not provide an acceptable account of being in a hierarchy.

At this point, I hope to have convinced the reader that the existing idiosyncratic and ontological approaches to a general account of hierarchies *qua* hierarchies are unsatisfying. But these are negative results; they bring us no closer to understanding what a hierarchy *is*. So far, we only have reason to think that hierarchies and their levels exhibit a mereological nature. In the next two sections we will add to this understanding by examining two contenders for a robust positive account: the partial order view, and the strict partial order view.

3.0 The Partial Order View

The partial order view has three main defenders. Let us begin with Herbert Simon and his commentators, and then proceed to the refinements of Burton Voorhees and Mario Bunge.

3.1 *Simon on Hierarchies*

Herbert Simon, in his address to the American Philosophical Society, characterized a hierarchical organization as “a system that is composed of interrelated subsystems, each of the latter being, in turn, hierarchic in structure until we reach the lowest level of an elementary subsystem.” (Simon 1962, p. 468; quoted in Turney 1989, p. 518)

This initial characterization is weak. First and foremost, it is circular; we are left to wonder what it is for a structure to be “hierarchic.” And secondly, it fails to characterize what can be called *bottomless* hierarchies: those hierarchical structures for which there is no lowest-level, implying that the hierarchical decent proceeds into perpetuity.

That his analysis is weak does make it entirely uninteresting. Simon was, as far as I can tell, among the first to point out that contemporary use of the term ‘hierarchy’ bears scant resemblance to its etymological origins. He writes,

Etymologically, the word “hierarchy” has had a narrower meaning than I am giving it here. The term has generally been used to refer to a complex system in which each of the subsystems is subordinated by an authority relation to the system it belongs to... each system consists of a “boss” and a set of subordinate subsystems. Each of the subsystems has a “boss” who is the immediate subordinate of the boss of the system.

Thinking of hierarchies in terms of “bosses” or *hierarchs* (to use another term) limits discussions of hierarchies to a very narrow and restricted conception, a conception that many people find offensively value-laden. Although superiority and inferiority can be

modeled hierarchically, it is a mistake to think that hierarchies are, by their nature, models of superiority and inferiority. In the ecological hierarchy, for example, ecosystems are higher than organisms, but there is no reason to think that ecosystems are in any way superior to organisms. Even in social hierarchies (the sorts of hierarchies with which Simon was most concerned) being higher does not imply being superior. Ranking employees by their respective seniority yields a hierarchy, but it is inappropriate to infer from this ranking that the old guard is in any way superior to the new.

An acceptable general account of hierarchies *qua* hierarchies, then, should be far more general than what the etymology of ‘hierarchy’ will allow. Those who have refined Simon’s approach seem to agree. Yet, as I will argue, the conception of hierarchy that they use is nevertheless too narrow.

3.2 *Webster on Hierarchies*

Simon’s original account of hierarchies was later refined in his 1973 paper, “The Organization of Complex Systems.” His refinement is far more influential. Jackson Webster (1979), for example, adopts this refinement when investigating the appeal to hierarchies within ecology.⁵⁸ According to Simon and Webster, a hierarchy is a partial ordering of some set. Webster writes,

That is, a hierarchical organization of a set, U , with subsets A, B, C, \dots , is formed by ordering the subsets by a relation, R , which specifies that the

⁵⁸ Note that Webster’s work on the ecological hierarchy precedes that of Allen and Starr’s, as discussed above and in the first chapter.

elements of A are higher than the elements of B which are higher than the elements of C. The relation R is a binary relation in U. (1979, p. 120)

According to Webster, this binary “higher than” relation, R, is antisymmetric and transitive.⁵⁹ R, he says, partitions the set U into *levels*. Levels, for Webster, are simply the relata of R, the *higher than* relation. He also notes that there are many sets, and many relations which meet these conditions. He borrows his main example of a hierarchy, “the Chinese box” from Simon (1962, 1973). The series of increasingly smaller boxes are the elements of U, and R is interpreted as *containment*: one box is higher than the other if it contains the other. (*ibid.*)

The box example is quite apt when it comes to the set/subset and system/subsystem characterization of hierarchies. But not all hierarchies are anything like these nested boxes. School rankings, for example, are hierarchical: there are top schools, mid-range schools, and others below those. Washington Adventist University is not a subsystem of the Oxford University system, although I would guess that it ranks lower. The Simon-Webster general characterization of hierarchies only suits a specific concern. It is therefore not a general analysis: it is far too narrow.⁶⁰

3.3 *Hierarchies: Set Theory vs. Mereology*

The Simon-Webster characterization of hierarchies is too narrow for another reason. On their account a hierarchy is, by definition, a set. They are not alone in their set-theoretic

⁵⁹ If xRy and yRx , then $x=y$ (antisymmetry); and if xRy and yRz , then xRz (transitivity). Webster says nothing about reflexivity, but since he claims R is a partial order, we can presume that R is reflexive: xRx , for all x in R's domain.

⁶⁰ To his credit, Simon was more interested in the question of “why are there hierarchies?” than in asking *what*, exactly, they are. His answer: because hierarchical structures are more stable (less liable to accidental breakage) than other complex structures. (1962, p. 470)

approach. Most available accounts of hierarchies, including those discussed below, are described as sets. I agree that set-theory offers a convenient precision when reasoning about hierarchies, but to *define* a hierarchy to be a set is a bit too quick. Such set-theoretic accounts do not offer a fully general characterization of hierarchies *qua* hierarchies. Such accounts are too narrow.

The problem is that not every hierarchical structure *need* be a set: to say otherwise betrays a common philosophical bias for set theory. Such a bias is unproblematic, and even useful, in many cases, but it strikes me as inappropriate in the present circumstances. Why inappropriate? There are four reasons.

First, the standard understanding of sets takes them to be *abstract objects*. It must be *argued* that concrete structures are necessarily not hierarchical to motivate an abstracta-based general characterization. To claim, without argument, that hierarchies are necessarily abstracta is to beg the question against Realists, such as Stanley Salthe.

Second, even if one avoids begging the question against Realism by endorsing a view according to which sets *are* (or can be) concrete (e.g., Maddy 1990), we are still without a *reason* to think that a hierarchy must be a set. Set theory is a very useful method for representing and reasoning about puzzling situations in a precise way, but here we have gone beyond convenient *representations* of what hierarchies are; we are after what hierarchies *are*. Using an ontology of sets to investigate what hierarchies *are* is nowhere near as innocuous as when one uses sets to *represent* hierarchies for purposes of precise reasoning. Sets are theoretically useful, but they are also ontologically controversial regardless of whether or not we take them to be abstracta or concreta. Since

we are, at present, inquiring into the metaphysics of hierarchies, rather than the innocuous representation of hierarchies, to maintain the view that hierarchies *are* sets in the face of *any* sort of controversy, a reason must be given.

Third, by reflecting on the arguments found above (in section 2.2), I hesitate to assume that a convincing reason to take hierarchies to *be* sets can be offered, for the set-based approach lacks theoretical virtues possessed by the competing approach. Given our stated commitment to mereology, for any set S_H that is a hierarchy, there is a more *general* mereological collection,⁶¹ C_H , such that each level of S_H is identical to some part of C_H , and each part of C_H is identical to some level of S_H , but C_H is not a set (by stipulation). Despite not being a set, C_H is nevertheless hierarchical: it has, as we just noted, parts arranged hierarchically. By virtue of being something that has a hierarchical arrangement of its parts, it certainly *seems* appropriate to call C_H a hierarchy despite that it is not a set. Where ‘{ }’ designates a set and ‘[]’ designates a general mereological collection, insofar as {1, 2, 3} is a hierarchy of numbers, so is [1, 2, 3]. By all appearances, taking hierarchies to be sets makes them out to be a more specific kind of thing than necessary; that there are non-sets which are appropriately labeled hierarchies implies that a set-based approach lacks *generality*.

Fourth, with respect to the other theoretical virtue mentioned earlier, namely, parsimony, take note of the difference between the set-based approach and the mereological approach: the empty set, and a plethora of subsets. On the set-based approach, every hierarchy has its subsets, and thus the empty-set (the empty-set is a

⁶¹ More general since every set is a mereological whole, but not every whole need be a set. Sets, as I understand them here, are special kinds of wholes.

subset of every set), as parts; not so for the mereological approach. Despite appearances to the contrary, the empty set is not *nothing*; so, hierarchies defined using the set-based approach do not simply consist in their levels and the contents of those levels (and collections thereof). They have extra bits: their subsets, including the empty set. Thus, strictly speaking, the mereological approach is more parsimonious than the set-based approach, if only by a hair. The mereological characterization of hierarchies, then, benefits from *two* theoretic virtues over its set-based competitor: parsimony and generality.

In short: there appears to be no clear reason to impose the requirement that hierarchies *must be* sets. To do so makes hierarchies out to be a more specific kind of thing than necessary; granted that some hierarchies are sets, not all of them need to be. Moreover, a set-based approach is less parsimonious than its mereological competitor (if only by a hair). A purely mereological characterization of hierarchies, such as that found in section 2.2 above, therefore benefits from greater *parsimony* and greater *generality* than the set-based characterization.

3.4 *Voorhees and Bunge on Hierarchies*

Improving upon the Simon-Webster account, Burton Voorhees (1983) generalizes the partial order analysis. He refines it by generalizing away from the “subset” aspects of their characterization:

An hierarchical structure on a system S is a partial ordering of the variables of S. (1983 p. 26)

By framing the analysis in terms of a partial ordering – reflexive anti-symmetric, and transitive – Voorhees regains the precision and clarity of the Simon-Webster analysis without “Chinese boxing” himself in a corner. Voorhees’ analysis appears to be a substantial improvement on the partial order account.

The term ‘variable’ is the most obvious drawback to Voorhees’ analysis. He is upfront with his intention to take the notion as a primitive, but gives no reason for doing so. As was argued in a previous section, unnecessary primitives are to be avoided when giving a substantive analysis. There should be a reason for taking something to be primitive.

Mario Bunge (1969), has what is probably the most rigorous of all the partial order accounts. He offers an axiomatic treatment of hierarchies, beginning with a definition of ‘hierarchy’:

an ordered triple $H = \langle S, b, D \rangle$ where S is a non-empty set, b a distinguished element of S , and D a binary relation in S , such that... S has a single beginner, namely, b [That is, H has one and only one supreme commander]... b stands in some power of D to every other member of S [That is, no matter how low in the hierarchy an element of S may stand, it still is under the command of the beginner.]... For any given element y of S except b , there is exactly one other element x of S such that Dxy . [That is, every member has a single boss.]... D is antisymmetric and transitive... D represents [mirrors] domination or power. [That is,... the behavior of each element of S save its beginner is ultimately determined by its superiors.] (1969, p. 18)

Bunge's appeal to 'commander' and 'boss' are merely illustrative. He is not claiming that hierarchies are conceptually bound to the notion of authority. What matters, for Bunge, is that you have a topmost level, with branches extending downward, but never upward. Anything that fails to meet these conditions, says Bunge, is not a hierarchy (by definition).

Voorhees' and Bunge's respective definitions face a familiar problem. In Voorhees' case it is hidden behind the term 'system'. The problem is that their respective accounts share in the Simon-Webster set-theoretic bias. A system, for Voorhees, is "a set of variables in interaction" (1983, p. 25). For Bunge, the reference to sets is explicit.

Their characterizations then rest on shaky ground. We have every reason to think that there are non-set-theoretic individuals that, exhibiting a hierarchical structure, deserve the label 'hierarchy'. It is senseless to define them away, simply to exploit the convenient precision of set-theory. Some sets might indeed be hierarchical: the hierarchy of pure sets, for instance. But still, it is not obvious that all hierarchies need to be related to sets. And there is no sense deciding the matter on an *a priori* basis, not when a more general characterization of 'hierarchy' is available.⁶²

A quick adjustment to the partial order view can accommodate this insight, and remove the set-theoretic bias. Using Bunge's formulation, if we replace 'set' with 'collection', which we understand in terms of the parthood relation, then the set-theoretic bias disappears. We are left with a clear, simple, and fully general theory of hierarchies. Or so it appears.

⁶² Recall: since all set-theoretic structures are parthood structures, but not all parthood structures are set-theoretic structures, it is evident that the mereological characterization of 'hierarchy' is the more general characterization.

3.5 Rejecting the Partial Order Account

Unfortunately for us, the partial order account is not *fully* general. At least, it is not general enough for our purposes. The partial order view rules out, *a priori*, structures that are commonly called hierarchies.

Consider an organizational hierarchy for a corporation consisting of *the marketer* (M), *the financier* (F), and *the employee* (E). The upwards-directed arrows indicate that M and F are both higher than E.




Figure 4: An Organizational Hierarchy for a Limited Partnership

Organizational theorists do not hesitate to label the sort of organizational structure we see in Figure 4 as a hierarchy. It is a command hierarchy, ordered by virtue of some having power over others.

Bunge's analysis is obviously at odds with Figure 4: that there is exactly one "boss" for every subordinate is an axiom of his system. What we see in Figure 4 fails to respect that.

The incompatibility extends more generally to all versions of the partial order view. If the partial order view of Simon, Webster, Voorhees, and others, is correct, then the arrow in Figure 4 designates an antisymmetric and reflexive ordering. But there is a

tension between reflexivity and antisymmetry. If we agree that *higher than* is reflexive, then in Figure 4 we must say that *higher than* is *has as much or more power than*, rather than the asymmetric (and irreflexive) *has more power than*. Notice, however, that in the organization depicted in Figure 4, M and F have the same amount of power – they both have power over E and no one else – it follows from the antisymmetry of *higher than* that $M=F$: since M has as much or more power than F, and F has as much or more power than M, an antisymmetric *higher than* implies that they are identical. The Marketer and The Financier, however, are not identical; they are different people with different talents. If we take Figure 4 to be a genuine hierarchy, ordered by the power one has, then the partial order view cannot provide a general account of hierarchies *qua* hierarchies.

Also note that organization hierarchies are not merely a special case. Jaegwon Kim (2002) and William Wimsatt (1976) are recalcitrant that hierarchies in science can, and often do, branch upward. For example, Wimsatt’s hierarchical model of the scientific domain (presented in the fourth chapter of this dissertation) is massively upward-branching. Kim’s model is similar, but involves a much more modest degree of branching.

Figure 4 is, in fact, suggestive of a more serious problem. Let us flip the image in Figure 4 upside-down.⁶³ In this new picture, let M denote Mary, let F denote Frank, and let E designate the executive in charge. In such an organization, Mary and Frank have as

⁶³ Note that the inversion principle mentioned earlier is incompatible with the partial order view as we have presented it. If certain upward-branching hierarchies are ruled out *a priori*, then not every hierarchy can be flipped upside-down to produce an inverse hierarchy. Downward branching hierarchies, for example, cannot be flipped. Our present appeal to the inversion principle, however, is not problematic in this way: we are appealing to inversion to produce a downward branching hierarchy from an upward branching one. Downward branching hierarchies are, we can assume, acceptable to the partial order theorist. Inverting Figure 1 *should be* unproblematic from the partial order perspective.

much or more power than one another, since they have exactly as much power in the organization – absolutely none whatsoever – it follows from the antisymmetry constraint that Frank and Mary are identical. Not only must the partial order theorist rule out upward branching hierarchies. They must rule out downward branching hierarchies as well. Such a consequence is unacceptable; the partial order view is to be rejected.

4.0 The Strict Partial Order View

Leaving the partial order view behind us, let us proceed to its main competitor: the *strict* partial order view. The main difference between the two is ever so slight. The partial order view takes higher than to be antisymmetric, whereas the strict view takes it to be asymmetric: if x is higher than y, then y is not higher than x.

Aside from the problem of generality that we raised for Simon and his followers, usage of the term ‘hierarchy’ within philosophy and the sciences, in fact, favors the strict interpretation. Oppenheim and Putnam (1958), Kim (2002), and Craver (2007, Chapter 5), for example, are all upfront with their prejudice that higher than is asymmetric. None of these thinkers, however, cares to flesh out this prejudice in the pursuit of a fully general theory of hierarchies qua hierarchies. In fact, there are only two strict partial order accounts that I am aware of. One comes from Richard Dawkins, and the other comes from the international society for systems science (henceforth, ISSS).

4.1 Dawkins on Hierarchy

Richard Dawkins (1976 p. 9) attempts to give the notion of hierarchy a precise foundation using the language of set theory. Although I maintain that this is a problem for

a fully general theory of hierarchies, we already know how to forgive someone of set-theoretic bias: we simply read ‘set’ as ‘collection’.

For Dawkins, a hierarchy is a set that meets two conditions: first, no element of the set is superior to itself; and second, there is one element of the set, the *hierarch*, which is superior to all other elements of the set. *Superior to* is an asymmetric and transitive relation.

There are two problems for Dawkins analysis, however. Both of them parallel the problems faced by Simon. First, the analysis of ‘higher than’ is empty. The superiority relation that defines the hierarchy is, for Dawkins, understood in terms of a primitive relation ‘is boss of’: x is superior to y iff either x is boss of y, or x is boss of something superior to y. (1976, p. 8) While it is possible that Dawkins intends to define the superiority relation recursively, he has not expressed that intention clearly. His account of *higher than* is clearly circular. We are left with no understanding of what “higher than” is supposed to mean. Moreover, Dawkins is lead to take ‘boss of’ as a primitive relation; and moreover, one which we are expected to grasp by metaphor. Those who are interested in the conceptual underpinnings of hierarchical science will find this very frustrating.

The second problem is that, as with partial order theorists, his account is too narrow. He limits the class of hierarchies to those with a *hierarch* – a top level. Where Simon faces the problem of bottomless hierarchies, Dawkins then faces the problem of topless hierarchies. Any structure without a top-level is *ipso facto* not a hierarchy on Dawkin’s analysis.

What makes this narrowing especially problematic is that Dawkins fails to justify this restriction. There appears to be no reason to disqualify a structure as a hierarchy simply because it has no top level: the natural numbers, for example, could be ordered hierarchically, where each successor is superior to its predecessor. The hierarchy of pure sets, for another example, is generated by a recursive process in which the number of levels increases with every step of the recursion.

Dawkins anticipates these sorts of worries. To resolve them, he suggests that topless hierarchies are not hierarchies at all, but are rather an *infinite collection* of finite hierarchies. The idea is that we can give an infinite series of finite hierarchies, where for each hierarchy, H, there is another finite hierarchy, H*, such that H's hierarchy is the immediate subordinate to H*'s hierarchy. In simpler terms: for every hierarchy there is another hierarchy with one extra level on top.

A nice trick, I can agree; but I cannot agree that it gives Dawkins a genuine solution to the problem. Consider, again, the set of natural numbers. Hierarchies become *a priori* finite at the top level, on Dawkins' view. The set of natural numbers, however, is infinite. So, the set of natural numbers is not hierarchical (nor is the set of integers, for that matter). Dawkins allows for an infinite number of finite hierarchies corresponding to the natural numbers, but that is not the point. The point is that the set of natural numbers must have a hierarchical structure, insofar as the hierarchy of pure sets does: for every countable iteration for the hierarchy of pure sets, there is an isomorphism from the natural numbers to the levels of the pure-set hierarchy, qua countable. Thus, if the countably infinite hierarchy of pure sets has a hierarchical structure – and it does – then so does the set natural numbers, qua a self-evident structure preserving mapping from the

levels into the set of numbers. Dawkins view cannot be correct. No account of hierarchies that requires the presence of a hierarch can be correct. Not if one's hope is to provide a conceptual underpinning to the usage of hierarchies within science and philosophy.

4.2 *Hierarchies as Directed Acyclic Graphs*

There is another way to interpret Dawkins' view, independently of his explicit definition. In describing his concept of a hierarchy, Dawkins makes use of the language of Graph Theory. His usage suggests the following alternative characterization of hierarchies: *a hierarchical structure is an acyclic directed graph*. Nothing about this definition demands the presence of a hierarch. There can be infinite hierarchies insofar as there are infinite acyclic directed graphs.

A directed graph is a mathematical structure $\langle V, E \rangle$ where V is a set of vertices and E is a set of edges. Typically we visualize these vertices as objects and the edges as lines with arrows, which represent direction. What makes a particular directed graph *acyclic* is that it contains no completed circuits: you cannot follow the arrows in any way that forms a circle.

There is a problem with the graph-theoretic definition, one that we raised earlier for set-theoretic definitions: we are confronted with abstracta-bias; this time, more straightforwardly so than in the case of sets.⁶⁴ If we understand hierarchies *as* graphs, then hierarchies, by definition, cannot be found in the physical world. We have defined away the issue of Realism, which is exactly what we *do not* want to do.

⁶⁴ While some might think of sets as concrete things, I know of no one who defends graphs as concreta.

The fix is simple. Let us then define a hierarchy as anything (physical or mathematical) that can be *represented* using an acyclic directed graph. Generality appears to be restored. But there is still one minor problem. We have not given an analysis of what hierarchies *are*; we have merely said what they are *like*. We can all agree, I hope, that hierarchies correspond to acyclic directed graphs; but that does not tell us what hierarchies are, in the most general sense.

There is, in fact, another problem with the directed graph view. It is a problem that also applies to the standard view of hierarchies for the International Society for System Science (ISSS). Let us introduce their view before issuing the criticism.

4.3 *The ISSS on Hierarchies*

The ISSS has an official answer to the question: What are hierarchies? They say that “a hierarchy is a collection of parts with ordered asymmetric relationships inside a whole.”⁶⁵ The reference to asymmetry identifies the ISSS as a class of strict partial order theorists for that reason. Although the ISSS appears to take no official stance regarding the transitivity of hierarchies, such a stance is unnecessary. Everyone takes *higher than* to be transitive. There is no need for them to make it explicit.

This, I take to be evident, is the best theory of hierarchies that we have seen so far. It is mereological in nature, contra Poli. It takes *higher than* to be asymmetric rather than antisymmetric. And it tells us what a hierarchy is: a collection of things exhibiting an asymmetric (and transitive) higher than structure.

⁶⁵ <http://www.iss.org/hierarchy.htm> (accessed 4/13/12). The web-document was compiled by none other than Timothy Allen (the ‘Allen’ in Allen and Starr), who thus appears to have abandoned his idiosyncratic 1982 view.

Alas, there are two lingering problems. First, the account is overly general. There are structures consistent with the ISSS definition and the acyclic graph definition that are quite clearly not hierarchies (although one might still say that they are *hierarchical* in some sense). Consider Figure 5, below:




Figure 5: A Broken Hierarchy: a is higher than b, c is higher than d.

In Figure 5 we see a counterexample to the strict partial order view: the disconnected, or “broken”, hierarchy. If a collection or system S , for example, has four components, a , b , c , and d , and the ordering relation R is such that aRb , cRd , $\neg bRa$, and $\neg dRc$, and neither a nor b stands R to either c or d , then R satisfies asymmetry and transitivity. Thus, according to the best strict partial order views we have – the graph-view and the standard ISSS view – S is a hierarchy. However, intuitively, S is not a hierarchy: it is two disconnected hierarchies. An extra constraint must be added to rule out the broken hierarchy. I will offer a suggestion in the next chapter.⁶⁶

The other problem is one of motivation. Granted that there are all these different accounts of hierarchies *qua* hierarchies, why should we think that the strict partial order view, however it is articulated, is really getting at what scientist and philosophers mean

⁶⁶ Within the language of graph theory, we can simply add a *connectedness* constraint. That is, we can say that a hierarchy is anything that can be represented by a *connected* acyclic directed graph. Such a modification does nothing to resolve the other problems mentioned with the graph approach to hierarchies, and so, in the next chapter a different solution is offered.

when they talk about hierarchies? Although I think that strict partial order theorists *are* getting at how the term ‘hierarchy’ is used within science, it is not obvious. Nor can we say that their view is justified on the grounds of an argument-by-cases. Although the partial order and idiosyncratic contenders face serious concerns, there might be a host of alternative accounts that we did not cover.

So we should not trust the strict partial order view. We have not refuted it, as we did for its known competitors. It might even seem intuitive. But we have not argued for it. So we have no reason to think that the strict partial order view gives an adequate account of hierarchies and hierarchical structures mentioned in science and philosophy.

5.0 A Superficial Account of Hierarchies

Notice that most of the accounts we have discussed here share a similar ambition: to *make precise* the scientific notion of a hierarchy.⁶⁷ Within all of the positive accounts we have seen, notice, there is an *imprecise* notion of hierarchy lurking in the background.

One that everyone appears to share:

A hierarchy is a collection of things (call them “levels”) linked together by the higher than relationships among them.

This conception is not precise. It is superficial. It says nothing about what a level is, what linkage is, and most importantly, what *higher than* is.

To *argue* for the strict partial order view, I propose that we begin by abandoning it altogether. Striking the view from our minds, let us then, in the next chapter, try to

⁶⁷ Poli is an exception. Poli thinks that we should ignore hierarchies and shift our attention to levels.

recover it via argument. Such a recovery, however, requires that we begin *somewhere*. I will begin with the superficial account stated above. This one sentence, it seems, is what lies at the heart of all these accounts. The problem is that no one can agree on what it means. What is a level? What is a higher than relationship? What is it to be linked by one? These are all good questions. The answers are not trivial.

6.0 Conclusion

In this chapter we have reviewed the main theories of hierarchies qua hierarchies that have been influential to scientists and philosophers over the past fifty years. The idiosyncratic view of Allen and Starr is shockingly specific. They think that hierarchies, by their nature, have something to do with rates of behaviors. I cannot help but to agree with Potochnik and McGill (from the previous chapter) that this is a mistake.

The idiosyncratic view of Roberto Poli, which ignores theories of hierarchies in favor of a theory of levels, presents a stark contrast to Allen and Starr's approach. Although Poli had not formulated any sort of general definition, his work nevertheless serves as a productive platform for investigating the connection between hierarchies and mereology. Using Poli as a foil, it was argued that (i) levels can be objects or classes, depending on the circumstance; (ii) that levels are parts of hierarchies; (iii) that those things "at" each level are part of it; and (iv) that each part of a level is "at" that level, in the unrestricted sense, but can always be ignored by fiat, depending on the practical goals of the hierarchical analyst. These features will be taken into account in the next chapter, where I develop a general theory of hierarchies from the ground up.

We also examined the partial order views of Simon, Webster, Voorhees, and Bunge. All these thinkers take ‘higher than’ to designate an antisymmetric relation, and define hierarchies as a sort of set. The partial order view, I argued, was too narrow to serve as a general theory of hierarchies qua hierarchies. First, set-based approaches to the *metaphysics* of hierarchies (rather than the *representation*) of hierarchies suffer drawbacks that are not to be taken lightly; second, branching hierarchies pose a problem for their account of *higher than*.

Lastly, we investigated the *strict* partial order views of Richard Dawkins and the International Society for Systems Science. Although they provide the most intuitive definition it suffers from two problems: it is too general, since it makes “broken” hierarchies out to be actual hierarchies; and it is also undermotivated. Building a case for the strict partial order view is my main goal in the next chapter.

Chapter 3

Core Hierarchy Theory

1.0	INTRODUCTION	118
2.0	CORE HIERARCHY THEORY: SUPERFICIAL AND NAÏVE.....	120
2.1	<i>WHAT IS A HIERARCHY?</i>	121
2.1.1	<i>A Collection of Levels</i>	121
2.1.2	<i>... Linked Together by</i>	125
2.2	<i>... HIGHER THAN</i>	126
2.2.1	<i>Transitivity and Complex Non-branching Hierarchies</i>	127
2.2.2	<i>Higher vs. Lower</i>	130
2.2.3	<i>Hierarchical Identity and Distinctness</i>	132
2.2.4	<i>The SPO Analysis of Higher than</i>	134
3.0	OBJECTIONS TO CORE HIERARCHY THEORY	145
3.1	<i>THE “TOO GENERAL” OBJECTION</i>	145
3.2	<i>THE ARBITRARINESS OBJECTION</i>	147
3.3	<i>THE EMPTINESS OBJECTION</i>	148
4.0	CONSEQUENCES OF CORE HIERARCHY THEORY	150
4.1	<i>AN ARGUMENT FOR HIERARCHICAL REALISM</i>	151
4.2	<i>TRANSITIVITY OF ‘IS A LEVEL OF’</i>	154
4.3	<i>ASYMMETRY, IRREFLEXIVITY, AND ‘IS A LEVEL OF’</i>	155
4.4	<i>COMPOSITIONAL HIERARCHIES ARE LEVELS OF THEMSELVES</i>	157
5.0	CONCLUSION.....	161
	APPENDIX A: A FIRST-ORDER RECONSTRUCTION OF CORE HIERARCHY THEORY	164
	APPENDIX B: THEOREMS FOR UNRESTRICTED HIERARCHIES	170

1.0 Introduction

At this point we have only the most superficial grasp of hierarchies. Our understanding is restricted to a mere surface-description of the basic concept, the concept that the theorists described in the previous chapter were hoping to capture. Hierarchies, we said, are

collections of levels linked by a *higher than* relation. It is a description that everyone appears committed to, but it is not a useful one. It is vague. Making it precise has been difficult. We have seen that the most influential attempts to do so are inadequate.

The main point of disagreement among hierarchy theorists, and the source of the inadequacy in their accounts, surrounds the nature of *higher than*. Our first task in this chapter is to come to an understanding of this relation. After justifying, and building upon our surface description, what we will end up with is a version of the strict partial order view, a version which provides a suitable foundation for a robust formal theory of hierarchies. Herein is our secondary purpose: to develop a formal theory of hierarchies that can be made available to scientists, philosophers, computer scientists, ontologists, engineers, organizational theorists, and anyone else who might find hierarchies useful. The formal development can be found in Appendices A and B.

Having formulated and defended a general theory of hierarchies, our second task is to explore its implications for the central questions of the dissertation, specifically, the question of hierarchical Realism (as it was characterized in the introduction). I argue that the account we developed, call it *Core Hierarchy Theory*, supports a Realist interpretation of hierarchies, insofar as we are realists about scientific phenomena to begin with. In this regard, I am arguing in favor of Stanley Salthe's (1985) position over the standard view of Allen and Starr (1982). Yet, on the other hand, I also argue that Allen and Starr's approach to the ecological hierarchy, as well as Simon and Webster's approach to hierarchies – viz., the nested approach that we encountered in the first and second chapters – is a direct consequence of Core Hierarchy Theory when the “is at”

relation is interpreted as the primitive-parthood relation, rather than a restricted φ -part relation. I call such hierarchies “unrestricted” hierarchies.

This chapter is structured as follows. I begin with the superficial understanding of ‘hierarchy’ that we ended off with in the previous chapter, systematically removing the superficiality surrounding each of its substantive lexical components: collection, linkage, and higher than. I then proceed to an analysis of higher than. I argue that every such relation is an asymmetric and transitive relation; and that every asymmetric and transitive relation can be given a *higher than* interpretation. The argument for the latter claim is regrettably quite complicated. I show, in two different ways, that any asymmetric and transitive relation suffices for the creation of a *higher than* structure that is isomorphic to the asymmetric and transitive ordering we began with. I conclude with some consequences of Core Hierarchy Theory. Within a formal reconstruction of Core Hierarchy Theory, I show that this simplest and most general theory of hierarchies implies that, when the “is at” relation is unrestricted (i.e., when ‘is at’ denotes primitive-parthood), every compositional hierarchy is a level of itself and every such compositional hierarchy is a nested hierarchy. This result vindicates the coherence of Simon (1962, 1973), Webster (1979), Allen and Starr (1982), and Salthe’s (1985) nested approach to compositional hierarchies.

2.0 Core Hierarchy Theory: Superficial and Naïve

In this section, I carve out a general theory of hierarchies qua hierarchies. Such a theory requires four things: (i) an account of ‘hierarchy’; (ii) an account of ‘*higher than*’; (iii) an account of ‘level’; and (iv) an account of what it is to be “at” a level.

Two of these needs have already been satisfied. We satisfied (iv) within our interrogation of Roberto Poli on mereology: to be at a level to be some part of it that fits the interests of the hierarchical analyst. We satisfied (iii) with our examination of Poli and the partial order theorists: a level is the relatum⁶⁸ of *higher than*,⁶⁹ whatever that means. In what follows, then, we need only focus our attention on (i) and (ii).

2.1 *What is a Hierarchy?*

Our superficial description – a hierarchy is a collection of levels linked together by *higher than* – gives us a substantial head-start regarding the satisfaction of (i) – an account of ‘hierarchy’. We cannot say that we have satisfied it just yet. The description is only a useful starting point.

We must first ensure that we understand the terms used to formulate it: ‘collection’, ‘linked’, ‘level’, and ‘higher than’. ‘Level’ is a notion we already understand. The notion of ‘higher than’, however, is not nearly so simple. Until we turn our attention to ‘higher than’, it is taken as a primitive.

2.1.1 *A Collection of Levels ...*

Let us begin by clarifying ‘collection’, a notion that I introduced earlier. Most hierarchy theorists, recall, are in the grip of a set-theoretic bias, and tend to frame things in terms of *sets* rather than collections.

⁶⁸ Either an individual or Poli’s categories/classes.

⁶⁹ Recall that, in Chapter 1, we saw fit to draw a distinction between *is a higher level than* and *is at a higher level than*. The former holds between levels, the latter holds between the contents of levels. By convention, I use the term ‘higher than’ to designate the former, rather than the latter, relation. A justification for this convention can be found later in the chapter. Briefly put, *is at a higher level* is more easily defined in terms of *is a higher level* than vice versa. By ‘higher than’ I mean to designate the more general relation: *is a higher level than*.

We argued against that bias in the previous chapter. If the set-membership and subset relations are a sort of parthood, as was argued, then hierarchies *are* collections in our sense of ‘collection’. We are understanding ‘collection’, recall, to be anything with proper parts. Following David Lewis, we argued that subsets are parts of sets, and *pace* Lewis, we argued that its members are too. Insofar as ‘hierarchy’ applies to sets, then, it also applies to collections. Since calling them *collections* is metaphysically neutral, whereas calling them *sets* is not, we must then prefer the mereological characterization to avoid begging any questions regarding hierarchical Realism.

But perhaps we are unwittingly begging another question. One might object to ‘collection’ on the grounds that there is an even more general term: ‘plurality’. If we use ‘collection’, then we must maintain, on *a priori* grounds, that no “uncollectable” plurality has a hierarchical structure. If it had such a structure, then it would appear to count as a hierarchy.

Here it is useful to draw a distinction between hierarchies and hierarchical structures. A hierarchy is a thing that has a hierarchical structure, which is a sort of relational structure. The relational structure that makes hierarchies *hierarchical* is simply the *higher than* relation. Uncollectable pluralities, then, can still have a hierarchical structure insofar as each thing in the plurality is linked to the others by *higher than*. These pluralities are not hierarchies, but they are still hierarchical.

One might question the usefulness of this distinction. The arguments found in the previous chapter involve two constraints on acceptable accounts of hierarchies: parsimony and generality. One might argue that taking hierarchies to be uncollected

pluralities offers a more parsimonious and a more general account than taking them to be individuals. More parsimonious because it eliminates the commitment to anything aside from levels and their contents, and more general because it accounts for potential hierarchies whose levels cannot be collected into a single individual.

I maintain that the distinction is a useful one. There is good reason to think that a hierarchy is not a plurality, and there is a good reason to think that it is not simply a sort of relational structure (i.e., a relation) defined on a plurality or collection.

First, let us say why hierarchies are not pluralities. There is no doubt that one can individuate hierarchies and differentiate between them: ECOHIERARCHY is one hierarchy, a corporate hierarchy is another. If hierarchies were pluralities, then each and every one of these hierarchies would, in fact, be many things (viz., levels). But, as David Lewis (1991) points out, one thing is not many things. Therefore, hierarchies are not pluralities. A hierarchy is one thing; pluralities are many things.

A collection, as I use the term, is simply one thing made from many. Anytime one attempts to count pluralities, “one plurality, two pluralities, three pluralities, four...” one is counting collections. I simply cannot make sense of the idea that we are enumerating the many when counting in this way: not only is it odd to say, “one many, two many, three many, four...”, it seems outright incoherent. To argue for its coherence is tantamount to arguing that the one can in fact be many, *pace* Lewis and basic common sense. There are few maxims as self-evident as “one is not many”.

Now let us say why it will not do to define a hierarchy as a sort of relational structure (i.e., a relation) defined on a plurality or collection.

If hierarchies are relations, then they must have relata. What are their relata? In the previous chapter it was argued that hierarchies consist only in their levels and the contents of those levels. Taking this argument for granted, we then have three options: the relata are either levels, contents of levels, or both. Note, however, that we have already designated the relationship between levels as the *higher than* relationship⁷⁰, the relationship between the contents of levels as the *is at a higher level than* relation, and the relationship between levels and their contents as the *is at (or is in)* relation.

If we say that the *hierarchy* relation is any one of these three (or some other relation I might have neglected to mention) then we are left with a denotational gap: what do we call the collection of levels itself?⁷¹ We can call it whatever we want. Why not ‘hierarchy’? Such a label is surely not arbitrary: since we already have labels for the three aforementioned relations, the term ‘hierarchy’ seems more appropriately assigned elsewhere. Insofar as a hierarchy consists only in levels and their contents, the assignment of ‘hierarchy’ to the collection of levels certainly appears to be the most appropriate option for filling the denotational gap. It is inappropriate, recall, to assign it to any plurality.

Also, take note of how ‘hierarchy’ is commonly used. If hierarchies are collections of levels, then one would expect ‘hierarchy’ to be used as a monadic

⁷⁰ Or: the *is at a higher level than* relation, which we shortened to *higher than* by convention.

⁷¹ I neglect to mention the three-place relation holding between contents of levels, the levels they are at, and the levels that they are higher than. To think of a hierarchy as this three-place relation is quite inappropriate. For one, it makes hierarchies out to be much more complicated than they appear, or need to be. For another, to identify the hierarchy with this three-place relation still leaves a denotational gap with respect to the collection on which the *hierarchy* relation is defined.

predicate, whereas if they were relations, one would expect common usage to reflect this; we would expect to hear "... is a hierarchy of (levels) $L_1 \dots L_n$ " whenever a hierarchical reference is made. Since the word 'hierarchy' is most commonly used as a monadic predicate, rather than a relation, usage of the term 'hierarchy' lends no support to the relational view. On the contrary, common usage supports the collection-view. The ecological hierarchy, for example, *is a hierarchy*.

Having argued that hierarchies are best taken to be collections, rather than pluralities or relations, we are ready to move forward with the analysis. Let us then proceed to clarify 'linked together'.

2.1.2 ... *Linked Together by...*

The problem with strict partial order view of hierarchies, defended by Dawkins and the ISSS, is that their respective accounts do not rule out a broken hierarchical structure: two disjointed hierarchies, taken together in a single *disconnected* acyclic graph, satisfy the strict partial order definition. Fortunately, there is a straightforward way to solve the problem. We simply add a linkage constraint: a constraint which ensures that there are no gaps in the hierarchical structure.

One way to obtain an internal linkage amongst the levels is to introduce a hierarch. But, recall that we are no longer under the spell of the hierarch. To include a hierarch is to necessitate a top level, ruling out the topless hierarchy of pure sets, which is a very important hierarchy within science.

There is another way to ensure the appropriate sort of linkage using nothing but the resources of mereology. We have already argued that levels are parts of hierarchies.

Let us now exploit this mereological feature. Where H is a collection, ϕ -part specifies a kind, ϕ , of part⁷² (e.g., direct-part, functional-part, etc.), and ‘R’ is any relation at all (including *higher than*), we can specify our *linkage* constraint generally as follows:

(LINKAGE) The ϕ -parts of H are *linked* by relation R iff for every way of partitioning H into two sub-collections of ϕ -parts, H* and H**, without remainder, there is an x in H* and a y in H** such that x and y stand in the R relation.

Where ϕ picks out *levels*, and R picks out *higher than*, the LINKAGE conditions, when satisfied, ensure that the levels of the hierarchy are joined together to form a single unit: it requires that for every way of cutting the collection in two, there remains at least one R (i.e. *higher than*) connection between parts of the two sub-collections. This linkage condition, when satisfied, thus rules out the possibility of broken hierarchies without introducing the problematic *hierarch*.

2.2 ... Higher Than

So far, we have been taking the *higher than* relation as primitive. Now we must give it substance, preferably in the form of necessary and sufficient conditions.

The argument for this analysis comes in two stages. First, the *necessity* stage: I argue that *higher than* is an asymmetric and transitive relation. Second, the *sufficiency* stage: I argue that all asymmetric and transitive relations are to be considered *higher than*

⁷² Note that ϕ can be given a null interpretation, if, for example, one wants to test whether or not a collection of parts is linked, without having a *kind* of part in mind. To give the null interpretation, simply turn your gaze away from the ‘ ϕ ’ and go about your business as if it does not exist.

relations. These two arguments taken together imply the strict partial order view discussed in the previous chapter.

Let us first argue that *higher than* is transitive. To do so, it is useful to make a distinction between branching and non-branching hierarchies.

2.2.1 Transitivity and Complex Non-branching Hierarchies

Non-branching hierarchies require that for any pair of levels in that hierarchy, one member of the pair must be higher than the other. Branching hierarchies do not have such a requirement: levels residing on different branches are incomparable with respect to the *higher than* relation.

The common way of representing non-branching hierarchies reveals an important feature of *higher than*: it is transitive. Consider Figure 6 below. It is a simple depiction of a three-level non-branching hierarchy: ‘A’, ‘B’, and ‘C’ designate levels, and the upwards-directed arrow designates *higher than* (by convention, the arrows always point to the higher level).




Figure 6: A Simple Non-branching Hierarchy

But there is something missing. Given our understanding of non-branching hierarchies (still taking *higher than* as a primitive), it must then be that, in Figure 6, either level C is higher than A, or level A is higher than C, even though this connection has not been made explicit using an arrow. A, however, is not plausibly higher than C: such a chain would be circular as opposed to hierarchical. Our naïve concept of hierarchies allows for branching and non-branching hierarchies, but it does not allow for circular hierarchies. If it did, then Figure 6 would lead to an ambiguity between the following two representations (Figure 7):

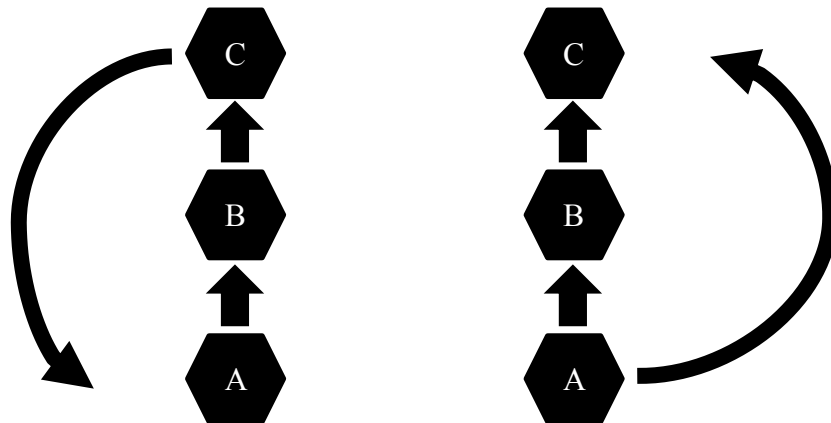


Figure 7: The ambiguity of the missing "higher than"

I know of no one who, when confronted with Figure 6, recognizes the ambiguity depicted in Figure 7.⁷³ As soon as one is told that Figure 6 represents a non-branching hierarchy (in which the arrows point to the higher level), there is little doubt that C is higher than A.⁷⁴

⁷³ Note that the ambiguity is not three-way: “levels A and C are not related by *higher than* at all” is not a third possible disambiguation. Non-branching hierarchies, such as what we see in the above figure, must satisfy the claim ‘for any two levels in the hierarchy, they stand in a *higher than* relationship with one another’.

⁷⁴ Adding arbitrary branches to Figure 6 appears to make no difference to the disambiguation. Since branching and non-branching hierarchies exhaust the logical space of hierarchies, the transitivity disambiguation should then generalize to all hierarchies.

This strongly suggests that transitivity is a necessary condition for *higher than*. If higher than were not transitive, then one would not naturally rule out the circular hierarchy as a possible disambiguation of Figure 6: a paradigmatic expression of a non-branching hierarchy. Rather, we would *ask* for disambiguation if transitivity were not inherent in our concept of higher than; for all the time I have spent showing these pictures to others, no one has ever asked for disambiguation.

One might complain that, as it stands, this “ambiguity” argument suffers from a lack of rigor. I have not performed any meticulous psychological studies to support my hypothesis. One might even suggest that there could be a cultural bias at play in my informal experiment.

Even so, this is no reason to *doubt* the hypothesis that the leftmost structure in Figure 7 would be ruled out by test subjects, were such an experiment to be performed. In the first chapter I noted that an investigation into hierarchical concepts is an active area of research within cognitive science (e.g., Winer 1980; Greene 1989, 1994; Murphy 2002; Deneault and Ricard, 2006; and Slattery et al 2011). The consensus in this literature is that the relationship between levels of hierarchical classification is both transitive and asymmetric (Slattery et al 2011, p. 243; Murphy 2002, p. 200-2); and it has been observed that such hierarchical structure “appears to be a universal property of all cultures’ categories of the natural world.” (Murphy 2002, p. 204)⁷⁵ If one were to perform a more rigorous experiment to test the above hypothesis, and the results were

⁷⁵ Although it is irrelevant to our analysis, it is worth noting that Murphy (2002) assumes (without argument) that hierarchies cannot branch upwards. This, as we will see, is contrary to Jaegwon Kim and William Wimsatt’s view of hierarchies.

contrary to that hypothesis, then that would indeed be a very interesting and surprising result.

Taking *higher than* to be transitive is thus warranted. Let us then state our first necessary condition for *higher than*. For any hierarchy, H, and levels, x, y, and z of H:

(TRANSITIVITY) If x is higher than y and y is higher than z, then x is higher than z.

2.2.2 *Higher vs. Lower*

Although noting the transitivity of *higher than* is important for our understanding of hierarchies, it does not provide us with a complete analysis. To search for an informative definition of *higher than*, it is useful to contrast it with its opposite: the equally notorious *lower than* relation.

The *lower than* relation is no less mysterious than *higher than*. Even so, there is much we can say. For instance, we can say why these relations are equally mysterious: the relations are directly inter-definable. No one denies that the *lower than* relation is simply the converse of *higher than*.

(INTER-DEFINABILITY) x is higher than y iff y is lower than x.

This definition is, of course, trivial: it adds nothing to our understanding of either relation. Still, there is more we can say about the relationship between “higher” and “lower”. For any hierarchy, H, and any two levels x and y of H,

(INCOMPATIBILITY) if x is higher than y, then x is not lower than y.⁷⁶

Why should we endorse INCOMPATIBILITY? Because to deny it is to claim that there is at least one level that is both higher *and* lower than some other level, and we have already ruled out this sort of case. If there are hierarchies that contain such levels, then Figure 6, for all we know, might be such a hierarchy. It is consistent with a denial of INCOMPATIBILITY that *both* of the disambiguations that we see in Figure 7 are correct. But as noted above, no disambiguation is needed. The downward arrow connecting C to A in Figure 7 is immediately ruled out, as soon as one is told that Figure 6 depicts a hierarchy.

Taken together, INTER-DEFINABILITY and INCOMPATIBILITY imply an interesting feature of *higher than*: that it is asymmetric. Given a hierarchy, H, For every pair of levels x and y of H,

(ASYMMETRY) if x is higher than y, then y is not higher than x.

The inference is straightforward: If x is higher than y and y is also higher than x, then INTER-DEFINABILITY tells us that y is lower than x and INCOMPATIBILITY tells us that y is not lower than x; a contradiction. Insofar as we accept these two basic principles, asymmetry cannot be denied.

What we are left with, then, is the strict partial order view. We have not presupposed it. We have argued for it from more basic claims and testable psychological hypotheses. We can, after all, test the hypothesis that Figure 6 is universally unambiguous between the two images in Figure 7.

⁷⁶ Since x and y might be the same level, the conditional does not hold in the reverse direction.

Taking the success of such experiments for granted, we are still only half-way to a complete account of *higher than*. Having argued that *higher than* is asymmetric and transitive, we must now ask: which asymmetric and transitive relation? Here we face a dilemma: either *higher than* is a relation unto its own, or it is a designator for an entire class of relations, each of which count as *higher than* in some context or another.

In the next two sections, I argue for the latter horn of the dilemma. In the next section, I argue that ‘higher than’ designates a class of relations, and then, in the section that follows, I argue that being an asymmetric and transitive relation suffices to count as a *higher than* relation. With that, our analysis will be complete.

2.2.3 *Hierarchical Identity and Distinctness*

We have, so far, been looking only at a single hierarchy. Doing so has been fruitful for our investigation into *higher than*. Contrasting the *higher than* and *lower than* relations has also been fruitful. To press our conclusions further, let us now see what we can glean from contrasting two *different* hierarchies.

Given our naïve (but no longer superficial) definition of hierarchies, it is appropriate to infer that hierarchies H_1 and H_2 are the same if and only if they have exactly the same levels ordered in exactly the same way. For two hierarchies to be distinct, then, is for them to either differ in at least one level, or for them to be ordered in a different way, or both.




Figure 8: Two Distinct Hierarchies H_1 and H_2 – either $B \neq Y$, $A \neq Z$.

In Figure 8, H_1 and H_2 are distinct only if either $B \neq Y$ or $A \neq Z$. If B were Y and A was Z , then both hierarchies would contain the same levels, in the same *higher than* order: H_1 and H_2 would be identical.

This implies that there are different *higher than* relations: if two hierarchies are distinct, but have all the same levels, then those levels must be ordered differently. For example, if we let $B=Z$ and $A=Y$, then what would make these hierarchies distinct is that the ordering would differ. But since they are both hierarchies, they must both be ordered by a *higher than* relation. And so, there is no single relation, *higher than*. There must be many.

But this is to be expected. In all of the hierarchical contexts we care about for the purposes of this dissertation, there is some way to understand the higher than relation in terms of another relation. Things are never just plain *higher than*: in a hierarchy of importance, for example, they are (typically) ordered by the relation *more important than*; in power-hierarchies, they are ordered by *has more power than*. In the upwards-nested ecological hierarchy, the ordering relation is a bit more complicated; after each instanceable and individual has been situated within the appropriate level via AT-LEVEL, it is evident that level L is higher than level L^* only if L contains something that

is not contained within L^* . Being a *higher than relation*, then, is a second-order property: some relations count as higher than, where others do not.

2.2.4 The SPO Analysis of Higher than.

Which relations count as *higher than* relations? At the very least, they must be strict partial ordering (SPO) relations. In fact, it is enough: it can be shown that any such relation can be given as a *higher than* relation.

In what follows, I give two different proofs. The first is much longer than the second. The second is short, less-mechanical, but requires a lot more thinking.

Let me begin with a brief summary of what I am doing with the longer proof. I am showing that for any asymmetric and transitive relation, R , there is a corresponding *higher than* structure. Doing so shows that asymmetric and transitive relations are *higher than* relations, giving us an analysis of *higher than*. To accomplish this, I describe the most basic *higher than* structure, and then use this notion to exhaustively partition the class of *higher than* structures into four groups: those that are basic, those that branch, those that do not branch, and those that are one of the first three with at least one disconnected bit.⁷⁷ I show that any strictly partially ordered set with two (the smallest such set⁷⁸) or three elements can be placed in one of these four groups. I then show that adding an element to the structured set produces a structure that falls into one of the four

⁷⁷ We could exhaustively partition the class into less than four groups, but doing so adds complications.

⁷⁸ Proof: if a set has only one element, a , then any binary relation defined upon it, $\langle a, a \rangle$, is reflexive. Strict partial orders are, by definition, asymmetric. Asymmetry implies irreflexivity. Therefore, the smallest strictly partially ordered set has two *distinct* elements.

groups. The procedure can be repeated ad infinitum, and so all strict partial orders, including infinite ones, can be given as *higher than* structures.⁷⁹

Here are the definitions. A branching *higher than* structure is any structure in which something *x* stands either higher than, or lower than, at least two distinct things that do not stand higher than the other.⁸⁰ A non-branching structure is one that does not branch.

Fixing all that in your mind, here is the longer version of the argument. Consider a set of things, *C*, and a strict-partial order, *R*, on that set. If the ordered set has only two elements – *x* and *y* – that stand in relation *R* to one another, e.g., xRy , we can represent this as, what we see in Figure 9 (below), a basic hierarchy: *x* can be represented as being higher than *y*, for example.




Figure 9: The Most Basic Higher Than/Hierarchical Structure

⁷⁹ Take note that these arguments can be run with respect to collections and their proper parts rather than sets and their elements. The presentation is analogous with ‘collection’ substituted for ‘set’ and ‘proper part’ substituted for ‘element’. I opt for conventional terminology here.

⁸⁰ Note that there are two kinds of branching structures: an upward-branching higher than structure, let us say, is any structure that has at least three distinct elements *x*, *y*, and *z*, such that *x* and *y* are higher than *z*, but neither is higher than the other. A downward-branching higher than structure, by contrast, is any structure that has at least three distinct elements *x*, *y*, and *z*, such that *x* is higher than both *y* and *z*, but neither *y* nor *z* is higher than the other.

By ‘basic hierarchy’ I mean that H cannot be understood to be a combination of any smaller hierarchies. Adding a third element, z, to our set, then it must be either that

- (I) xRz ,
- (II) zRx , or
- (III) neither (I) nor (II).

Presume (I) for conditional proof. There are three options regarding z’s relation to y: (i) zRy , (ii) yRz , and (iii) neither. If (i), then we can depict the situation as a three-level hierarchy such as that found in Figure 6: the transitivity of R implies that z takes the middle place between x and y ($xRzRy$).

If (ii) then either (a) yRx , (b) xRy , or (c) neither. If (ii) and (a) hold, then the sequence $yRxRz$ results in another instance of Figure 6. If (ii) and (b) hold, then $xRyRz$ results, again resulting in an instance of Figure 6. If (ii) and (c) hold, then yRz and xRz , but y and x are incomparable with respect to R: a situation that we can depict using an upward- branching hierarchy, as we see in Figure 10 below.

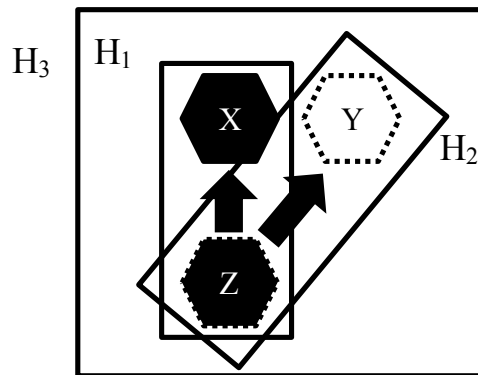


Figure 10: An Upward Branching Hierarchy H_3 as Two Basic Hierarchies H_1 and H_2 That Have Level z In Common

If (iii), then either (a) yRx , (b) xRy , or (c) neither. (I), (iii), and (a), however, are not consistent with R being a strict partial order: (a) yRx , (I) xRz , and the transitivity of R imply (ii) yRz , which implies not-(iii).

If (iii) and (b), then all three of the following are true: xRz , xRy , z and y are not R -related. We can depict this situation as a downward-branching hierarchy (Figure 11):

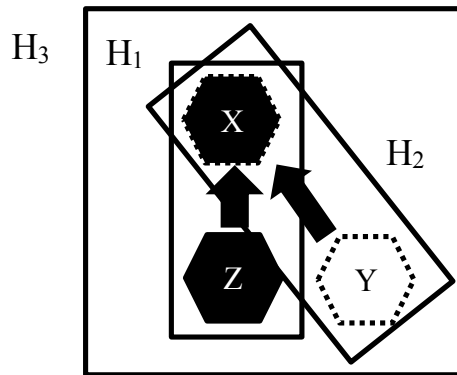


Figure 11: Downward Branching Hierarchy H_3 as Two Basic Hierarchies H_1 and H_2 That Have Level x in Common

If (iii) and (c) hold, then y does not bear R to anything at all, and what we are left with is a higher than structure consisting of a basic hierarchy and a dangling y left over.

Closing our first conditional proof, we have shown that every strict partial ordering of three elements that satisfies (i) yields a higher than structure: either a hierarchy, or a hierarchy with a dangling bit left over.

Presume (II), zRx , for conditional proof. We are confronted with three familiar options: (i) zRy , (ii) yRz , and (iii) neither. If (i), then either (a) yRx , (b) xRy , or (c) neither. If (i) and either (a) or (b) hold, then a non-branching hierarchy (as in Figure 6) results. $zRyRx$ is the sequence for (a); and $zRxRy$ is the sequence for (b). If (i) and (c)

hold, then a downward-branching hierarchy results: we have zRx , zRy , but x and y do not stand R . That is, we get Figure 11, except with x and z having switched places.

If (ii), then the transitivity of R , given our present assumption that zRx , implies that yRx . Again, we see the non-branching Figure 6. This time ordered (from top to bottom): $yRzRx$.

If (iii), neither zRy nor yRz , then there are three possible options: (a) yRx , (b) xRy , or (c) neither. If (iii) and (a), then we have an upward-branching hierarchy: yRx , zRx , with y and z being incomparable with respect to R gives us Figure 10 with z and x switching places.

If (iii) and (b), then R is not a strict partial order. By our presumption that (II), we have zRx , by (b) we have xRy , and so, if R is transitive (which it must be if R is a strict partial order), then zRy must be the case, which implies not-(iii). (iii) is the claim ‘neither zRy nor yRz ’.

If (iii) and (c), then we have a higher than structure consisting in a basic hierarchy – z ’s being higher than x (from II) – and a bit left over: y .

Closing our second conditional proof: every strict partial order that satisfies (II) corresponds to some *higher than* structure: either a hierarchy or a higher than structure with a bit left over. Only one more presumption remains.

Presume (III) for conditional proof. If (III), then again there are three options: either (i) zRy , (ii) yRz , and (iii) neither. If (i), then we face three more options: either (a) yRx , (b) xRy , or (c) neither.

If (i) and (a) hold, then R is not a strict partial order. If zRy and yRx hold, and R is transitive, then zRx follows. zRx implies not-(III).

If (i) and (b) hold, then from our presumption that (III), we derive that zRy , xRy , but z and x do not stand R . This situation can be depicted as an upward-branching hierarchy: Figure 10 with y and z switching places.

If (i) and (c) hold, then zRy , but x does not stand R with anything. That is, we have a higher than structure consisting in a basic hierarchy in which z is higher than y , with an extra bit left over: x .

If (ii), then we face three options. Either (a) yRx , (b) xRy , or (c) neither.

If (ii) and (a) hold, then a downward-branching hierarchy is the result: y is higher than both z and x , but x and z are not comparable with respect to R . That is, we have Figure 11 with x and y switching places.

If (ii) and (b) hold, then R is not a strict partial order. (b) xRy , (ii) yRz and a transitive R imply xRz , which implies not-(III). Maintaining the presumption that (III), (ii), and (b), it follows that R is not transitive; and therefore, not a strict partial order.

If (ii) and (c) hold, then we have yRz , but x is R -related to neither y nor z . So, again we have a higher than structure consisting of a basic hierarchy – y being higher than z – and an extra bit left over: x .

If (iii), then neither zRy nor yRz . Given (III), we can add that neither xRz nor zRx . That is, z is not R -related to either x or y . We have three more options: (a) yRx , (b) xRy , or (c) neither.

If (iii) and either (a) or (b), then we have a *higher than* structure consisting of a basic hierarchy and a bit, z , left over: if (a), then y is higher than x ; if (b), then x is higher than y .

If (iii) and (c), maintaining our presumption that (III), then nothing in our set, C , stands R to anything else. There is no R relation at all.

Closing our third, and final, conditional proof, we see that for every strict partial order that satisfies (III), there is a corresponding higher than structure: either a hierarchy or a hierarchy with extra bits left over. This is the same result we get for both (I) and (II), which exhausts the logical space. Therefore, for every relational combination of three items that satisfy a strict partial ordering, we can depict the situation using a higher than structure: either a basic, branching, or non-branching higher than structure, with or without disconnected bits leftover.

From this, we can generalize to an R -ordered set with n elements. If there is a fourth element, f , we merely compare its R -relatedness to each of the existing elements. We can depict it as being lower-than all those elements α for which $\alpha R f$; and we can depict it as higher than all those elements β for which $f R \beta$. If there is no α such that $\alpha R f$, then either f is at the top of its hierarchy, or it is an extra disconnected bit of the higher than structure we already have; if there is no β such that $f R \beta$, then either f is at the bottom of its hierarchy, or it is an extra disconnected bit of the higher than structure we already have. Under all of these options, the addition of an extra element nevertheless yields a higher than structure, constructed out of one or more basic hierarchies, perhaps with some extra bits. Since we can repeat the process described in this paragraph ad infinitum,

we have the generalization we need. We have deduced that for every strict partial order relation R , there is a *higher than* structure that depicts it. Therefore, being a strict partial order relation suffices for serving as a *higher than* relation.

Now for the shorter proof. What is different here is that it begins with a strictly partially ordered collection of any size. To make our lives simpler, we will begin with *linked* ordered sets. We will generalize to unlinked ordered sets afterwards.

(Preliminaries) Let O be a set with cardinality n (call it the ‘origin set’) and let R be a binary strict partial order (SPO) relation on O that *links* O ’s elements (in the sense of *linkage* defined above). Let H be an empty collection that is filled⁸¹ with elements according the recursive process below (called ‘inductive step’); and let \bar{H} denote a *higher than* relation on H .

1. Take any two arbitrary elements of O , x and y , such that xRy and let H now contain them such that $x\bar{H}y$.
2. Now take any v that is in O , but not in H , and then for every w that is found in both O and H , compare its R -relatedness with v , via the following *inductive step*: (i) if wRv , then let H contain v such that $w\bar{H}v$; (ii) if vRw , then let H contain v such that $v\bar{H}w$; (iii) if neither wRv or vRw , then leave

⁸¹ This process of “filling a collection” is unconventional. To restore convention, we can say that this “filling” is short for an iterative sequencing on H . That is, we can say that ‘ H ’ is an abbreviation for a sequential *class* of collections $H^1 \dots H^n$ (where $1 \leq n \leq \infty$; and there exists a structure-preserving mapping from each H^n to H^{n+1}). Under this convention, adding elements to H is to be understood as an implicit shift in the denotation of ‘ H ’: that is, my use of the expression ‘let H contain...’ actually means “let ‘ H ’ now designate the next H^n in the H -sequence, which is structurally identical to H except that it contains...”. Doing things this way allows us to avoid the messy metaphysical issues surrounding the claim that collections might change their parts. But, it also forsakes readability.

v out of H (for now) and move on to the next w that is in both O and H, repeating (i)-(iii) until there are no more w's to compare.

3. Repeat 2 until there is no v that is in O but not in H.

From 1-3 we must conclude that any SPO relation can be used to produce a higher than structure (in fact, a hierarchy). What you end up with by going through this process is an internally-linked *higher than* structure that is isomorphic to the SPO-structure you started with.

Although the recursive process used involves finite structures, we can nevertheless generalize the result to those that are infinite. The finite limitation only prevents the *actual* generation of an infinite higher than structure; but that was never my intent. The process shows that for any n, if O is a collection of size n, and R is a strict partial ordering relation linking the elements of O, then there is a collection H with a relation \bar{T} on H, such that R is isomorphic to \bar{T} . Although n is presumed to be finite, we can nevertheless induce an infinitely persistent isomorphism.

Another way to challenge the generality of these arguments is to complain that we began the recursive process with a collection that is linked by the SPO relation. Since we have not begun with any arbitrary (possibly unlinked) set, the result is not fully general: we have only shown that SPO relations that *link* the origin-set can be considered higher than relations.

But this is misleading. All we need to show is that any SPO relation can be used to create a *higher than* structure. If we were to be given an unlinked SPO-ordered collection to start with, we can simply run the algorithm on each of its linked sub-

collections. The result will be the generation of multiple *higher than* structures. If there are disconnected bits in the origin-set, then we simply add those to the *higher than* structure with one sweeping stipulation at the end of the recursion: “Let H contain all the extra disconnected bits!” After which, full generality is restored. Since we can construct multiple *higher than* structures from an SPO relation, then the fact remains: every SPO-structured set can be used to create a (i.e., at least one) *higher than* structured collection. Being a strict partial order, therefore, is both necessary and sufficient to serve as a *higher than* relation.

And so, we can see that every binary relation with the properties of asymmetry, and transitivity can be treated as a *higher than* relation. At last, an analysis of *higher than* has been delivered.

(HIGHER) A relation, R, is a *higher than* relation iff R is a strict partial order (i.e., is asymmetric, and transitive)

In other words: ‘higher than’ is then just another name for ‘strict partial order’. Every SPO is a *higher than* relation: *causes, is ancestor to, is greater than, is less than, is proper part of* are all *higher than* relations, insofar as they are asymmetric and transitive.

We can use this understanding of *higher than* to define a notion of *being at a higher level than*.

(AT-HIGHER) x is at a higher level than y iff y is at some level L, and for every level L such that y is at L, there is a level L* such that x is at L* and L* is higher than L.

This is, of course, merely one way of specifying such a relation. It differs from our HIGHER-LEVEL analysis from Chapter 1, and it might differ from other accounts as well. It is an inessential addition to Core Hierarchy Theory, so I will leave the AT-HIGHER analysis out of it. The hierarchical analyst must make themselves clear what they mean by ‘at a higher level than’ whenever they use the expression. It is enough that the expressive resources of Core Hierarchy Theory facilitate such clarifications. AT-HIGHER is simply an example.

Our informal Core Hierarchy Theory is now complete. There are four questions a general theory of hierarchies qua hierarchies must answer: What is a hierarchy? What is a level? What is higher than? And what is it to be at a level? We have arrived at an answer to all of them.

A hierarchy is a collection of levels linked together by *higher than*. A *higher than* relation is nothing other than a strict partial order. Levels are nothing special: they are simply the relata of higher than. To be “at” a level is to be a part of it that falls within the contextual interest of the hierarchical analyst. That is, the things “at” each level are determined by the parts of that level, plus contextual fiat:

(φ -AT) Something x is (φ -)at a level L iff L is a level, x is part of L , and x satisfies φ .

Now that we have completed our informal analysis of hierarchies, levels, and *higher than*, we are able give it a regimented reconstruction using the relational predicate calculus (Appendix A) or a very basic model theory (Appendix B). I intend the regimented reconstruction to be used by computer modelers, scientists, ontologists, and

philosophers who hope to use hierarchies in their work and to prove things about hierarchy-based theories. Let us consider some objections to Core Hierarchy Theory before we proceed to this reconstruction.

3.0 Objections to Core Hierarchy Theory

For some, this analysis might seem suspiciously simple. I anticipate three main objections. First, critics might insist that the analysis is overly general: that the set of definitions provided imply that intuitively non-hierarchical structures turn out to be hierarchical. Second, it might be objected that the analysis provided implies that *higher than* is arbitrary. And third, it might be objected that the understanding of *level* provided by Core Hierarchy Theory is too deflationary; it becomes an empty notion.

3.1 The “Too General” Objection

The first reaction one might have to the SPO account is that it is overly general. If being an SPO is sufficient for being a higher than relation, then a collection that is linked by the *to the left of* relation counts as a hierarchy. For example, Mary’s being to the left of John. Such collections are not common examples of hierarchies.

In fact, there is empirical evidence which supports the idea that *to the left of* is a hierarchical relation. Consider Mary’s being to the left of John. Replacing ‘to the left of’ with an left-directed arrow ‘ \leftarrow ’, we have:

Mary ← John

Our empirical evidence is acquired from the following procedure, which I rely on you to perform: physically turn this page clockwise by ninety degrees. Voila! We encounter a hierarchy. Although the names are written sideways, we can plainly see that Mary is higher than John in this hierarchy (when maintaining the convention that the *higher than* arrow points to what is higher). We can perform the same procedure with respect to *is to the right of*. Of course, we have not *created* anything by turning the page; we simply change our perspective. The hierarchy, evidently, was there all along.

These are only simple examples. There are others. It is possible to string together a collection of branching hierarchies in a way that does not seem to yield a hierarchy at all. This sort of structure, the horizontal zigzag structure, can be represented as follows:




Figure 12: A Zigzagging Strict Partial Order.

For those who deny that the zigzagging Figure 12 represents a hierarchy, the SPO account will be unsatisfying.

But what argument justifies this denial? There are, in fact, straightforward examples of hierarchical organizations that take the zigzag structure: some terrorist organizations have no singular hierarch (i.e. an agent at the top-level). Instead they have a collection of cells which have contact with, and take orders from, only a handful of higher level coordinators, none of whom answer to anyone in particular. Terrorist

hierarchies, then, are often “short and wide” in the sort of manner depicted in Figure 12. The possibility of extremely complex zigzagging strict partial orders thus does not appear problematic for the SPO analysis of *higher than*, or the Core Hierarchy Theory it informs.

3.2 *The Arbitrariness Objection*

One noteworthy feature of the process used to establish *higher than* as a strict partial order is that an analogous process can be used to establish the same thing with respect to *lower than*; any hierarchy can simply be flipped upside-down. That is the inversion principle from the previous chapter.

This gives rise to another potential problem for the SPO analysis: it appears that *higher than* is an arbitrary relation. This arbitrariness might be problematic in one of two ways. First, if there is no substantive difference separating *higher* from *lower*, that might be taken to straightforwardly entail that there is nothing substantive about *higher than*. It is, one might think, a useless or meaningless relation.

But there is a substantial difference between these two relations: one is the converse of the other. If having a converse means a relation is meaningless, then why bother talking about relations at all? Every binary relation has a converse.

The second way to deride the arbitrariness resulting from Core Hierarchy Theory is to spot a potential reductio: if *lower than* is a strict partial order, and any strict partial order is a higher than relation, then it straightforwardly follows that anything lower than something else is also higher than that something. But that appears to conflict with one of the foundational principles of Core Hierarchy Theory: the incompatibility principle

requires that if Jill is higher than Bob, then Jill is not lower than Bob. Core Hierarchy Theory might then appear to contain an internal contradiction.

But the contradiction is only an appearance. The SPO account tells us that any strict partial order yields a higher than structure, but that does not mean that every strict partial order yields the same *higher than* structure. In fact, throughout the development of the Core Hierarchy Theory, I have been assuming that things stand *higher than* only relative to some particular *higher than* structure; ‘higher than’ can be interpreted in many ways. If, for example, Jill is higher than Bob in one interpretation there could nevertheless be a different hierarchy in which Jill is lower than Bob.

Although it is true that the *lower than* relation is a strict partial order, this only means that the *lower than* ordering can always be reinterpreted as a *higher than* ordering. As long as two levels are not both higher *and* lower than each other relative to the same interpretation of ‘higher than’, the reductio can be avoided. With respect to some hierarchy, H_1 , Bob might indeed be lower than Jill; and yes, since *lower than* is also a strict partial order, there is indeed a hierarchy, H_2 , in which Bob is higher than Jill. But in this case, Bob’s being higher than Jill only holds relative to H_2 . We are dealing with different hierarchies: since their levels are ordered differently, $H_1 \neq H_2$. With respect to H_1 , Bob is still the lower one; with respect to H_2 Bob is the higher one. The reductio can therefore be avoided.

3.3 *The Emptiness Objection*

Two notable aspects of Core Hierarchy Theory are its account of *higher than* and its account of *level*. It was argued that any asymmetric and transitive relation can be

considered a *higher than* relation, and the notion of *level* was characterized at the relation of any such relation. This account of hierarchical levels might be cause for some concern. The concern is not that the account is false, but rather that it is uninteresting.

To think of a level simply as the relation of *higher than* is indeed a deflationary account of level. But it is not the purpose of Core Hierarchy Theory to reveal any deep insights into the metaphysics of hierarchical *levels*. The purpose is only to capture what is meant by ‘hierarchy’ in the simplest and most general terms, and to provide a backdrop for the investigation of more substantive questions about theories and models that utilize hierarchical structures. In fact, there is a risk of including *too much* substance within Core Hierarchy Theory. To give a more substantive analysis of *level* risks making Core Hierarchy Theory overly specific and unnecessarily complicated; the foundational purpose of Core Hierarchy Theory is thereby sacrificed. Staying non-committal on the notion of level is one way of mitigating that risk.

Also note that the account of level inherent in Core Hierarchy Theory is not entirely trivial; it was argued, recall, that levels are sorts of collections (in the mereological sense of ‘collection’). Specifically, they are collections of those things said to be “at” a level, which is partially determined by whoever is offering the hierarchical analysis. Maintaining a deflationist account of ‘level’ within Core Hierarchy Theory allows more leeway for the hierarchical analyst: the more substantive aspects of any theory or model that utilizes hierarchical structure are to remain unspecified in Core Hierarchy Theory. With all this in mind, the triviality concern appears misplaced; the emptiness of ‘level’ is, in fact, intentional and should be viewed as an advantage.

4.0 Consequences of Core Hierarchy Theory

Recall from the introduction that the central thesis of this dissertation involves three claims: that hierarchical Realism is true, that Metaphysical-MLE appears false, and that Reductionism is false. Having formulated a precise notion of what a hierarchy is, we can now put it to work. Core Hierarchy Theory, I will now argue, supports Realism; it implies that hierarchies are indeed to be found in the world insofar as the subjects of interest in ecology actually exist.⁸²

There are other notable consequences as well. I also argue that Core Hierarchy Theory implies that the ‘is a level of’ relation is not generally transitive, intransitive, asymmetric, or irreflexive. These results sound tame, but they are informative.

Lastly, an adherence to Core Hierarchy Theory settles the issue of whether or not *nested* hierarchies, in the sense described by Allen and Starr (1982), Allen and Ahl (1996), Simon (1962), and Webster (1979), are coherent, motivated, and well-suited for a characterization of the ecological hierarchy. Nested hierarchies, I argue, are all these things, since it follows from Core Hierarchy – the simplest and most general account of hierarchies – that every (unrestricted) compositional hierarchy⁸³ is a level of itself, which suggests that every (such) compositional hierarchy *is* nested. I acquire these results using standard philosophical argumentation below, and then I reacquire them more rigorously

⁸² Also recall from the introduction that it is a fundamental assumption of this dissertation that the subjects of scientific discourse (hierarchies notwithstanding) really do exist. Keller and Golley call this assumption “metaphysical-epistemic realism.” (2000, p. 12)

⁸³ Recall that I add the ‘unrestricted’ qualification to pick out only hierarchies in which the contents of levels have not been restricted by φ -constraints imposed on the parthood (i.e., “is at”) relation. This qualification allows that there are compositional hierarchies which are not subject to our results: those hierarchies in which the contents of levels are specified using a φ -part relationship rather than the primitive parthood relation. Recall that these “restricted” hierarchies are non-basic; they are nothing more than unrestricted hierarchies in which some of the contents are ignored by fiat.

within Appendix B, which is a formal reconstruction of the Core Hierarchy Theory developed over the course of the past two chapters. Having shown that Core Hierarchy Theory implies that every (unrestricted) compositional hierarchy is nested, there is nothing at all wrong with the idea that the ecological hierarchy, *qua* a compositional hierarchy, is a nested hierarchy.

4.1 *An Argument for Hierarchical Realism*

Hierarchies, I contend, are collections of levels that are internally linked by a relation with certain formal characteristics – viz., asymmetry and transitivity. If we suppose that there is a reified collection that is linked by a relationship which has these characteristics, then it follows that there is a reified hierarchy.

There are many such collections. Although there is some dispute regarding the properties of the relationship of proper-parthood, the classical characterization takes it to be a strict partial order. Since we are idiomatically using ‘collection’ to pick out those individuals with proper-parts, it follows that composites are collections. Composites are not identical to their proper-parts (Lewis 1991, is convincing when he points out that the composite is a single thing while the parts are many), and proper-parthood is a strict partial order. Core Hierarchy Theory then tells us that there are hierarchies everywhere there is a composite object.

We can also take relative size as an example. *Is-bigger-than* is a strict partial order, and there are plenty of examples of things being bigger than others. Find three books of differing sizes, and remove them from the shelf with one hand. There you are; you have a simple three level hierarchy. You are holding it right in your hand. Now take

a look at your bookshelf. More than likely, each shelf is further from the ground than some other shelf, except for the bottom. *Further-from-the-ground-than* is another strict partial order. And so, again, we see that these objectors to the reality of hierarchies are harboring them in their very own living quarters!

Finally, there is something to be said when it comes to the role of hierarchies within scientific practice. The *is-a-sub-system-of* and *is-a-sub-mechanism* of relation is also a strict partial order. Biological, neural, and medical sciences make frequent use of the concept of a system and of a mechanism (MDC 2000, Glennan 1996, 2002). On all of these accounts, it is important for the mechanistic scientist to describe the entities, activities, properties, and relations that come together to make a sufficient difference to the occurrence of a phenomena.

Some systemic and mechanistic phenomena are not so simply produced. Neurologists must look at relations between cellular and molecular phenomena, as well as their connection to conscious experience. For a patient with renal troubles, she should hope to find a doctor interested in the source of renal failure at many levels. It is important to know whether it was a physical blow, a vascular problem, a sort of poison, or whether it was genetics that was responsible for the kidney problem. Each of these sources corresponds to a different compositional level of the renal system: a karate kick to the kidney is a high-level source; narrowed blood vessels are lower level components of the renal system; poisons are found within those blood vessels; and genetic abnormalities are the lowest source of failure. Knowing the source of the problem is helpful in treating it. This establishes kidney function as a multi-level phenomenon and the renal system as a multi-level system (or mechanism).

We can conclude, then, that Core Hierarchy Theory implies that there are hierarchies to be found in the world. We see it in the standard ontology of parts and wholes; we see it within our scientific practices and explanations, and we even see it in our own homes, should we care to look. Taking basic scientific realism for granted, since there are objects of scientific interest that can be collected such that those collections stand in asymmetric and transitive relations, it straightforwardly follows from Core Hierarchy Theory that there are hierarchies. Of course, it does not follow from Core Hierarchy Theory that *every* hierarchy used within science actually exists, but insofar as ECOHIERARCHY is a collection of real things ordered by composition, its existence is justified by Core Hierarchy Theory.

To challenge this defense of Realism requires one of three things: a skeptical stance towards basic scientific realism, lingering doubts that the objects of scientific interest can be collected such that there is an asymmetric and transitive relation holding between those collections, or a skeptical stance towards Core Hierarchy Theory itself.

I have little to say to those who doubt basic scientific realism. It was listed as one of our fundamental assumptions in the introduction; to defend it takes us beyond the scope of the present work. But for the other skeptical concerns, something must be said in effort to dispel them. For those who doubt that the objects of scientific discourse can be collected and ordered in the manner required by our argument, let us be reminded of the ecological hierarchy in the first chapter. Unless one can object to the “collect and order” process used to characterize ECOHIERARCHY in that chapter, those doubts should be quelled. For those skeptical of Core Hierarchy Theory itself, let us be reminded that it is the result of an extended argument occupying this and the previous chapter. Does the

source of doubt involve the claim that hierarchies are mereological collections of levels, that levels are themselves mereological collections, or that *higher than* is simply a name given to asymmetric and transitive relations? Without a clear reason to cast doubt on the specific arguments given in favor of these three contentious aspects, there is little reason to cast doubt on Realism. Nor is there reason to cast doubt on the direct consequences of Core Hierarchy Theory yet to be mentioned.

4.2 *Transitivity of 'is a level of'*

Let us now turn to some of the other consequences of Core Hierarchy Theory. Since they rely heavily on the notion of a *higher than* relation, let us henceforth designate this relation using the h-bar symbol: \bar{h} . The formal properties of \bar{h} have already been investigated within our hierarchy theory (Appendix A contains further developments), but this still leaves us to inquire into the formal properties of the *is-a-level-of* relation. Is it transitive? Is it symmetric? Is it reflexive?

It is doubtful that the 'level of' relation is generally transitive. To say *x is a level of H* is transitive would be to say that levels of one hierarchy H are, *ipso facto*, also levels of any hierarchy that has H as one of its levels.⁸⁴ For example, if Doug is the lowest member of the accounting department, and the accounting department is the lowest level of the organizational hierarchy, then given transitivity, Doug is the lowest member of the organization. Superficially, the reasoning appears sound.

Although there are contexts in which transitivity is satisfied, there are cases in which it fails. Consider two hierarchies A, and B, and suppose that A has two levels, d

⁸⁴ Where Lxy is 'x is a level of y': for any x, y and z, if Lxy and Lyx , then Lxz .

and e, and that B has three levels f, g, and h. Since the A and B can be collected together and ordered by the *has a greater number of levels than* relation – a strict partial order – the collection of hierarchies A and B is thus itself a hierarchy, H, with A and B as levels: B is higher than A by virtue of B’s having one extra level.

But let us now suppose that the letters ‘a’ through ‘h’ designate things that are not hierarchies: things that have no levels. Since a through h cannot be ordered by the *has a greater number of levels than* relation, a through h are therefore *not* levels of the main hierarchy, H; although they are indeed parts of H. Transitivity thus fails in this case. The ‘is a level of’ relation is therefore neither generally transitive nor generally intransitive.

4.3 *Asymmetry, Irreflexivity, and ‘is a level of’*

We can also see that an asymmetry constraint also fails for the *is a level of* relation. If asymmetry were to hold, then so would irreflexivity. That is, if x’s being a level of H implies that H is not a level of x (asymmetry), then it would immediately follow that x could not be a level of itself (irreflexivity). But, having argued that levels are parts of hierarchies, it can then be shown that there are hierarchies which break the irreflexivity constraint. And so, the *is a level of* relation cannot be generally asymmetric. I describe one such reflexive hierarchy in what follows.

Let UNIVERSE be the total summation of everything there is in the domain of quantification, for some such domain. By ‘summation’ here, I mean to invoke the mereological operation of *fusion*. Summations, or *sums*, are the result of fusion operations on things called *parts*. Although it is contentious whether or not the fusion

operation has a restricted or unrestricted application, we can stipulate that we are beginning with a domain of quantification that permits all of its contents to be fused.

Philosophers distinguish between two types of parts. They distinguish parthood from proper-parthood. Parthood is a simple partial ordering: reflexive, anti-symmetric, and transitive. Proper-parthood, on the other hand, is a *strict* partial ordering commonly defined in terms of the parthood relation (though you can always do things the other way and define parthood in terms of proper-parthood). Proper-parthood is irreflexive, asymmetric, and transitive.

This distinction is important. It allows us to associate a hierarchy with UNIVERSE. Since proper-parthood is a strict partial order, Core Hierarchy Theory implies the existence of a hierarchical structure. And so, since UNIVERSE is the summation of everything in the domain, it follows from Core Hierarchy Theory that there is a hierarchy H with UNIVERSE as the top level of H: UNIVERSE is the terminal relation in a (rather large) chain of proper-parthood.⁸⁵

However, since UNIVERSE is the summation of everything within the domain of quantification, and since we have derived that H is something that “is” (i.e., we are quantifying over it), it follows that H is a part (either proper or improper) of UNIVERSE. But, since levels are parts of hierarchies (as argued in the previous chapter), and Core Hierarchy Theory implies that UNIVERSE is a level of H, it follows that UNIVERSE is part (either proper or improper) of H. And so, we have derived both that UNIVERSE is part of H and H is part of UNIVERSE.

⁸⁵ Or perhaps the bottom level of H, invoking the inversion principle. Whether UNIVERSE is the top or bottom level in H is irrelevant here.

As we are already aware, the proper-parthood relation is asymmetric: UNIVERSE and H thus cannot stand in the proper-parthood relation. So, they must stand in the general parthood relation. Parthood, we know, is anti-symmetric: if x is part of y and y is part of x, then $x = y$. Therefore, it follows that $H = \text{UNIVERSE}$. UNIVERSE, it turns out, is both the top level⁸⁶ of its parthood hierarchy, as well as that parthood hierarchy itself. We can only conclude, then, that it is possible for a hierarchy to be a level of itself. The *is a level of* relation is not generally asymmetric.

4.4 *Compositional Hierarchies are Levels of Themselves?*

What we see above is only a single example. It has not been shown that Core Hierarchy Theory implies that compositional hierarchies are, in general, levels of themselves. We have only shown that *there are* hierarchies that are levels of themselves. But, in fact, we can generalize this argument with respect to any composite object, not just our UNIVERSE toy-example.

First let us take for granted Core Hierarchy Theory: hierarchies are collections whose levels are (primitive) parts of it which stand *higher than* other (primitive) parts of it. Levels are collections whose (primitive) parts are their contents. I use the primitive interpretation of parthood, rather than a φ -part (e.g., functional part, \in -part) interpretation in the spirit of parsimony, but at the expense of full generality: set-theory based hierarchies are beyond the scope of everything that follows.

To generalize the result from the preceding section, we merely crop everything from the domain of quantification except the composite object in question and its parts.

⁸⁶ Or bottom level, if you choose to reason upside-down.

By doing so, the composite object O becomes a UNIVERSE of sorts. Since O has proper-parts, Core Hierarchy Theory implies that *there is* also a hierarchy H^* associated with O , which consists only of O and its parts. Since we are only quantifying over O and its parts, and we have just now quantified over H^* , it follows that H^* is part of O : either a proper or improper-part. Moreover, since O is a level of H^* , and Core Hierarchy Theory implies that levels are parts of hierarchies, O is thus a part of H^* . Since O is part of H^* and H^* is part of O , the anti-symmetry of primitive-parthood entails that $O=H^*$. H^* , is therefore a level of itself. Since we can do this for any composite object, it follows that everything with proper-parts is a level of itself with respect to its unrestricted parthood, or “compositional”, hierarchy.

This, of course, does not imply that *every* possible compositional hierarchy is a level of itself: set-based hierarchies are beyond the reach of this sort of argument. There are no sets in the domain, by stipulation, and so we have shown nothing about set-based hierarchies. The argument applies only to the simplest sort of compositional hierarchy: the sort that consists *only* in the composite and its parts.

Even with this qualification, the result is unexpected and puzzling. Perhaps too puzzling; it might cause concern for Core Hierarchy Theory. The notion of a reflexive level is something that appears *prima facie* incoherent. When reasoning diagrammatically, it appears to entail a contradiction. If H is a hierarchy consisting of levels A , B , and C , such that $H = A$; and if H^* is a hierarchy consisting of levels B and C , and $H^*= B$, then it appears to follow that B is higher than itself (relative to H): it is found within both levels A and B . But, given the asymmetry of $\bar{\text{h}}$ this is impossible. The situation is depicted in Figure 13.

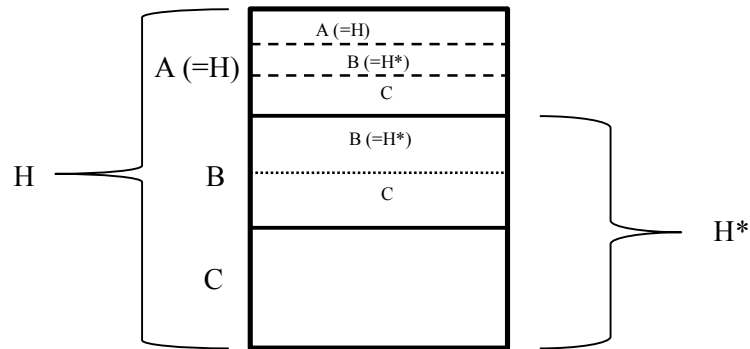


Figure 13: A Reflexive Hierarchy – H has levels A, B, C; H* has levels B and C. B appears to be higher than itself in H.

The diagram is misleading. It is made to appear that that B in level A is higher than B in level B because the diagram portrays two hierarchies stacked one on top of the other, yielding what appears to be six levels in total for H. But, recall from above, the *is a level of relation* is not generally transitive. H only contains three levels: the three levels separated by a solid line in H, which are the same as the three separated by dashed lines in A (Figure 13). The B-level that is designated by short-dashed lines in B in Figure 13, is within another hierarchy altogether: H*, a two-level hierarchy. B, therefore, is not higher than itself relative to H, nor to H*.

Another factor giving rise to the appearance of incoherence is that B is contained in a level higher than itself. But again, the incoherence is only an appearance. Since the higher than relation explicitly relates levels, being *contained* within a higher level than oneself does not entail being higher than oneself. Recall from the first chapter that we drew a distinction between the *higher than* (i.e., *is a higher level than*) and the *is at a higher level than* relations. With that distinction in mind, all the diagram really suggests is that hierarchies can stand in the *higher than* relation with one another (e.g., H is higher than H*), and that this sometimes happens when one is part of the other.

Besides, Figure 13 is not the best way to visualize the situation. Visualizing hierarchies diagrammatically can be useful in simple cases, but this does not mean that a visual paradigm is the best way to investigate the nuances of the structure. In fact, there is a familiar image that depicts the situation a bit better than Figure 13. Consider Figure 14, below.




Figure 14: The Nested Hierarchy

Figure 14, recall, is how we depicted the nested ecological hierarchy, ECOHIERARCHY, in the first chapter. We might then consider the coherence of our original depiction of it to be vindicated. If all basic compositional hierarchies (“basic” in that they consist only of the composite and its parts) are levels of themselves,⁸⁷ viz., the top-level, then a two-level compositional hierarchy will be the second level of itself, a three-level hierarchy the third level, and so on. A three-level basic compositional hierarchy is simply a two-level hierarchy with an extra level on top; or, more generally, an n -level (where $n > 2$) basic compositional hierarchy is a $n-1$ level hierarchy with an extra level on top. Since levels are primitive-parts of hierarchies and primitive-parthood is transitive, those levels that are themselves hierarchies are such that *their* levels are also

⁸⁷ For a proof of this generalization that utilizes an even more general understanding of ‘compositional hierarchy’, see Appendix B. But again, note that the generalization is only intended to apply to compositional hierarchies that utilize the *unrestricted* (i.e., undefined) notion of being “at” a level (viz., parthood).

primitive-parts of the $n+1$ level. Every level of these bare compositional hierarchies, then, is also found at all levels higher than it. For a generalized presentation of this result, see Theorem 4 of Appendix B.

These basic compositional hierarchies, therefore, are nested hierarchies qua top-levels of themselves. Although we acquired this result informally, a formal presentation can be given. In the appendices that follow this chapter, I give two formal reconstructions of Core Hierarchy Theory. The first uses nothing but the resources of the standard predicate logic (Appendix A). I use it to show how one might expand the vocabulary of Core Hierarchy Theory.

The second formulation (Appendix B) is far more compact and useful for theorem-proving. I take it to be the definitive formulation of the unrestricted Core Hierarchy Theory, and I use it to rigorously recover the main results from our (largely informal) investigation. When I refer to Core Hierarchy Theory from this point forward, please note that I am referring to the articulation given in Appendix B, and no other.

5.0 Conclusion

The central purpose of this chapter was to develop and defend a general theory of hierarchies qua hierarchies, as the notion is used within philosophy and the sciences. I began, in the last chapter, with a review of the available theories of hierarchies. All of them were problematic. I began this chapter with a superficial description of the notion of hierarchy that they were trying to capture precisely. Building and refining each aspect of this superficial notion proved fruitful. Not only were we able to repair the problems with

the standard view of hierarchies, the strict partial order view, we were also able to *argue* for it.

The argument for the view came in two parts. First, I argued that asymmetry and transitivity are necessary conditions for a relation to count as a *higher than* relation. Second, I argued that these properties are also sufficient to count as a *higher than* relation; I showed that, given any strict partial order on a set, an isomorphic *higher than* structure can be given via recursion.

The account of hierarchies provided helped to make progress towards the central thesis of this dissertation by providing a compelling reason to favor the hierarchical Realism of Salthe over the Anti-realism of Allen and Starr (under the assumptions specified in the introductory chapter). We argued that since hierarchies are collections of things ordered by a *higher than* relation, and since there are real collections of things related by real strict partial orders – e.g., turtles are related to their proper-parts via a *proper-parthood* relation – it follows that there are, at least some, “real” hierarchies.

To end the chapter, I argued that in certain compositional hierarchies the *is a level of* relation is reflexive; moreover, under the unrestricted interpretation of “is at”, Core Hierarchy Theory implies that all compositional hierarchies (that consist only of the whole and its parts) are levels of themselves. This feature of hierarchies is surprising. But nevertheless, it clarifies and vindicates the coherence of Allen and Starr’s bold characterization of the ecological hierarchy, which is so often misunderstood. It also adds credence to the views of Simon and Webster (see Chapter 2) who draw an analogy between hierarchies of systems and an assortment of nested “Chinese boxes” in which the

larger systems/boxes contain the smaller. (Webster 1979, p. 120) In appendices A and B to this chapter, I formalize Core Hierarchy Theory, and use that formalization to generalize the main results of this chapter as theorems (but with a more general characterization of ‘compositional hierarchy’). The definitive statement of Core Hierarchy Theory, as well as its theorems, can be found in Appendix B.

In the coming chapters I put these theorems to work. In the next two chapters I use Core Hierarchy Theory to argue for the remaining conjuncts of our central thesis. In the next chapter, I discuss what is commonly known as “the layered worldview” with respect to the philosophical positions of Reductionism and Holism (which, recall, is a sort of anti-reductionism). I argue that Core Hierarchy Theory, insofar as we accept it, provides a reason to favor Holism over Reductionism.

Appendix A: A First-Order Reconstruction of Core Hierarchy Theory

To reconstruct our naïve Core Hierarchy Theory (CHT) in formal terms, there is no need to specify which relation counts as *higher than*. We can take it as a primitive, in the same way ‘is part of’ is treated in formal mereology. Mereologists, however, take parthood as primitive because they have no account of it; here we take *higher than* as primitive out of convenience. The main chapter gives reason to think that it is simply the name given to a strict partial order relation of particular interest. We can account for this insight using a second-order extension to CHT, as I demonstrate below.

1.0 Preliminaries

Let \underline{H} be a domain of individuals and h be an interpretation that assigns the symbol ‘ \mathbb{H} ’ to any relation on \underline{H} that satisfies the axioms listed below. Let ‘@’ designate the general parthood relation on \underline{H} . Where ‘ x ’ and ‘ y ’ are variables, and ‘ \mathbb{H}^h ’ denotes an arbitrary interpretation of the *higher than* relation on \underline{H} .⁸⁸ We can express ‘ x is higher than y (on interpretation h)’ as:

$$(P1) \quad x\mathbb{H}^h y \quad \text{(Higher than)}$$

With this primitive, we can define its converse, ‘is lower-than’:

$$(D1) \quad x\mathbb{L}^h y =_{df} y\mathbb{H}^h x \quad \text{(Lower)}$$

We can also use \mathbb{H} to define the ‘is a level of’ relation that stands between levels and hierarchies (D2).

⁸⁸ I use collection rather than hierarchy here since there are *higher than* structures that are not hierarchies: e.g. the broken hierarchy from the second chapter.

$$(D2) \quad Lxh =_{df} \exists y(x\bar{H}^hy \vee y\bar{H}^hx) \wedge x@h \quad (\text{Level})$$

D2 captures the idea that a level (relative to some hierarchy) is a part of that hierarchy that stands in the *higher than* relation with another part of the hierarchy. Recall from the main chapter our deflationary account of levels: levels are simply the relata of *higher than*. The contents of those levels, it was argued, stand in the *is at a higher level than* relation.

The \bar{H} relation is also useful in characterizing the notion a hierarchy itself (D3):

$$(D3) \quad Hh =_{df} \exists xLxh \wedge \forall h^*\forall h^{**}(\forall x((Lxh^* \vee Lxh^{**}) \equiv Lxh) \wedge \neg\exists x(Lxh^* \wedge Lxh^{**}) \rightarrow \exists x\exists y(Lxh^* \wedge Lyh^{**} \wedge (x\bar{H}^hy \vee y\bar{H}^hx))) \quad (\text{Hierarchy})$$

The second conjunct of D3 is simply a formalization of the *linkage* requirement from our naïve account. It ensures that for every way of partitioning the collection into two sub-collections without remainder, those sub-collections are linked by the *higher than* relation.⁸⁹

This \bar{H} relation is also sufficient to characterize the axioms of incompatibility and transitivity:

$$(A1) \quad x\bar{H}^hy \rightarrow \neg xE^hy \quad (\text{Incompatibility})$$

$$(A2) \quad (x\bar{H}^hy \wedge y\bar{H}^hz) \rightarrow x\bar{H}^hz \quad (\text{Transitivity})$$

From this we can, of course, derive the irreflexivity and asymmetry of \bar{H} :

⁸⁹ The second conjunct of D3 can be omitted by those who affirm broken hierarchies as genuine hierarchies. On such a view, a hierarchy is simply a *higher than* structure, rather than a linked *higher than* structure.

(T1) $\neg x\mathbb{H}^hx$ (Irreflexivity)

(T2) $x\mathbb{H}^hy \rightarrow \neg y\mathbb{H}^hx$ (Asymmetry)

And this completes the lexical core of our naïve hierarchy theory, taken as a first-order theory. $\text{CHT}_1 = \text{D1} + \text{D2} + \text{D3} + \text{A1} + \text{A2}$. Given this CHT, it follows from D2, D3 and T1 that:

(T3) $\text{Hh} \rightarrow \exists x\exists y(\text{Lxh} \wedge \text{Lyh} \wedge x \neq y)$ (Pairs)

which says that every hierarchy has at least two distinct levels.

2.0 Extensions to CHT

One advantage of stating our CHT using quantificational logic is that it is easily extended. Although I will not attempt to provide a *complete* theory of hierarchies here, I will suggest a few different ways in which our theory can be extended.

2.1 The Second-Order Extension

As a first-order theory, we must take an interpretation of \mathbb{H} as a primitive. To extend CHT in a way that eliminates the primitiveness, we can give a second-order characterization of \mathbb{H} itself.

(D4) $\exists ix\mathbb{H}^iy \equiv \exists R(xRy \wedge \forall m\forall n\forall o((mRn \rightarrow \neg nRm) \wedge ((mRn \wedge nRo) \rightarrow mRo)))$

In other words, there is an interpretation on which x is higher than y if and only if there is some relation, R , such that x stands R to y , and R is a strict partial order: it is asymmetric

and transitive. This second-order extension is given as, $CHT_2 = D1 + D2 + D3 + D4 + A1 + A2$.

2.2 *Axiomatic Extensions to CHT*

With the predicates comprising the lexical core of CHT – $x\bar{H}y$, $x\bar{L}y$, Lxh , Hh – we have a way to describe a variety of different hierarchies. The variable designating the interpretation of ‘ H ’ will be henceforth suppressed out of convenience. Besides, we have now defined it.

$$(E1) \quad \forall x(Lxh \rightarrow \exists y(Lyh \wedge y\bar{H}x)) \quad (\text{Topless Hierarchies})$$

$$(E2) \quad \forall x(Lxh \rightarrow \exists y(Lyh \wedge x\bar{H}y)) \quad (\text{Bottomless Hierarchies})$$

To describe a hierarchy with a top – that is, a highest level – one merely negates E1. Negating E2 gives us a hierarchy with a fundamental level. And to describe a finite hierarchy, one merely gives the negation of both E1 and E2.

These predicates also suffice to describe the two major sorts of hierarchies: branching and non-branching (linear) hierarchies:

$$(E3) \quad \forall x\forall y((Lxh \wedge Lyh) \rightarrow (x\bar{H}y \vee y\bar{H}x)) \quad (\text{Linearity})$$

$$(E4) \quad \exists x\exists y(Lxh \wedge Lyh \wedge \neg x\bar{H}y \wedge \neg y\bar{H}x) \quad (\text{Branching})$$

3.0 Definitional Extensions to CHT

By extending the list of definitions, we can describe relations between levels and even describe relations that hold between different hierarchies. Let us begin with relations that hold between levels.

$$(D5) \quad CP_{zxy} =_{df} z\bar{H}x \wedge z\bar{H}y \quad (\text{Common Parent})$$

$$(D6) \quad CD_{zxy} =_{df} x\bar{H}z \wedge y\bar{H}z \quad (\text{Common Descendant})$$

$$(D7) \quad CI_{zxy} =_{df} (x\bar{H}z \wedge z\bar{H}y) \vee (y\bar{H}z \wedge z\bar{H}x) \quad (\text{Common Intermediary})$$

In other words, the class of common parents of x and y are all those levels which are higher than both x and y, the common descendants of levels x and y are all those levels which are lower than both x and y, and the class of common intermediaries of x and y are all those levels which form a chain between x and y.

The ‘level of’ relation that holds between levels and hierarchies (or perhaps between two hierarchies) can also be used to define hierarchical correlates of well-known mereological notions:

$$(D8) \quad O_{xy} =_{df} \exists z(Lzx \wedge Lzy) \quad (\text{Level-overlap})$$

$$(D9) \quad U_{xy} =_{df} \exists z(Lxz \wedge Lyz) \quad (\text{Level-underlap})$$

$$(D10) \quad OX_{xy} =_{df} O_{xy} \wedge \neg L_{xy} \quad (\text{Over-cross})$$

$$(D11) \quad UX_{xy} =_{df} U_{xy} \wedge \neg L_{xy} \quad (\text{Under-cross})$$

$$(D12) \quad HC_{xy} =_{df} \forall z(Lzx \rightarrow Lzy) \quad (\text{Hierarchical Containment})$$

Overlap occurs when two hierarchies share a level in common. Underlap picks out situations in which two things are both levels of the same hierarchy. Over-cross describes situations in which two hierarchies have a level in common, but the first hierarchy is not a level of (not contained within) the second hierarchy – they cross-over one another. Under-crossing describes situations when two things are levels of the same hierarchy but the first thing is not a level of the second thing. Lastly, hierarchical containment describes what it is for one hierarchy to be contained within another.

By now it is clear that further extensions to CHT are possible. As noted before, however, fleshing out a complete hierarchy theory is a daunting task, and my purposes here do not require it. It should be obvious, however, that the relational predicate calculus can be used to underwrite a rich, and perfectly general, theory of hierarchies.

There is one deficiency of the first-order approach. The theory we have given ignores the contents of levels. A more robust formal apparatus is therefore needed. I provide one in Appendix B.

Appendix B: Theorems For Unrestricted Hierarchies

The terminology here is idiomatic, but consistent with the rest of the dissertation. Here ‘individual’ picks out those things within the grasp of our singular quantifiers; and ‘collection’ picks out a special sort of individual that has other individuals as proper-parts.⁹⁰

Taking for granted the first order predicate calculus, we can define a model of hierarchical structure H (of kind o)⁹¹ to be an ordered triple that satisfies (i), (ii), (iii), Axiom 1 and Axiom 2 found below:

$$H^o = \langle H, \langle \mathfrak{H}, @, \rangle, o \rangle$$

Where:

- (i) H is a collection (viz., the hierarchy) and every proper-part of H is an individual found in the domain of discourse.
 - a. Let HT^H = the sub-class of all asymmetric and transitive binary relations on H .⁹² Call them the *higher than* relations on H .
- (ii) $\langle \mathfrak{H}, @ \rangle$ is a signature of symbols consisting of, in the first place, the two-place predicate x *higher than* y , and in the second place, the two place predicate *is-at*, to be taken as a primitive-parthood relation.

⁹⁰ Although I take ‘proper-part’ as a primitive within the meta-language; for purposes of proof, I later define proper-parthood using the resources of the object language being developed. Since the definition is the standard one in mereology, this subtle overlap should not lead to any problems.

⁹¹ Take note the difference between a model of a hierarchical structure and a hierarchical structure. A model of a hierarchical structure is an n -tuple consisting of a domain, a signature of symbols, and an interpretation function. A hierarchical structure (in the most general sense) is that which we are modeling.

⁹² Although relations are typically defined on sets, we must alter this conception in order to avoid problems with collections that have no set of parts, for fear of antimony. Instead, let us define a relation on a collection C as a set of ordered pairs (x,y) such that x and y are both parts of C .

- (iii) \circ is an interpretation function, from the symbols in the signature into HT^H :
 - a. $\circ(\mathfrak{H})$ is exactly one element of HT^H , called the *specified ordering relation*.
 - b. $\circ(@)$ is the parthood relation on H , called the *association relation*.

The interpretation function \circ determines what type of hierarchy is being modeled: a parthood hierarchy, an asymmetric supervenience hierarchy, a “has power over” hierarchy, etc.

For those puzzled by (iii)-(b), recall from Chapter 2 that we argued that levels are parts of hierarchies on the grounds that a mereological characterization offers the simplest and most general account of hierarchies *qua* hierarchies. One might think that set-membership is a more familiar way to characterize the relationship, but again recall from Chapter 2 the argument that the set-based characterizations of hierarchies is *not* the most general approach to understanding hierarchical structure: we argued that elementhood is a sort of parthood: a specific and restricted sort of parthood. The purpose of this Appendix is to explore the logical consequences of the most *unrestricted* sort of hierarchy. Nevertheless, the above definition can be used to model any restricted hierarchy one finds favorable; one need only specify the sort of φ -parthood relevant for that restriction. For example, $\circ(@)$ might have set-membership as its value. Of course, the theorems derived here, under the *unrestricted* interpretation, should not be taken to apply to these restricted interpretations of ‘@’.

From the recursion arguments found in the main chapter, we say that any strict partial order is a *higher than* order; and so, for any member of HT^H , there is a corresponding \mathfrak{H} . This does not imply that there is a corresponding hierarchy. Not just

any *higher than* collection counts as a hierarchy. All and only those collections which are \mathcal{H} -linked count as a hierarchy. To formalize this requirement, let us, relative to an arbitrary interpretation, define what it is to be a level of H , viz. the *higher than* collection, and use that definition to specify two axioms: the aforementioned Axiom 1 and Axiom 2.

(Definition 1) $LxH =_{df} \exists y(x\mathcal{H}y \vee y\mathcal{H}x) \wedge x@H$ (is a level of H)

(Axiom 1) $\exists xLxH$ (Existential axiom)

(Axiom 2) $\forall h_1 \forall h_2 (\forall x((Lxh_1 \vee Lxh_2) \equiv LxH) \wedge \neg \exists x(Lxh_1 \wedge Lxh_2) \rightarrow \exists x \exists y (Lxh_1 \wedge Lyh_2 \wedge (x\mathcal{H}y \vee y\mathcal{H}x)))$ (Linkage axiom)

Definition 1 captures the idea that levels are parts of collections which stand *higher than* other parts of that collection. Contents of levels, recall, are not taken to stand in this relation; according to Core Hierarchy Theory, they stand in the *is at a higher level than* relation. The *is at a higher level than* relation is left unspecified; there are a variety of different ways that such a relation can be specified. One's preferred specification can be given as an axiom added to the two provided above.

The existential axiom ensures that the hierarchy has a level. And the linkage axiom requires that for every way H can be partitioned into two non-overlapping sub-collections, h_1 and h_2 , there is always some part x of the first sub-collection and some part y of the second sub-collection such that one stands *higher than* the other. To distinguish structures that satisfy only the first axiom from those that satisfy both, let us say that H° is considered a model of *higher than* structure iff H° satisfies Axiom 1 but not Axiom 2. Keeping with our original definition, let us maintain that H° is a model of *hierarchical*

structure iff H^0 satisfies clauses (i) - (iii), as well as Axiom 1 and Axiom 2. Having defined a model of hierarchical structure, we can finally say, formally, what a hierarchy is: it is the collection found within a model of hierarchical structure.

(Definition 2) for any c , c is a hierarchy iff for some model of hierarchical structure H^0 ,

$$H^0 = \langle c, \langle \mathfrak{H}, @ \rangle, o \rangle$$

Having said what a hierarchy is, there is still something missing. As stated, a model of a hierarchical structure involves some individual, namely, a collection. Pluralities, however, might also be said to exhibit a hierarchical structure, so they too must be accounted for by our core hierarchy theory. Let us say that a plurality, the x s, exhibits a hierarchical structure iff there is some c and a model of hierarchical structure H^0 such that $H^0 = \langle c, \langle \mathfrak{H}, @ \rangle, o \rangle$ and every one of the x s is a part of c . That is, a plurality, the x s, is hierarchically structured when the x s are collected together into a hierarchy. Of course, qua plurality, the pluralistic hierarchical structure is not a hierarchy.

The above is a structural reconstruction and re-articulation of the Core Hierarchy Theory given in the main chapters, plus an extension of that theory to account for hierarchical pluralities. Note that for the present purposes, our ambition is not to develop a complete semantics of ‘higher than’. We only need a precise reconstruction of our hierarchy theory that facilitates deduction, using standard bivalent deductive reasoning. Truth in a model and theoremhood will be defined in a complete semantics for ‘higher than’, but there is no need to add such complications just yet. We already have enough resources to demonstrate some interesting facts about unrestricted-@ hierarchies.

Let us then use our hierarchy theory to demonstrate these facts. To begin, let us use our structural reconstruction of Core Hierarch Theory to prove our simple “pairs” theorem: a constraint on any acceptable theory of hierarchies.

(Theorem 1) For every hierarchy, there are at least two unique levels.

$$\text{I.e., } \forall z(z \text{ is a hierarchy} \rightarrow \exists x \exists y(Lxz \wedge Lyz \wedge x \neq y))$$

Proof: Assume some arbitrary h is a hierarchy. Thus, from Definition 2, for some model of hierarchical structure H^* , $H^* = \langle h, \langle \mathfrak{H}, @ \rangle, * \rangle$. Since H^* is a model of hierarchical structure, by clause (iii) of its definition ‘ \mathfrak{H} ’ is a predicate designating a strict partial order, as determined by the interpretation function, $*$. Since \mathfrak{H} is a strict partial order, it is asymmetric, irreflexive, and transitive. Since \mathfrak{H} is irreflexive nothing satisfies $x\mathfrak{H}x$. Moreover, since H^* is a model of hierarchical structure, Axiom 1 is satisfied with respect to h . Thus, $\exists xLxh$. Existentially instantiating a for x , from Definition 1 we then derive $\exists y(a\mathfrak{H}y \vee y\mathfrak{H}a)$.

Existentially instantiating for y , we get $a\mathfrak{H}b \vee b\mathfrak{H}a$. No matter which disjunct we assume, it follows from Definition 1 that both Lah and Lbh , i.e. that both a and b are levels of h . Given that \mathfrak{H} is irreflexive, it also follows, assuming either disjunct, that $b \neq a$. Existentially generalizing over a and b , we derive $\exists x \exists y(Lxh \wedge Lyh \wedge x \neq y)$. Since this derivation holds no matter which disjunct, $a\mathfrak{H}b$ or $b\mathfrak{H}a$ is true, the derivation must hold within the scope of our original assumption.

Closing our original assumption under conditional proof while universally generalizing over the arbitrary h , we derive our theorem: $\forall z(z \text{ is a hierarchy} \rightarrow$

$$\exists x \exists y(Lxz \wedge Lyz \wedge x \neq y)).$$

QED

Simple enough. To prove the compositional reflexivity result from the main chapter, as well as generalize it, things get a bit more complicated. We must first specify some definitions. Definitions 3 and 4 distinguish two sorts of compositional relationships:⁹³ direct and indirect. Roughly put, a direct compositional relationship is taken to be the proper-parthood relationship itself; whereas an indirect compositional relationship is one that holds between x and y whenever x and y have parts that stand in a proper-parthood relationship with one another. If directness picks out proper-parthood, then indirectness picks out proper-parthood once removed. Definition 5 then combines these definitions into a general definition of a compositional relationship. Lastly Definition 6 specifies what it meant by ‘compositional hierarchy’: a hierarchy for which *higher than* is interpreted as either a direct or indirect compositional relationship.

(Definition 3) For any relation R , R is a direct compositional relationship iff for all x and y , $(xRy \vee yRx) \equiv (x \text{ is a proper-part of } y \text{ or } y \text{ is a proper-part of } x)$.

(Definition 4) For any relation R , R is an indirect compositional relationship iff for all x and y , $(xRy \vee yRx) \equiv \exists z \exists v (z@y \wedge v@x \wedge (z \text{ is a proper-part of } v \text{ or } v \text{ is a proper-part of } z))$ ⁹⁴

(Definition 5) For any relation R , R is a compositional relationship iff R is either a direct or indirect compositional relationship.

⁹³ The standard view of composition, as I understand it: the x s compose y iff each x is proper part of y , no two (distinct) x s properly overlap (i.e. no two x s have a *proper*-part in common), and every proper part of y overlaps (i.e. has a *part* in common) at least one of the x s. Since the notion of composition can be understood entirely in terms of parthood and proper-parthood, these will be the relations (aside from our logical connectives and quantifiers) that we will use to define *compositional relationship*.

⁹⁴ I use existential quantifiers here to ensure that we are dealing with the weakest sort of compositional connection between levels. The definition is intended to capture the idea that an indirect compositional relationship is a direct compositional relationship “once removed”.

Having specified that the type of hierarchical structure (e.g. compositional, supervenience, etc.) is given by the interpretation function o , it is appropriate to define an (unrestricted) compositional hierarchy as follows:

(Definition 6) for any h , h is a compositional hierarchy iff $\exists H^o$ ($H^o = \langle h, \langle \mathcal{H}, @ \rangle, o \rangle$ and H^o is a model of hierarchical structure and for every R , if $o(\mathcal{H}) = R$, then R is a compositional relationship).

In other words: a compositional hierarchy (the ‘unrestricted’ should henceforth be taken to be implicit, for sake of convenience) is defined as a collection found within some model of hierarchical structure where the interpretation of $@$ in this structure is the unrestricted parthood relationship and the interpretation of \mathcal{H} in this structure is a compositional relationship: either direct or indirect. A hierarchy ordered directly will produce a hierarchy where each level is a relata of the proper-parthood relation; if the hierarchy is ordered indirectly, it will produce a hierarchy in which the ordering of the



Figure 15: A Hierarchy Ordered by a Direct Compositional Relationship (left) vs One Ordered by an Indirect Compositional Relationship (right). ‘a’ and ‘b’ designate levels; ‘higher than’ is indicated by a solid arrow (pointing to what is higher); ‘proper part of’ is represented by a dashed arrow (pointing to the whole). The relationship is direct when the arrows are coextensive; indirect when it holds in between things found *within* the levels (i.e. some parts of each level).

levels depends on the proper-parthood relations holding between some of the contents (i.e. parts) of those levels (Figure 15).

We can now prove our compositional reflexivity theorem from the main chapter.

(Theorem 2) Every compositional hierarchy is a level of itself.

I.e., $\forall h (h \text{ is a compositional hierarchy} \rightarrow Lhh)$

Proof: Assume for some arbitrary k : that k is a compositional hierarchy. Then, given Definition 6, it follows that $\exists H^o (H^o = \langle k, \langle \bar{h}, @ \rangle, o \rangle)$, that H^o is a model of hierarchical structure and that for every $R \in HT^k$, if $o(\bar{h}) = R$, then R is a compositional relationship. Existentially instantiating for H^o , we get (1) $H^{\text{comp}} = \langle k, \langle \bar{h}, @ \rangle, \text{comp} \rangle$ and (2) H^{comp} is a model of hierarchical structure, and (3) for every R , if $\text{comp}(\bar{h}) = R$, then R is a compositional relationship. Letting $F = \text{comp}(\bar{h})$, i.e. the relation \bar{h} is interpreted as designating in H^{comp} , it then follows from (3) that F is a compositional relationship.

Open a second assumption (for reductio): $\neg Lkk$. Then, given Definition 1, it follows that (4) $\neg \exists y (k\bar{h}y \vee y\bar{h}k)$ or $\neg k@k$. Parthood, of course, is reflexive, and so k must be part of itself.⁹⁵ Since clause (iii) of our original definition tells us $@$ is parthood, it follows that $k@k$; and therefore, $\neg \exists y (k\bar{h}y \vee y\bar{h}k)$: nothing stands \bar{h} to k and k stands \bar{h} to nothing. Since, we have given the label F to the ‘ \bar{h} ’ in H^{comp} , it follows that $\neg \exists y (kFy \vee yFk)$ and (equivalently) $\forall y \neg (kFy \vee yFk)$. Since

⁹⁵ Some might be tempted to think that the appeal to parthood here is simply a methodological choice, and that ‘ $@$ ’ is better interpreted as *proper-parthood* or some other ϕ -refined notion of parthood (i.e., ϕ -parthood). Since the proof does not work with ‘ $@$ ’ designating proper-part, the proof is problematic. To alleviate such concerns, note that ‘ $@$ ’, as shorthand for ‘is at’ or ‘is in’, must be taken to be reflexive: some hierarchies might have levels without any proper-contents. Consider a hierarchy of mereological atoms. Since mereological atoms have no proper-parts, but are nevertheless levels of the hierarchy, the only things at each level are the levels (i.e., the atoms) themselves. Besides, one of the purposes of Core Hierarchy Theory is to give the simplest and most encompassing (i.e., general) account of hierarchies upon which all other hierarchies can be specified; with that purpose in mind, it does not seem appropriate to retreat to a restricted sense of parthood (e.g., proper-part, functional-part, direct-part, etc.) when characterizing the most basic hierarchical relations.

F is a compositional relationship, from Definition 5 it follows that F is either a direct or an indirect compositional relationship.

(Lemma 1) F is not a direct compositional relationship

Proof of Lemma: Suppose (for reductio) that F is a direct compositional relationship. Then, from Definition 3, it follows that for all x and y, $(xFy \vee yFx)$ iff (x is a proper-part of y or y is a proper-part of x). Universally instantiating k for x and a for y, we now have $(kFa \vee aFk)$ iff (k is a proper-part of a or a is a proper-part of k). Since we already have $\forall y \neg(kFy \vee yFk)$, it follows from universal instantiation that $\neg(kFa \vee aFk)$. From the biconditional we derived a moment ago, it follows from modus tollens and DeMorgan's that k is a *not* proper-part of a and a is *not* a proper-part of k. Since our constant 'a' appeared via universal instantiation within our most recent assumption, we can universally generalize: for all x, x is not a proper-part of k. But, from our undischarged assumption that k is a hierarchy, it follows from clause (i) from our original definition of a model of hierarchical structure that k must have proper-parts. A contradiction! We have established both that k has, and does not have, proper-parts. Closing our most recent assumption under reductio, it thus follows that F is not a direct compositional relationship.

Applying our lemma to the preceding disjunction that F is either a direct or indirect compositional relationship, it then follows that F is an indirect

compositional relationship. From this and Definition 4 it follows that for all x and y , $(xFy \vee yFx) \equiv \exists z \exists v (z@a \wedge v@x \wedge (z \text{ is a proper-part of } v \text{ or } v \text{ is a proper-part of } z))$. Universally instantiating k for x and a for y , we now have that $(kFa \vee aFk) \equiv \exists z \exists v (z@a \wedge v@k \wedge (z \text{ is a proper-part of } v \text{ or } v \text{ is a proper-part of } z))$. Since we have established (just prior to the statement of our lemma) that $\forall y \neg(kFy \vee yFk)$, we instantiate a for y , giving us $\neg(kFa \vee aFk)$. Applying this result to the left hand side of our most recent biconditional, it follows that $\neg \exists z \exists v (z@a \wedge v@k \wedge (z \text{ is a proper-part of } v \text{ or } v \text{ is a proper-part of } z))$; which is equivalent to $\forall z \forall v ((z@a \wedge v@k) \rightarrow \neg(z \text{ is a proper-part of } v \text{ or } v \text{ is a proper-part of } z))$. Universally instantiating for both z and v , results in: $(a@a \wedge k@k) \rightarrow \neg(a \text{ is a proper-part of } k \text{ or } k \text{ is a proper-part of } a)$. Since $@$ is parthood, which is reflexive, it follows that $a@a \wedge k@k$; allowing us to apply modus ponens to our above conditional. Once we do so, we arrive at $\neg(a \text{ is a proper-part of } k \text{ or } k \text{ is a proper-part of } a)$; which straightforwardly implies that a is not a proper-part of k . Since ‘ a ’ was introduced via universal instantiation within the current assumption, we can universally generalize: for all x , x is not a proper-part of k . That is, it follows that nothing is a proper-part of k . But again, given clause (i) of our definition of a model of hierarchical structure, it must be that k has proper-parts. A contradiction! Since, within the scope of our second assumption, we have established that k both has, and does not have, proper-parts, we can thus discharge that assumption and conclude, via reductio, that Lkk .

Now that we have derived the consequent of our theorem under the scope of our only remaining undischarged assumption, we close that assumption under conditional proof and derive: if k is a compositional hierarchy, then Lkk . Since k was chosen arbitrarily, we simply generalize this conditional to derive our theorem: for all x , if x is a compositional hierarchy, then Lxx . That is, every compositional hierarchy is a level of itself.⁹⁶ QED

One interesting corollary of this can be seen by combining Theorem 2 with Lemma 2 below.

(Lemma 2) For everything x that is a level of itself, there is a level that contains everything found on any level of x .

I.e., $\forall x(Lxx \rightarrow \exists y(Lyx \wedge \forall v \forall z((Lvz \wedge z@v) \rightarrow z@y)))$.

Proof: assume (for an arbitrary h) that $\neg \exists y(Lyh \wedge \forall v \forall z((Lvz \wedge z@v) \rightarrow z@y))$.

Exchanging the existential for universal quantifier, we push the leading negation throughout the formula using the appropriate elementary inferences, finally resulting in: $\forall y(Lyh \rightarrow \exists z \exists v((Lvz \wedge z@v) \wedge \neg z@y))$. Universally instantiating h for y , we arrive at: $Lhh \rightarrow \exists z \exists v((Lvz \wedge z@v) \wedge \neg z@h)$.

Assume (for reductio): that Lhh . Then, from modus ponens, we get $\exists z \exists v((Lvz \wedge z@v) \wedge \neg z@h)$. Existentially instantiating for b z and a for v , we derive $Lah \wedge b@a \wedge \neg b@h$. Given that Lah , from Definition 1 it follows that $a@h$. Since $@$ (primitive mereological parthood) is transitive, and we have $b@a \wedge a@h$, it follows that $b@h$. We have derived a

⁹⁶ Recall that ‘compositional hierarchy’ here, and in all other theorems, refers only to unrestricted compositional hierarchies.

contradiction: $b@h$ and $\neg b@h$. Closing our assumption under reductio ad absurdum, it follows that $\neg Lhh$.

Closing our original assumption under conditional proof, we have $\neg\exists y(Lyh \wedge \forall u\forall z((Lvh \wedge z@v) \rightarrow z@y)) \rightarrow \neg Lhh$. Contraposing the conditional, and then universally generalizing over the arbitrary h , we arrive at our Lemma: $\forall x(Lxx \rightarrow \exists y(Lyx \wedge \forall v\forall z((Lvx \wedge z@v) \rightarrow z@y)))$. QED

(Corollary 1) For every compositional hierarchy h , there is a level of that hierarchy which contains the contents of any level of h .

I.e., $\forall x(x \text{ is a compositional hierarchy} \rightarrow \exists y(Lyx \wedge \forall v\forall z((Lvx \wedge z@v) \rightarrow z@y)))$

Proof: the proof is straightforward. Theorem 2 tells us that every compositional hierarchy h is a level of itself and Lemma 2 tells us that for everything h that is a level of itself, there is a level of h that contains everything found on any level of h . From Hypothetical Syllogism, we get our Corollary. QED

We have not yet shown that every (unrestricted) compositional hierarchy is a nested hierarchy and so far we have not concerned ourselves with the direction of the compositional ordering. To proceed to our next step let us eliminate the generality we have enjoyed so far, and turn the investigation to those (unrestricted) compositional hierarchies where wholes are associated with higher levels than their parts. Let us call them *upwards directed* compositional hierarchies. Henceforth letting ‘ \mathfrak{h} ’ designate its own interpretation; they can be defined as follows.

(Definition 7) For any h , h is an upwards directed compositional hierarchy =df h is a compositional hierarchy and either (a) or (b) hold:

(a) $\forall x \forall y (x \text{Th} y \equiv y \text{ is a proper-part of } x)$

(b) $\forall x \forall y (x \text{Th} y \equiv \exists v \exists z (v @ y \wedge z @ x \wedge v \text{ is a proper-part of } z))$

NB: To define a downwards directed compositional hierarchy, we merely replace the order of the variables flanking the ‘is a proper-part of’ in (a) and (b). Wholes are lower than their proper-parts in such a hierarchy.

With this definition we can show that every upwards directed compositional hierarchy is the top level of itself.

(Theorem 3) Every upwards directed compositional hierarchy h is the top level of h .⁹⁷

I.e. $\forall h (h \text{ is an upwards directed compositional hierarchy} \rightarrow (Lhh \wedge \neg \exists x (Lxh \wedge x \text{Th} h)))$

Proof: first, assume k is an upwards directed compositional hierarchy. Then, from Definition 7, it follows that k is a compositional hierarchy and it follows that either (I) or (II) hold.

(I) $\forall x \forall y (x \text{Th} y \equiv y \text{ is a proper-part of } x)$

(II) $\forall x \forall y (x \text{Th} y \equiv \exists v \exists z (v @ y \wedge z @ x \wedge v \text{ is a proper-part of } z))$

Since k is a compositional hierarchy, it follows from Theorem 2 that Lkk .

⁹⁷ A top level of a hierarchy, recall from the main chapter, is a level of the hierarchy for which there is no higher-level.

Now open a second assumption (for reductio): $\exists x(Lxk \wedge x\bar{H}k)$. Existentially instantiating a for x , we derive both Lak and $a\bar{H}k$. Since $a\bar{H}k$ holds, from clause (i) of our original definition of a model of hierarchical structure, a must be part of k : a is in the “domain” of \bar{H} , and since according to clause (i), \bar{H} is a relation on k (where $\langle x,y \rangle$ is a relation on collection k iff x and y are parts of k), it follows that a is part of k . Given clause (iii), it then follows that $a@k$.

Lemma 3: if $a@k$, then either a is a proper-part of k or $a=k$.

Proof: The lemma follows from clause (iii) of our original definition of a model of hierarchical structure and the standard definition of proper-parthood: for all x and y , x is a proper-part of y =df x is part of y and y is not part of x . Let us assume, for conditional proof, that $a@k$. Then, given clause (iii), it follows that a is part of k .

Now assume, for reductio, that a is not a proper-part of k and $a \neq k$. Since a is not a proper-part of k , then by the definition of proper-parthood, either a is not part of k or k is part of a . Since we have already assumed that a is part of k , from disjunctive syllogism it follows that k is part of a . But, since we have established that a is part of k and k is part of a , it follows from the anti-symmetry of parthood (parthood is reflexive, anti-symmetric, and transitive), that $a=k$. That $a=k$ directly contradicts the second conjunct of our most recent undischarged assumption, and so that assumption must be false. Closing it under reductio, it follows that a is a proper-part of k or $a=k$. Closing the remaining open assumption, we get

the lemma we need for the current theorem: if $a@k$, then either a is a proper-part of k or $a=k$. QED

From Lemma 3 and the fact that $a@k$ (a fact established immediately before the statement of our lemma), it then follows that either a is a proper-part of k or $a=k$.

Now open a third assumption that $a=k$. Since we have already established $a\bar{H}k$, it follows that $a\bar{H}a$. But, from clauses (iii-a) and (i-a) of our original definition of a model of hierarchical structure, it follows that \bar{H} is asymmetric. Since the asymmetry of \bar{H} implies the irreflexivity of \bar{H} ,⁹⁸ it follows that $\neg a\bar{H}a$. Since we have derived by $a\bar{H}a$ and $\neg a\bar{H}a$ within the scope of our most recent assumption, we can close our third assumption under reductio: $a\neq k$.

Looking at the disjunction that precedes our most recently closed assumption, it thus follows that a is a proper-part of k .

With the third assumption discharged, let us open a fourth assumption (still within the scope of the second assumption). Assume: $\forall v\forall z((v@a \wedge z@k) \rightarrow v$ is not a proper-part of $z)$. Universally instantiating a for v and k for z , we arrive at $((a@a \wedge k@k) \rightarrow a$ is not a proper-part of $k)$. Since we have already noted that $@$ is reflexive, it straightforwardly follows that $a@a \wedge k@k$. Modus ponens then allows us to derive that a is not a proper-part of k . Since we have established that a is proper-part of k within the

⁹⁸ I suppress the proof because it is obvious: \bar{H} is asymmetric when and only when all x and y satisfy $(x\bar{H}y \rightarrow \neg y\bar{H}x)$. If we assume for some arbitrary a that $a\bar{H}a$, then it follows from the asymmetry definition that $\neg a\bar{H}a$; a contradiction.

scope of our first two assumptions (but no others), we have arrived at a contradiction. Discharging our fourth assumption under reductio, we have $\neg \forall v \forall z ((v@a \wedge z@k) \rightarrow v \text{ is not a proper-part of } z)$, which is equivalent to $\exists v \exists z (v@a \wedge z@k \wedge v \text{ is a proper-part of } z)$.

We now have all the resources we need to complete the proof. Since we have derived, at the beginning of the proof, that either (I) or (II) hold, and we have derived that $a \overline{h} k$ within the scope of our undischarged assumptions, it follows from (I) that k is a proper-part of a . (I), then must be false, since we have already derived that a is a proper-part of k , and proper-parthood is an asymmetric. Since the first disjunct, (I) is false, it must be that (II) is true. That is,

$\forall x \forall y (x \overline{h} y \equiv \exists v \exists z (v@y \wedge z@x \wedge v \text{ is a proper-part of } z))$. Instantiating k for x and a for y , we get $k \overline{h} a \equiv \exists v \exists z (v@a \wedge z@k \wedge v \text{ is a proper-part of } z)$. Since we were able to derive $\exists v \exists z (v@a \wedge z@k \wedge v \text{ is a proper-part of } z)$ from our most recent reductio assumption, it follows that $k \overline{h} a$. Since we derived $a \overline{h} k$ near the beginning of our second assumption, we have then derived both $a \overline{h} k$ and $k \overline{h} a$. Since \overline{h} is a strict partial order, it is asymmetric (i.e. $a \overline{h} k \rightarrow \neg k \overline{h} a$), and we have therefore arrived at a contradiction: $k \overline{h} a \wedge \neg k \overline{h} a$. Closing our second assumption under reductio, we arrive at: $\neg \exists x (L x k \wedge x \overline{h} k)$. Having already established that $L k k$ at the start, we have $L k k \wedge \neg \exists x (L x k \wedge x \overline{h} k)$. Closing our original assumption under conditional proof, it follows that, if k is an upwards directed compositional hierarchy, then $(L k k \wedge \neg \exists x (L x k \wedge x \overline{h} k))$. Since our choice of k was arbitrary, we can generalize to establish our third theorem: $\forall h (h \text{ is an upwards$

directed compositional hierarchy, then $(Lhh \wedge \neg \exists x(Lxh \wedge xTh))$). That is, every upwards directed compositional hierarchy is the topmost level of itself. QED

To obtain the most interesting results, that is, to vindicate the coherence of Allen and Starr's vision of the nested ecological hierarchy by showing that *every* unrestricted upwards directed compositional hierarchy reflects this vision, we must remind ourselves what it means to be upwards-nested.⁹⁹

(Definition 8) h is an upwards nested hierarchy =df for any level x of h , everything at that level is also at every higher level.

$$\text{I.e., } \forall h(h \text{ is upwards nested} \equiv \forall x \forall y \forall z((Lxh \wedge y@x \wedge Lzh \wedge zTx) \rightarrow y@z))$$

With this definition and Definition 7, we can prove:

(Theorem 4) Every upwards directed compositional hierarchy is upwards nested

$$\text{I.e., } \forall h(h \text{ is an upwards directed compositional hierarchy} \rightarrow h \text{ is upwards nested})$$

For some arbitrary k , assume (for reductio) that k is an upwards directed compositional hierarchy and that k is not upwards nested. Then it follows from Definition 7 that k is a compositional hierarchy and that either (a) or (b) hold:

$$(a) \forall x \forall y(xThy \equiv y \text{ is a proper-part of } x)$$

$$(b) \forall x \forall y(xThy \equiv \exists v \exists z(v@y \wedge z@x \wedge v \text{ is a proper-part of } z)),$$

⁹⁹ I henceforth write only about upwards directed compositional hierarchies, since the upward direction is how they are standardly presented, as witnessed in Chapter 1.

and it also follows from Definition 8 that $\exists x \exists y \exists z (Lxk \wedge y@x \wedge Lzk \wedge z\bar{H}x \wedge \neg y@z)$. Existential instantiating a for x , b for y , and c for z , it follows that $Lak \wedge b@a \wedge Lck, \wedge c\bar{H}a \wedge \neg b@c$.

Assume for reductio, that (a) holds. Then, instantiating c for x and a for y we derive $c\bar{H}a \equiv a$ is a proper-part of c . Since we have $c\bar{H}a$, biconditional modus ponens tells us that a is a proper-part of c . Since a 's being a proper-part of c implies that a is part of c ,¹⁰⁰ and from clause (iii) of our original definition, 'part of' is the $@$ relation, it follows that $a@c$. Thus, we have $b@a \wedge a@c$. Since $@$, qua parthood, is transitive, it follows that $b@c$. Since we have derived a contradiction – viz. $b@c$ and $\neg b@c$ – in the scope of our assumption, we can close it and conclude that (a) does not hold.

Since (a) does not hold, but we have derived that either (a) or (b) must hold, it follows that (b) must hold: $\forall x \forall y (x\bar{H}y \equiv \exists v \exists z (v@y \wedge z@x \wedge v \text{ is a proper-part of } z))$. Since \bar{H} is asymmetric, nothing is higher than itself, and so $\neg k\bar{H}k$. Given (b), it then follows that $\neg \exists v \exists z (v@k \wedge z@k \wedge v \text{ is a proper-part of } z)$, which is equivalent to $\forall v \forall z ((v@k \wedge z@k) \rightarrow v \text{ is not a proper-part of } z)$. Since k is a hierarchy, by Definition 2 it follows that there is a model of hierarchical structure $H^o = \langle k, \langle \bar{H}, @ \rangle, o \rangle$. Given clause (i) of our original definition of a model of hierarchical structure, it follows that k has proper-parts.

¹⁰⁰ Given the standard definition of proper-parthood mentioned earlier, a is a proper part of c iff a is part of c and c is not part of a .

Assume for arbitrary m (for reductio) that m is a proper-part of k .

Universally instantiating m for v and k for z , we derive $(m@k \wedge k@k) \rightarrow m$ is not a proper-part of k . Since we have just assumed that m is a proper-part of k , we derive from modus tollens that either $\neg m@k$ or $\neg k@k$. Since $@$ is reflexive, it follows that $k@k$. And so, we have $\neg m@k$. Since x is a proper-part of y just in case x is part of y and y is not part of x , and clause (iii) tells us our interpretation of ‘@’ is parthood, it follows that m is a proper-part of k just in case $m@k \wedge \neg k@m$. On our most recent assumption, m is a proper-part of k ; and so, $m@k$. Since we have derived $m@k$ and $\neg m@k$ within the scope of our most recent assumption, that assumption must be false for any arbitrary m : for all x , x is not a proper-part of k . That is, k lacks proper-parts.

Since we have derived that k has no proper-parts, but we have derived (just prior to our assumption) that k has proper-parts, we must close our last remaining open assumption under reductio, giving us: either k is not an upwards directed compositional hierarchy or k is upwards nested. Converting this disjunction into a material implication that generalizes over our arbitrary k , we derive our theorem: for all x , if x is an upwards directed compositional hierarchy, then k is upwards nested. QED

Having proved Theorem 4, we have shown that our Core Hierarchy Theory entails that every (unrestricted) upwards directed compositional hierarchy¹⁰¹ (a compositional

¹⁰¹ It is important to note here that I am only referring to *unrestricted* compositional hierarchies. Hierarchies constructed from a restricted φ -@ conception of *is at* are not subject to these theorems.

hierarchy in which the @ relation is unrestricted and the wholes reside on higher levels than their parts) is upwards nested, just as Allen and Starr (1982) suggest. Nesting is not only a *coherent* feature of compositional hierarchies, such as ECOHIERARCHY, according to Core Hierarchy Theory it is a *necessary* feature of the simplest and most general sort of compositional hierarchy: the unrestricted compositional hierarchy.

Chapter 4

The Layered Worldview and the Metaphysics of Reduction

1.0	INTRODUCTION.....	190
2.0	THE LAYERED WORLDVIEW	193
2.1	<i>GLOBAL HIERARCHIES</i>	193
2.2	<i>LOCAL HIERARCHIES</i>	194
2.3	<i>GLOBAL VS. LOCAL HIERARCHIES</i>	195
3.0	HIERARCHIES AND THE METAPHYSICS OF REDUCTION	195
3.1	<i>THE “OLD” REDUCTIONISM</i>	196
3.2	<i>THE “NEW” REDUCTIONISM</i>	198
3.3	<i>REDUCTION, REALIZATION, AND HIERARCHY</i>	200
4.0	OPPENHEIM AND PUTNAM ON THE MICROREDUCTION HIERARCHY	201
4.1	<i>FREE-MOLECULES AND OTHER PROBLEMS</i>	204
4.2	<i>THE METAPHYSICS OF REDUCTION – OLD AND NEW</i>	207
5.0	WIMSATT ON THE MICROREDUCTIONIST HIERARCHY	209
5.1	<i>WIMSATT’S INITIAL ASSUMPTIONS ABOUT LEVELS</i>	209
5.2	<i>COMPOSITION, EXPLANATION AND SCALAR ORDERINGS</i>	210
5.3	<i>WIMSATT VS. OPPENHEIM AND PUTNAM</i>	214
6.0	THE NESTING PROBLEM FOR REDUCTIONISM -- OLD AND NEW.....	220
6.1	<i>OPTION 1: THERE IS NO HIERARCHY!</i>	225
6.2	<i>OPTION 2: REJECT CORE HIERARCHY THEORY</i>	226
6.3	<i>OPTION 3: REINTERPRET ‘@’ IN CORE HIERARCHY THEORY</i>	228
6.4	<i>OPTION 4: THE REDUCTIONIST’S HIERARCHY AS A φ-@ HIERARCHY</i>	231
6.5	<i>OPTION 5: ‘HIERARCHY’ IS AMBIGUOUS</i>	232
6.6	<i>OPTION 6: CHALLENGING THE ASSUMPTIONS</i>	235
7.0	CONCLUSION	237

1.0 Introduction

The defense of the existence of hierarchies given so far depends on two things: that there are real collections of things, and that there is a real relation – some sort of strict partial

order – that holds between them. Taking for granted that there are such things (e.g., elephants) and such relations (e.g., composition), it follows that there are hierarchies.¹⁰²

That there are hierarchies supports the characterization of Realism that was specified in the introductory chapter. Realism, recall, is the first component of the view I called *Multi-Grade Holism*. Two components of that view remain undefended. In this chapter I argue the *third* component; I argue that ecological Holism is true by arguing that Reductionism is false under the assumption that Core Hierarchy Theory is our best theory of hierarchies. The distinction, recall, is as follows:

[Reductionism] If the world is hierarchical in the sort of way portrayed within ecology, then wholes are there along with their parts at the bottommost level, whatever that might be.

[Holism] The world *is* hierarchical in the sort of way portrayed within ecology, but the wholes are *not* there along with their parts at the bottommost level, whatever that might be.

In defense of Holism, I argue that Reductionism is incompatible with Core Hierarchy Theory. Insofar as Core Hierarchy Theory is the simplest and most general available theory of hierarchies, Reductionism should therefore be rejected.

I begin with a short discussion of what one might mean by ‘layered’ worldview.

This discussion also provides the backdrop for both this and the next chapter. After laying the groundwork, I differentiate between the epistemological and metaphysical

¹⁰² Note that this sort of argument can only establish that *some* hierarchies are real. It does not justify the stronger thesis that *all* hierarchies of scientific interest are real. Only those hierarchies consisting in real things, ordered by real relationships are subject to the argument; it has not been shown that all hierarchies of scientific interest are ordered by reified relationships.

components of reductionism, as discussed within the philosophy of science, and then focus on the metaphysical component.

The review of reductionist metaphysics, and how it has evolved in tandem with the idea of the layered world, begins with the state-of-the-art. I review Carl Gillett's description of what is called the "new" reductionism (or *compositional reductionism*), which picks out the hard core of the reductionist metaphysics that has remained constant throughout its evolution: the claim that everything needed to account for any scientific phenomena is, in fact, found at the bottommost level of the world. In support of Gillett's analysis of metaphysical reductionism, I review two classic writings *on* the layered worldview itself: Oppenheim and Putnam's classic reductionist analysis, and William Wimsatt's (critical) refinements of it.¹⁰³ Both conform to Gillett's analysis of reduction, as well as the characterization of Reductionism found above, despite minor differences in their view of the layered world.

An examination of the interface between the hard core of Reductionism and the Core Hierarchy Theory developed in the previous chapter yields results that do not favor Reductionism. I argue that Core Hierarchy Theory is incompatible with Reductionism with respect to the most basic sort of hierarchical structure. I conclude the chapter by considering some Reductionist attempts to challenge, or modify, Core Hierarchy Theory. Each of the available options faces serious hurdles; overcoming the challenge posed by Core Hierarchy Theory is not a simple task.

¹⁰³ Note that Wimsatt should not be taken to endorse reductionism in the classic sense. In his 1976, he only means to describe an interpretation of reductionism that a traditional *antireductionist*, such as himself, can sympathize with. He ends the paper by writing, "If this constitutes a reduction of the mental to the physical, then so be it!" (p. 263)

2.0 The Layered Worldview

The layered worldview has a number of proponents, and has served as an implicit backdrop for a number of philosophical discussions including, but not limited to: the Unity of Science movement (Putnam and Oppenheim, 1958; Fodor 1974), the metaphysical reduction of special science theories or phenomena to microphysical theories (*ibid.*; Wimsatt 1976; 1994; 2007), the specific relationship between the mental and the physical (Kim 1998; 1999; 2002; Lycan 1995), and more recently, discussions surrounding the role of multi-level mechanisms within the life sciences (Craver 2001, 2002, 2007; Craver and Bechtel 2007; Glennan 2010).

Although it serves as a backdrop for these philosophical debates, the claim that the world is layered is ambiguous. One might mean a number of different things. One might mean to say that the entire universe is some sort of hierarchy; or one might mean to say that *everything* found within the universe is a hierarchy; or even weaker, that only *some* things found within the universe are hierarchies. The difference between these approaches has been made clear by Jaegwon Kim (2002). He distinguishes between *global* and *local* hierarchies.

2.1 Global Hierarchies

On the first disambiguation one would be asserting that the universe itself can be partitioned into a set of levels, and everything found within the world falls somewhere within this hierarchy. For example, a hierarchy of mereological objects with, perhaps, the entire universe as the terminal object at the top (Simons 1987, Lewis 1991, Varzi 1999); a hierarchy of properties, ordered by some asymmetric supervenience relation; a

hierarchy of events (e.g., where atomic events reside lower in the hierarchy than molecular events); a hierarchy of causal relations; a hierarchy of laws; and a hierarchy of scientific branches. Such hierarchies are the result of a *de dicto* reading of ‘the layered world’. Let us call this view *globalism*; and let us call these *de dicto* hierarchies *global hierarchies*. Oppenheim and Putnam (1958) and Lycan (1995) are among most prominent defenders of this sort of approach within philosophy;¹⁰⁴ but, in fact, this view appears to enjoy implicit acceptance within much of science.

2.2 *Local Hierarchies*

The remaining disambiguations are to be associated with the view that the universe contains hierarchies, without making the further claim that the universe itself is hierarchical. That is, individual objects and systems found within the universe – a tree, an ecosystem, a toilet, an alarm clock, to name but a few – might be considered hierarchical, but the universe is possibly not such a system. Such hierarchies are picked out by a *de re* reading of ‘layered world’. Call this view *localism*; and call these *de re* hierarchies *local hierarchies*.¹⁰⁵ Defenders of the local view include Jaegwon Kim (2002) and Carl Craver (2007).¹⁰⁶

¹⁰⁴ Lycan’s (1995) partitioning of things into levels is similar to the division found in Oppenheim and Putnam’s work. The most striking difference between them is that Lycan orders his levels along a continuum (rather than a predetermined set of discrete levels) by virtue of his notion of *levels-of-functional-abstraction*.

¹⁰⁵ Note that there is no inconsistency in holding both localist and globalist views. They are not in opposition. Localism is simply a weaker sort of hierarchical realism. Indeed, in a world containing only one global hierarchically organized object, the two views are equivalent.

¹⁰⁶ It should be noted that localism comes in two different varieties. On the stronger view, *everything* falls into a hierarchy of some sort or another. On the weaker view, one is only committed to the idea that some things are hierarchical. Note that Kim and Craver are only committed to the weaker view.

2.3 *Global vs. Local Hierarchies*

Within philosophy, globalism is thought to be contentious. William Wimsatt (1976), Jaegwon Kim (2002), among others, have argued for positions which suggest that a neatly ordered global scientific hierarchy is highly implausible, and does not fit well with actual science. Each of these authors, however, have a specific sort of global hierarchy in mind. Wimsatt (p. 252) synthesizes a number of different approaches to the global explanatory hierarchy: scalar hierarchies, compositional hierarchies, and hierarchies in which a variety of different scientific kinds populate levels. Kim concerns himself with a global supervenience hierarchy, where (roughly speaking) classes of events are taken as levels and are ordered by a relation of asymmetric supervenience.¹⁰⁷

Localism, on the other hand, appears to be relatively uncontentious. It is the sort of approach found in the writings (cited above) of Jaegwon Kim, Carl Craver, and Stuart Glennan. We will discuss their work in the next chapter. Here we focus on the layered worldview within the metaphysics of reduction.

3.0 Hierarchies and the Metaphysics of Reduction

Reductionism,¹⁰⁸ as it was characterized at the outset, involves both an epistemological component and a related metaphysical component. This bipartite aspect of reductionism is not unique to that characterization; in fact, all of the philosophical work on reduction

¹⁰⁷ Although Kim is a reductionist, we will leave him out of the discussion until the next chapter, where we discuss realization hierarchies.

¹⁰⁸ Recall that I use the uppercase 'R' to designate the characterization of reductionism given in the introduction; the lowercase is used to designate the basic notion of reduction that requires further philosophical analysis: the notion that Nagel (1961) and Schaffner (1967) are concerned with.

confronts this mixture of metaphysics and epistemology. In this section, I review both the classic efforts to characterize the epistemic component, with an eye for its metaphysics, as well as the more contemporary focus on the metaphysical component.

3.1 The “Old” Reductionism

The *epistemological* component of reduction, recall, is the idea that “lower level” knowledge (e.g., of the parts) suffices for “higher level” knowledge (e.g., of the whole). Much effort has gone into making this epistemic thesis more precise. Kemeny, Oppenheim, and Putnam (Kemeny and Oppenheim 1956; Oppenheim and Putnam 1958), for example, held that reduction is primarily a relationship between theories: theory T2 reduces to theory T1 iff (a) the vocabulary of T2 contains terms not in the vocabulary of T1, (b) any observable data which can be explained by T2 is explicable by T1, and (c) T1 is at least as well systematized as T2.¹⁰⁹

Ernest Nagel (1961), by contrast, omits the reference to observation data but retains the idea that reduction is a relationship between scientific theories. For Nagel, reduction involves only the logical relationship between the *laws* of the respective theories: T2 reduces to T1 iff the laws of T2 can be logically derived from the laws of T1. If the vocabulary of T2 contains terms not found in T1 (e.g., ‘gene’ is foreign with respect to the language of organic chemistry), then the reduction requires an appeal to a

¹⁰⁹ They also articulate a derivative understanding of reduction: as a relationship between scientific *branches*, where branch B2 reduces to B1 (at time t) iff there is some T1 in B1 (at t) such that T2 is reduced to T1 (at t). Since they take scientific disciplines to be composed of theories, an incomplete branch-reduction in which only some of the theories in B2 can be reduced is considered a *partial* reduction. (1958, p. 5)

“bridge principle” that connect T2’s special terms to (perhaps, a complex construction of) predicates found in the language of T1.¹¹⁰

Despite their differences, these two proposals share a common ground: the idea that the phenomena investigated within one science are to be explained by another, more basic, science. To use a hierarchical locution, the idea is that all of the so-called “higher level” phenomena are ultimately explicable by the behaviors and properties of the individuals found in the domain of “lower level” sciences.¹¹¹ Where those individuals are the entities described by microphysics, the sort of reduction is called a *microreduction*, though not all uses of the term ‘microreduction’ involve an appeal to microphysics.¹¹²

The *metaphysical* component of reductionism has received little attention by comparison. Rather than involving relationships between theories and theoretical terms, metaphysical (sometimes called *ontological*) reduction involves certain relationships between individuals and between properties: those things and properties in the domain of “higher level” sciences are either identical to or composed of (or can be eliminated in favor of) the things and properties in the “lower level” domain.

Although their concern is not exclusively metaphysical, Oppenheim and Putnam (1958) and Wimsatt (1976) are pioneers in the investigation of metaphysical

¹¹⁰ These bridge principles might take the form of a biconditional or identity claims (see Causey 1972, 1976 and Enc 1976). Schaffner’s (1967) account of reduction is similar to Nagel’s in that they both make use of derivation relations between the reduced and reducing theories. The main point of difference is that Schaffner adds the requirement that *all* of the statements – not just the laws – of T2* are derivable from T1 (where T2* is the *corrected version* of T2; that is, it is T2 appropriately modified such that it contains no false statements).

¹¹¹ This epistemic thesis pairs nicely with a *methodological thesis* that explanations in higher level sciences should proceed, in actual scientific practice, by examining the properties and behaviors of entities residing at lower levels; e.g., to explain the properties of the whole, one investigates the properties of the parts.

¹¹² E.g., Oppenheim and Putnam’s notion of microreduction allows that it might hold between two branches of science, neither of which are microphysics. Their definition of ‘microreduction’ is given in §3.2 below.

reductionism. Using the notion of a hierarchy as a foundation, they investigate the metaphysical and epistemic components as a mixed issue. Wimsatt (1976), for example, takes reduction to involve an explanatory relationship between a lower level theory (or lower level phenomena) and higher level phenomena (1976, p.224); the phenomenon of fire, for example, is explained by oxidation, a process described by, and found within the domain of, theoretical chemistry. Where Oppenheim and Putnam straightforwardly maintain that the ordering between reductive levels is a compositional ordering, Wimsatt entertains a variety of ordering options – size (p. 237), mereological composition (p. 243), or direction of explanation (pp. 245-52) – eventually settling on composition as being the most intuitive and least problematic. In any case, these metaphysical approaches clearly diverge from the sort of reductionism that Nagel had in mind.

3.2 *The “New” Reductionism*

Contemporary reductionists are in general agreement that Nagel’s classic epistemic approach to reduction is deeply flawed. The problems are well-known. Classic reductionism requires either attribute-identities between the properties picked out by the two theories (Causey, 1972; Enc, 1976), or at the very least attribute-coextensions between the predicates found within the two theories. The multiple-realizability arguments of Jerry Fodor (1974) and Hilary Putnam have convinced many that this sort of reductionism should be abandoned; giving rise to the most popular contemporary metaphysic for science, *non-reductive physicalism* – the view that special science properties (or their instances) are not identical to micro-physical properties, but are rather *realized* by them.

Accounting for this “realization” relation has been the subject of intense debate within the non-reductivist’s camp (see Melnyk 2003; Gillett 2002, 2007; Shapiro 2008; Shoemaker 2007, 2011). Contemporary reductionists look on in interest. They too are eager to embrace this notion of realization. Reductionists, such as Kim, see realization as a way to reincarnate the spirit of classic reductionism into a new, strictly metaphysical, sort of reductionism. Carl Gillett (2007) calls this “new” reductionism *compositional reduction*.¹¹³

According to this view, special science *terms* are not redefined using terms from any micro-based science. In this respect a compositional reduction is not Nagelian. What makes compositional reduction a genuinely metaphysical sort of reduction is that it retains the idea that there really is *nothing more* than the entities and properties of micro-physics, and configurations thereof. Composition does not add anything to the world, says the compositional reductionist, and neither does realization.

The austere realism espoused by the compositional reductionist should not be confused with an extreme eliminativism. There is a difference. The new reductionist retains and respects higher level explanations and descriptions, whereas extreme eliminativists approach the special sciences armed with a conceptual machete; they hack away at special sciences as if they were jungle-vines, doing away with macro-level descriptions and theories altogether. Eliminativists label them “folk” descriptions, and

¹¹³ Note that Gillett’s characterization differs somewhat from Kims. Kim’s (1999, p. 15) “reduction by functionalization” strategy upholds the classic reductionist picture by identifying special science property instances with micro-based based property instances. The difference between Kim and Gillett’s respective presentations makes no difference to our purpose here.

argue that the “folk” language is more or less empty, ingenuine, and about nothing in particular.¹¹⁴

According to the compositional reductionist, the predicates and other terms of our special sciences are in fact genuinely *about* micro-based properties and entities, and configurations thereof: if reality is exhausted by configurations of micro-particles, then insofar as ‘ecosystem’ picks out anything at all, it picks out some of those configurations. While the language of special science might rightly be criticized as a clumsy tool, delivering vague descriptions and *ceteris paribus* explanations of the world at best, this language is nevertheless considered indispensable by the new reductionist.¹¹⁵

In short, the thesis of compositional reductionism is this: configurations of fundamental individuals and properties are all that exists, and they explain everything without having to “reduce” (i.e., define) our indispensable special science *terms* to those of micro-based science. The “new” compositional reduction therefore upholds the characterization of Reductionism offered in the introductory chapter: if the world is layered into a hierarchy (e.g., as we see in ecology), then everything to be found in that hierarchy falls within the level of fundamental science.

3.3 *Reduction, Realization, and Hierarchy*

The advances in the philosophy of reduction since Nagel seem to have narrowed the gap between reductionism and anti-reductionism: anti-reductionists are steadfast in their

¹¹⁴ I am unsure whether or not anyone holds such an extreme eliminativism. But that is all the more reason not to confuse it with Reductionism.

¹¹⁵ For discussions on the indispensability of macro-level predicates see Gillett (2007 p. 214) and Kitcher (1984).

commitment to non-vacuous special science predicates and the macroscopic individuals in their extension; the new reductionists, on the other hand, also say that special-science predicates have non-vacuous extensions, except they say that their extensions consist entirely in things from the microphysical domain.

It should now be clear that the sort of reduction under dispute is best considered a sort framed by decomposition (i.e., parthood) relations. By virtue of the appeal to composition, the debate then occurs within a hierarchical framework: since the collection of micro property instances and individuals and the collection of macro property instances and individuals are linked by asymmetric and transitive parthood relationships (within the reduction-neutral *realization* framework mentioned above), they are therefore linked by a *higher than* relation; and that is sufficient to make the macro and micro collections *levels* of some hierarchy. Call it, the *microreduction* hierarchy.

The relationship between reductionism and hierarchies is well-known. When Oppenheim and Putnam first delivered their grand vision for the unification of science under the reductionist paradigm, they delivered it within an explicitly hierarchical framework. Having decided that the “new” reductionism falls within the characterization of Reductionism given in the introduction, let us now look back to the old.

4.0 Oppenheim and Putnam on the Microreduction Hierarchy

Although we see above that a hierarchical framework underwrites the metaphysics of the *new* reductionism, by virtue of its compositional formulation, it is in fact a very *old* idea. Even prior to Nagel’s work on inter-theoretic reduction, Oppenheim and Putnam make

explicit appeals to what we called the microreduction hierarchy as a way of making precise the notion, “unity of science”.

Oppenheim and Putnam, like the new reductionists discussed by Gillett, take the essential metaphysical feature of a (micro)reduction to be parthood, or more specifically, decomposition. They tell us that a scientific discipline or theory B1 is a microreduction of B2 iff B2 is reduced to B1 and the objects in the domain of B2 can be decomposed into parts, all of which belong to the domain of B1. (1958, p. 8)

Having framed their notion of microreduction in terms of composition, Oppenheim and Putnam notice that the *microreduction* relation has the following properties: it is transitive – if B1 microreduces B2 and B2 microreduces B3, then B1 microreduces B3 – and asymmetric – if B1 microreduces B2, then not vice-versa.

These properties, recall from the previous chapter, are indicative of a hierarchical relationship. Oppenheim and Putnam are perfectly aware of this and adopt a hierarchical framework for science without hesitation. After designating *microreduces-to* as *higher than*, they describe six levels, as well as six necessary conditions for a system of scientific levels. The six levels are (from the bottom going up): elementary particles, atoms, molecules, cells, multi-cellular living things, and lastly, social groups. The six necessary conditions are as follows:

- (1) There must be several levels.
- (2) The number of levels must be finite.
- (3) There must be a unique lowest level.

- (4) Anything of any level that is not the lowest must have a decomposition into things belonging to the next lower level.
- (5) Nothing on any level should have a part on any higher level.
- (6) Level selection must be justifiable from the standpoint of empirical science.

They do not take these propositions to describe general principles about hierarchies; but still, with exception of (1) and (6), their list is quite contentious. Ned Block (2003) and Jonathan Schaffer (2003), for example, discuss the possibility of a world with *no* fundamental level; a world of infinite descent calls into question both (2) and (3). Kim (2002) notes that (5) becomes problematic as soon as a key principle, the *downward-inclusion principle* (DIP) is introduced: “Any whole which possesses a decomposition into parts, all of which are on a given level, will be counted as also belonging to that level.” (pp. 9-10) In other words,

[DIP] Each level includes all higher levels.

DIP might seem odd at a glance, but in fact it is *central* to reductionist metaphysics. Insofar as composition and realization does not add anything to the world (a claim needed to uphold the idea that reality is exhausted by configurations of micro-particles), each level of the reductionist’s compositional hierarchy then consists in the individuals and instanceables that are characteristic of that level, as well as all “higher level” individuals and instanceables that are composed of (or realized by) those “lower level” individuals and instanceables. To use the hierarchical locution endorsed at the outset, if x is composed of the ys, then x is “nothing above” those ys.

The problems mentioned by Block and Schaffer are straightforward. We need say nothing more about them. Kim's criticisms are not so simple. Let us now examine them more closely.

4.1 *Free-Molecules and Other Problems*

Kim (2002) identifies a number of tensions within the preceding proposal. One problem, he says (p. 13 n. 18), is that DIP, which implies that I am located at all levels lower than my own, implies that I am found at the atomic-level. And so, it implies that I have parts at higher levels than mine: my molecules and cells are at a level higher than the atomic level at which I am said to reside. (5) must then be false, since it requires that I have *no* parts at higher levels than those at which I am found.

Of course, there is a simple reply to Kim. The highest level at which something is found is designated by Oppenheim and Putnam as the *proper*-level of the thing in question, giving us a way to accommodate Kim's insight: simply take (5) to mean (5').

(5') Nothing on any level has a part at any level higher than its proper-level.

Even with the adjustment, there is still a problem for those who want to read (5') as a general hierarchical principle, rather than as something peculiar to the microreduction hierarchy. The ecological hierarchy, an upwards-nested hierarchy in which everything that is part of the atomic level is also part of all higher levels, presents a simple counterexample to (5'): if everything is at the biospheric level, then every ecosystem has proper-parts on levels higher than its proper-level. (5'), then, is at best a necessary condition that is idiomatic to the Oppenheim and Putnam microreduction hierarchy. To restore generality, (5') should be replaced with

(5'') Nothing on any level L has a part p such that the proper-level of p is higher than the proper-level of L.

With respect to the upwards nested ECOHIERARCHY, we might say that the proper-level of a thing is the *lowest* level at which it is located, which would allow (5'') to apply to both downward and upward nested hierarchies.

Moving on to (4), notice that there is an ambiguity between two readings: complete vs. partial decomposition. Kim (2002) argues that the former is unacceptable. Since some of my molecules are not part of any of my organs (Kim, p. 14, calls them “free-molecules”), but they are nevertheless a part of me, I do not completely decompose into the next-lower level. To help Oppenheim and Putnam escape the problem, Kim suggests a modification to (4):

(4') Anything of any level L that is not the lowest must have a decomposition into things belonging to levels lower than L.

Alternatively, one might respond by appealing to *partial* decomposition. Instead of requiring that there is a complete decomposition, one might require that

(4'') Anything of any level that is not the lowest must have *some* proper-parts at the next lower level.

Both (4') and (4'') are consistent with DIP; but an advantage of (4'') over (4') is that (4'') rules out alternative orderings. As Brian Epstein (MS) points out, nothing implied by (4') rules out social groups as being lower than molecules, since both have complete decompositions into things found at levels lower than both molecules and social groups.

(4''), however, rules out this inverted ordering: since molecules do not have any social groups as parts, molecules cannot be higher than social groups in the hierarchy if one accepts (4'') as a necessary constraint.

Another problem for Oppenheim and Putnam's hierarchy is that it leaves no way to position certain aggregates within the hierarchy. For example, a-man-in-a-phonebooth consists both in something from the molecular level (viz., a phonebooth) and the multicellular level (viz., a man). Putnam and Oppenheim dismiss such examples as having no noteworthy place in the reductionist's hierarchy. As long as the explanatory relevance of aggregations which span multiple levels can be captured at the lowest level, a tidier hierarchy with a special place for such aggregates is unnecessary. The man in the phonebooth, taken as a singular unit, is found at the bottommost level; and that is really what matters to Oppenheim and Putnam. Otherwise, they could simply respond that gerrymandered aggregates are always found at all levels at which all of the listed parts are found.

Although either response is sensible, I think the first glosses over a subtle but very important point: if everything non-fundamental is indeed composed, as the reductionist appears to presume it is, then one would expect to find an appropriate place for it in a global compositional hierarchy. Hiding one's skeletons in the basement, after all, is no better than the proverbial closet. Notice, for example, that the upwards nested ECOHIERARCHY from the first chapter clearly has what is needed to sensibly position the man-in-the-phonebooth within the ecological hierarchy: any object composed out of things within the ecological hierarchy is "at" any level of that hierarchy at which *all* of its parts are found. In ecological terms, the man-in-phonebooth is an organism bounded by a

pitiful ecological gradient, so it is found at the organism, and all higher, levels. This, I take it, is an independent point in favor of upwards nested hierarchies, such as ECOHIERARCHY, over those which are downwards nested.

4.2 *The Metaphysics of Reduction – Old and New*

The similarities between Oppenheim and Putnam's form of reduction and Gillett's description of the new reductionism are quite striking. The compositional understanding of the metaphysical aspect of reduction is endorsed by both. And both versions take reality to be found entirely in the "basement", so to speak; that is, DIP, as stated above, is retained within the new reductionism.¹¹⁶ As Oppenheim and Putnam (1958) write,

it seems very doubtful, to say the least, that a branch B2 could be reduced to a branch B1, if the things in the universe of discourse of B2 are not themselves in the universe of discourse of B1 and do not possess a decomposition into parts in the universe of discourse of B1. ("They don't speak about the same things.") (p.8)

These similarities are striking, but what of the difference? Since the new reductionist is no less apt to describe things in terms of higher and lower levels, Oppenheim and Putnam's list of hierarchical constraints becomes a potential point of difference between the new and classic views.

Some of the list is plausibly retained. One of the uncontentious claims, namely (1), is trivial; and the other, (6), simply restricts the levels to those whose contents are

¹¹⁶ I fail to see how DIP can be denied by any sort of reductionist. If one thinks, as the reductionist does, that composition is an ontological free-lunch, then the intervening levels between the macro and the fundamental level should contain all levels higher than them, for sake of consistency. I know of no reason to think that the macro-level should jump straight to the bottommost level, bypassing all intermediary levels.

determined by their involvement in some nomic or natural kind. The new-reductionist is not likely to reject either of these claims. (2) and (3) are both available for dispute, although such a dispute strikes me as being subject to empirical investigation. (5'), on the other hand, appears to follow from the asymmetry of the (compositionally guided) microreduction relation, and is unlikely to be rejected.¹¹⁷

(4''), on the other hand, would be rejected if the new reductionist found reason to think some decompositions might skip levels entirely (i.e., there is no partial decomposition into the next lower level); however, it is unlikely that (4') would be rejected since it trivializes the decomposition. Not much differentiates Oppenheim and Putnam from the new reductionism in the end. The main difference, it seems, is in the new-reductionist's insistence on the indispensability of higher level language, as well as the noticeable absence of a clear delineation of levels in Gillett's (2007) discussion of the new-reductionism.

To summarize: Oppenheim and Putnam take a hierarchy to be an asymmetric and transitive higher than ordering of elements, where 'higher than' can receive a number of possible interpretations; for example, x microreduces y. In this respect, Oppenheim and Putnam defend a version of our Core Hierarchy Theory from the previous chapter.¹¹⁸ The levels include individuals involved in significant nomic kinds, and are finite in their number with some level being the bottommost "basement" level. Each level is connected to its predecessor by virtue of (at least) partial decomposition. What makes Oppenheim and Putnam's hierarchical metaphysics *reductive* is no different than from the new-

¹¹⁷ The same goes for (5').

¹¹⁸ In fact, the Core Hierarchy Theory developed in the previous chapter has its origins in a preliminary investigation of the similarities one finds between Oppenheim and Putnam's account of hierarchies, and those of Kim (2002), Craver (2007), as well as much older views, such as that of Conger (1925).

reductionists. They accept DIP, or at the very least the weaker claim that everything in the hierarchy is found at the bottommost level: the domain of fundamental physics.

5.0 Wimsatt on the Microreductionist Hierarchy

Although Oppenheim and Putnam might have been the first to give the reductionistic hierarchy a rigorous treatment, other philosophical accounts have been given. William Wimsatt (1976; and again in 2007), for example, conducts what might well be the most thorough investigation into hierarchical science as a supplement to his investigation of reduction. His verdict: there is no *global* reductionistic hierarchy in the classic sense, but rather a quasi-hierarchical tree of explanation and composition.

5.1 Wimsatt's Initial Assumptions About Levels

At the outset of his investigation, Wimsatt (1976) lists his assumptions about levels:¹¹⁹

I will assume that being at a given level is a property primarily of things in the world: phenomena, objects, properties, processes, causes and effects, etc., and derivatively of linguistic things relating to them: descriptions, law-statements, theories, predicates, etc. Intuitively, one thing is at a higher level than something else if things of the first type *are composed of* things of the second type, and at the same level with those things it interacts most strongly and frequently with or is capable of replacing in a variety of causal contexts. (2007, p. 215)

¹¹⁹ Years after stating these assumptions, Wimsatt is still without an analysis: "The notion of a compositional level of organization is presupposed but left unanalyzed by virtually all extant analyses of inter-level reduction and emergence." (2007, p. 203)

Here we find three assumptions: levels contain things in the world, the levels are ordered by *composition*, and two things are at the same level if they interact in the right sort of way. Wimsatt is also clear to note that he does not assume “that theories are limited to single levels, that levels are always well defined, or that two or more entities can always be unambiguously ordered with respect to level.” (1976, p. 215)

5.2 *Composition, Explanation and Scalar Orderings*

Although he explicitly assumes that levels are ordered by composition, he later complicates his approach by pairing this ordering with *causal-explanatory* relationships; he writes, “higher levels are in some sense *composed of* the entities at lower levels” (1976, p. 243), but then adds that “the arrows [in Figure 17, below] represent significant causal or explanatory connections” (*ibid*, p. 252), and finally pairs the compositional and explanatory together, “If we can’t fairly uniformly say what is composed of what or what explains what, then the things in question are not clearly orderable: they not only fail to be clearly at or not at the same level, but there *are* no clear levels.” (*ibid*, p. 254)¹²⁰

Further complications are introduced in Wimsatt’s discussion on *scale*. He notes that scalar orderings can be largely coincident with compositional orderings (wholes are, after all, typically larger than their parts), where “higher than” increases the scale (roughly) by orders of magnitude (1976, p. 241; c.f., Salthe 2009 p. 88). He offers the following caricature:

¹²⁰ Note that in Wimsatt’s most recent writings on the subject, he appears to drop the explanatory requirement: “By level of organization, I mean here compositional levels-hierarchical divisions of stuff (paradigmatically but not necessarily material stuff) organized by part-whole relations, in which wholes at one level function as parts at the next (and at all higher) levels.” (2007, p. 201)




Figure 1. Levels of organization—some conceivable waveforms.

Figure 16: Wimsatt's Waveforms Depicting Regularity and Predictability Across Levels (1976, p. 240-1)

Figure 16, note, is *not* representative of his view on the layered worldview: paramecia do not compose metazoan, and so it is clearly not a compositional ordering. Rather, he offers the diagram to help clarify the various ways regularity and predictability might map onto a scalar hierarchy, as a way of distinguishing the various alternatives to the reductionist hierarchy. In waveform c, (the middle one), we see that regularity and predictability tends to increase as you descend lower down the hierarchy; he labels this the *pro-reductionist*

waveform. By contrast, the anti-reductionist's waveform (d) has the regularity and predictability pushed towards the topmost level.¹²¹

Despite the rough correspondence to compositional orderings, Wimsatt's appeal to scale should only be taken as a rough *heuristic* with respect to the reductionist hierarchy, for he notes a clear weakness of scalar hierarchies. The same weakness was mentioned in the first chapter: there are many cases in which an increase in level is not coincident with an increase in size; we pointed out that some individual fungi are larger than entire ecosystems and, as Wimsatt points out, "black holes and bacteria can be the size of a dust mote, though the three will have radically different behaviors in similar circumstances." (1976, p. 237) To associate them with the same level seems inappropriate. Scale is no doubt important in both biological and physical sciences – Wimsatt claims that size is a major factor "in determining which physical forces are most central to the explanation of behavior" (p. 206) – but that does not imply that it makes for a good scientific ordering.

Although he argues that scalar orderings are inadequate (he gives an extended criticism of the black-hole-bacteria problem in his 2007, p. 209), Wimsatt nevertheless maintains that scale plays a role when characterizing the contents of each level. He demarcates his levels by virtue of a clustering of causal interactions. (1976, p. 242) In his most recent characterization of the demarcation of levels, he writes:

¹²¹ Noticeably absent from the space of waveforms is a regularly fluctuating, rather than random, pattern of regularities as the hierarchy increases. E.g., somewhat regular at the bottom, very regular near the middle, somewhat regular at the middle, very regular near the top, somewhat regular at the top. It goes without saying that Wimsatt's diagrammatic waveforms do not exhaust the space of logical possibilities regarding the mapping of regularity distributions to the layered worldview. Exactly how these regularities map to levels is an interesting question, but is not one that will be pursued here.

They are constituted by families of entities usually of comparable *size* and dynamical properties, which characteristically interact primarily with one another, and which, taken together, give an apparent rough closure over a range of phenomena and regularities. (2007, p. 204 emphasis mine)¹²²

This characterization appears quite sensible in light of his other views. If, as Wimsatt purports, scale is a major factor in determining what is responsible for a thing's behavior, then this clustering of causal interactability will bear a rough correspondence to similarities in relative size; and it is then no wonder that Wimsatt thinks levels will contain things of roughly the same size, despite that his hierarchy is not a scalar hierarchy.

Even so, Wimsatt acknowledges that the demarcation of levels is imperfect. He notes that "parts of roughly commensurate sizes as the whole system are treated as being at its level." (2007, p. 206) Although this imperfection appears innocuous to Wimsatt, many might think that it renders his account insensible. It might strike some as being *prima facie* inconsistent to maintain that levels are ordered by composition while also maintaining that there are levels at which the composition relation holds intra-level.

But this is an uncharitable criticism. Wimsatt's approach to the scientific hierarchy clusters things together into levels, *and then* orders them according to compositional relationships that are relevant to the explanation of higher level phenomena. The specification of levels comes first, and the compositional ordering is

¹²² This is not a final definition. Yet, oddly enough, it is more informative than the final statement he offers in lieu of a definition: "Level of organization can be thought of as local maxima of regularity and predictability in the phase space of alternative modes of organization of matter. This is the closest I will come to a definition..." (2007, p. 209)

then applied to those specified levels. The view becomes problematic only on the added assumption that ordering-relationships play a large role in both the ordering *and* specification of levels. Wimsatt is not committed to this assumption and so his view retains its coherence.

Indeed, there is a good reason for Wimsatt to want to maintain such a view. Throughout his career, he has been emphasizing what he calls *causal thickets* (1976; 1994; 2007), which he characterizes as a breakdown in clarity in the sense that they “indicate a situation of disorder and boundary ambiguities.” (1994, §IV). The disorder characterizing these causal thickets provide one reason to think the demarcation of levels is at best rough and imperfect.

5.3 *Wimsatt vs. Oppenheim and Putnam*

Wimsatt gathers together these claims – about composition, (reductive) explanation, scale, and complex bundles of interactions – into a singular reductive “hierarchy” (Figure 17 below), which nicely characterizes the difference between his view and that of Oppenheim and Putnam.

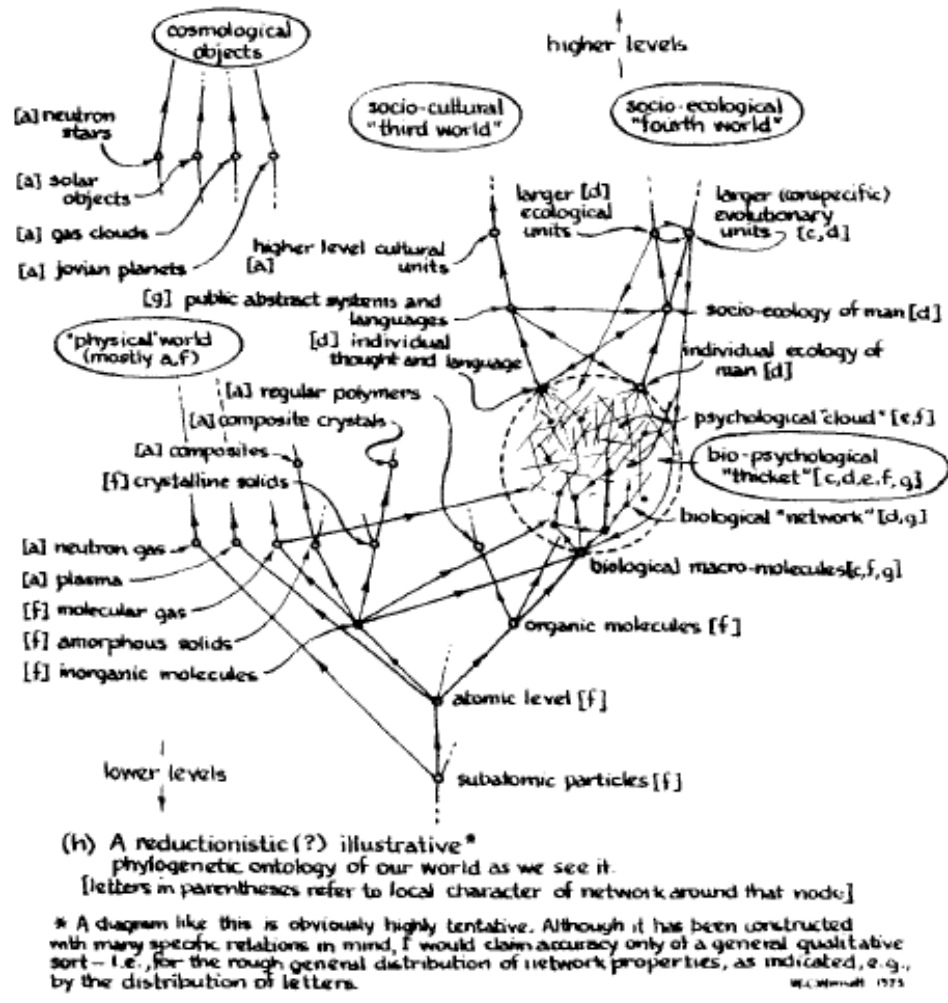


Figure 17: Wimsatt's Reductionist (Quasi)Hierarchy (1976, p. 253)

To hierarchically represent the explanatory relationship that characterize the reductionist's metaphysical epistemology, in Wimsatt's view, requires a much more complex organizational structure than what the classic reductionist ever imagined. As we see in the diagram, Wimsatt even expresses some doubt (viz., his "?"), as to whether his picture really is illustrative of the reductionist's global hierarchy.

He is correct to express doubt. His diagrammatic insight cannot be said to follow from reductionism; it follows only from his own idiomatic starting assumptions about how hierarchies relate to reductionist metaphysics, especially from the way he chooses to specify the contents of levels. For the classic reductionist, levels are specified to correspond to scientific branches. But they are free to specify their collection of levels however they see fit, as long as their specification is justified from the standpoint of empirical science (recall Oppenheim and Putnam's condition (6) mentioned above). Wimsatt selects his levels in a very narrow sort of way: by some function of interactability and relative size. But there are all sorts of ways to specify levels. Wimsatt's preferences for relative size, interaction, regularity, and predictability are perhaps justified from the standpoint of empirical science, but they are not sacrosanct.

The ordering of levels also differs between Wimsatt's diagram and classic reductionists. According to the classic reductionist, the levels are ordered using an asymmetric *microreduction* relationship, whereas for Wimsatt's graphical presentation, the ordering relation appears to be a more general explanatory relationship; both for the reason that the compositional ordering has been clearly been abandoned and that Wimsatt is *not* a microreductionist. In his version of reductionism (if it is rightfully called "reductionism")¹²³,

the point of reduction is not to get an... explanation for "eventually everything" in terms of "essentially nothing," but only to make sure that everything gets explained – at *some* level or other. This in fact allows for

¹²³ Wimsatt (2007, p. 202) acknowledges that his version of reductionism is not what philosophers typically think of as reduction; he thinks it falls in line with what scientists think of as reduction.

the possibility that some things may require explanation at *higher* levels (1976, p. 225).

The explanatory methodology that Wimsatt describes here is clearly *not* one a microreductionist would endorse.

The reason the compositional ordering must be abandoned is indicative of a more serious problem for, at least, his Figure 17 presentation of the reductionist layered worldview: near the top right corner of the diagram there are *higher than* sub-orderings that are either not asymmetric or not transitive; with respect to the edges connecting the “causal thicket” on the right side of the tree to the “socio-ecology of man”, the ordering transforms into a *cyclic* directed graph into and out of the thicket (the direction of the arrows might be hard to see in my reproduction). The cycle, presumably, designates feedback relationships; Wimsatt only tells us that the ordering breaks down.

That breakdown is a problem. Since *higher than* is an asymmetric and transitive relation, as was argued in the previous chapter, and Wimsatt’s ordering is not asymmetric and transitive, he cannot plausibly claim that Figure 17 depicts a hierarchy, or even that is a structure of levels (assuming that levels are the relata of *higher than*).¹²⁴ So what is it, this quasi-hierarchy? It is hard to say. It is a different breed of beast altogether; perhaps it is a hybrid of a hierarchy and some other sort of structural kind.¹²⁵ In any case, Figure 17

¹²⁴ Wimsatt (2007), in a footnote, makes it explicitly clear that his picture is not to be taken as a hierarchy, but rather as a more general *higher than* structure of levels. But we see here that his diagram does not even give him that. The view he expresses in writing, however, is not problematic in the same way. Although his writings do not stray far from the diagram (he acknowledges potentially conflicting orderings), when one ignores the perspectival and thicket-related claims, his view resembles something very much like the new reductionism Gillett describes.

¹²⁵ In fact, Wimsatt himself notes the problem. (1976, p. 254) I do not find his “perspectival” solution to be particularly satisfying. It completely trivializes the levels framework by allowing anything to be higher than

should not be taken to be indicative of any version of reductionism. His waveform characterization, as depicted in Figure 16, is far more illustrative; unfortunately, it is not representative of his actual view on the layered world. The hierarchy depicted by Figure 16 is ordered by size, whereas Wimsatt's reductionist hierarchy is ordered by compositional-explanatory relationships.

Perhaps what is most surprising about Wimsatt is that while he claims to be a kind of reductionist (2007, p. 203)¹²⁶, he nevertheless feels pulled towards, what he calls, holism. He writes,

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.
(*ibid.* emphasis his)

These are hardly the words one would expect a reductionist to write. The characteristic feature of reductionist metaphysics is that there is no *metaphysically* significant hierarchy: everything is found at the level of micro-physics, and the indispensable predicates associated with higher level scientific subjects, in fact, correspond to lower level things and their properties (or arrangements thereof). Higher levels are of little metaphysical importance for the true reductionist since they are superfluous: everything science needs, they say, is found buried "in the basement."

anything else in the same hierarchy, as long as it is permitted by some relevant and adoptable "perspective". From what I can tell, by 'perspective' Wimsatt does not mean *interpretation of 'higher than'*.

¹²⁶ He writes, "As a kind of reductionist, I want to get as much as I can about higher levels from the properties of lower ones." (2007, p. 203)

Although Wimsatt resists the epistemological thesis of microreduction whereby all phenomena are to be explained at the bottommost level, he nevertheless issues remarks which either imply, or strongly suggest, that he adheres to the same DIP endorsed by classic and “new” reductionists:

the same system will be found at a *number* of levels, if it has any reasonable degree of complexity, though it will of course be *a* system at only one level. At lower levels, it will be a multiplicity of systems in complex “ecological” interaction (1976, p. 242).¹²⁷

That is, a system at any given level, is found at all levels below it, not as *a* system, but as a plurality of systems; insofar as ‘it’ in the above quotation designates the system in question, Wimsatt clearly suggests that *it* is at lower levels. Wimsattian reduction, then, upholds the characterization of Reductionism given in the introductory chapter: if the world is layered into a hierarchy (e.g., as we see in ecology), then *everything* to be found in the domain of science falls within the level of fundamental science.

To summarize: Wimsatt, who claims to endorse a sort of reductionism, has a much different picture of the layered worldview than classic reductionists, and perhaps even those Gillett calls the “new” reductionists. His view is very complex. Levels are understood as collections of interacting things of roughly similar sizes. They are ordered by *either* composition relations, or (as depicted in Figure 17) a weaker explanatory relationship that has exceptions to its asymmetry. Classical reductionists view their

¹²⁷ He goes on to write “... and at higher levels, it will appear as a part of a system.” (1976, p. 242) I confess, I do not know what he means by this. Does he mean that the system is found at *all* levels of the hierarchy? If so, then there is a tension here with his earlier claim that levels are an important ontological feature of the world.

hierarchy as a neat pillar of well-chosen nomic kinds; Wimsatt, by contrast, heavily emphasizes a more natural ordering where the levels are not chosen, but are rather the result of similar size, dynamics, and regular interactions. Despite the differences between Wimsatt, the old, and the “new” reductionist they all have (at least) one thing in common: they uphold our initial characterization of Reductionism from the introduction.

6.0 The Nesting Problem for Reductionism -- Old and New

Insofar as Core Hierarchy Theory is our best (simplest and most general) account of hierarchical structures, none of these Reductionisms -- the old, new, or Wimsattian -- provides an adequate characterization of the layered worldview. If we combine the results from our previous chapter with some of the basic commitments of the Reductionist views described here, we expose a very serious tension within Reductionist metaphysics.

Consider the following two claims:

(A) Special sciences reduce to a more fundamental science.

(B) Everything (in the domain of natural science) is in fact found “in the basement” – i.e., in the domain of fundamental science.

Clearly, all of the reductionists described above are committed to both of these claims. (A) is the mantra of the microreductionist, of all sorts;¹²⁸ and (B) is straightforwardly implied by DIP, the downward inclusion principle, or a weakened version of it. For the Reductionist, (B) is trivially true; if it were not, then it would follow that there exists

¹²⁸ Note that Wimsatt is not committed to it. Even so, his reductionist hierarchy is still an upwards directed compositional hierarchy, and our earlier investigation into Wimsatt’s “reductionist” hierarchy indicates an apparent commitment to (B). Since the objection raised below draws out a tension between (B) and compositional hierarchies, that objection applies to Wimsatt’s “reductionism” inasmuch as it does the others.

something that is absent from the fundamental domain, which is contrary to all versions of reductionism we have examined here.

The problem for Reductionism is that (A) and (B) are jointly incompatible under Core Hierarchy Theory. Metaphysical reduction, as we have seen, is an asymmetric and transitive ordering that holds between scientific “levels” by virtue of composition relations holding between the contents of those levels. As an asymmetric and transitive ordering, Core Hierarchy Theory implies that the collection of special science levels (no matter how the reductionist opts to label them) forms a hierarchy. Moreover, it is a hierarchy of a well-known sort: in the Reductionist’s hierarchy, for one level to be higher than another is for it to contain something that has a part contained within the other. That is, the Reductionist’s *higher than* relation is a compositional relationship, in the sense described in the third chapter (See Definition 7 in Appendix B to Chapter 3). The Reductionist’s hierarchy then counts as a compositional hierarchy, as characterized in that chapter: specifically, an upwards directed compositional hierarchy.

Since the Reductionist’s hierarchy is an upwards directed compositional hierarchy, the theorems of our Core Hierarchy Theory require that (B) is false. From the fourth and final theorem found within the appendix, we learn that all (unrestricted)¹²⁹ upwards directed compositional hierarchies are upwards nested. That is, we learn that, *qua* compositional, everything found on any level of the (unrestricted) Reductionist hierarchy is *ipso facto* also on all levels higher than it. If (B) is true, then everything found on any level at all is *ipso facto* at the lowest level. But if everything is found at the

¹²⁹ Recall that by ‘unrestricted’ I mean that no ϕ restrictions have been placed on the @ relation described in Appendix B. ‘@’ is interpreted as pure primitive parthood.

lowest level, as implied by (B), and everything found at the lowest level is found on all levels higher than it, as implied by the notion of upwards nesting, then it follows that everything at any level is, in fact, at *all* levels. But, if there is no difference between the levels, then they are all the same level, which is not consistent with the first theorem of our Core Hierarchy Theory: every hierarchy has at least two unique levels.

Here, then, is the problem: since Reductionism implies both (A) and (B), and (A) implies that the Reductionist's hierarchy is upwards nested, which, when combined with (B), implies a contradiction – it then follows that Reductionism is false. The world *is* hierarchically layered (according to compositional relationships, as the ecologist holds) but it is *not* the case that everything is at the bottommost level of that hierarchy. Since the denial of Reductionism is simply Holism (as it was characterized in the introduction), we therefore have an argument for the *third* thesis of the dissertation: taking Core Hierarchy Theory for granted, Reductionism is false, and Holism is not.

This is no mean result! Let us look back at the argument more carefully. To draw out the inconsistency we inferred that there is only one level from the fact that the levels have identical contents. It might not be clear why this is the case. Perhaps one might deny that hierarchical structures are extensional. Why not, after all, have many levels with exactly the same contents?

There are two answers. First, to suggest that there are distinct levels with identical contents conflicts with the Reductionist's insistence that composition is ontologically innocent. To suggest that there is something more to levels than their contents goes against the very core of Reductionist metaphysics. Second, that there is only one level to

the Reductionist hierarchy straightforwardly follows from Upwards Nesting and the anti-symmetry of parthood. To see this, let us reconstruct (B) using the syntax of our Core Hierarchy Theory. Where the quantifier ranges over the individuals and instancesables in the physical domain, ‘B’ is a name for a particular level (viz., the fundamental level), we have:

$$(B') \forall x(x@B \wedge \neg\exists y(B\text{Th}y)) \quad [\text{Entities and properties are in the basement}]$$

(B’), however, straightforwardly entails that $H@B$, where H is the reductionist’s hierarchy itself (those who worry that H is not within the domain of \forall are hereby referred to section 5.1 below). But, since B is a level of H, it follows from Core Hierarchy Theory (viz., Definition 1 in Appendix B to Chapter 3) that $B@H$. Since the unrestricted ‘@’ designates primitive-parthood, which is antisymmetric, $H@B$ and $B@H$ implies $H=B$. The top level (recall that every compositional hierarchy is the top level of itself) is the bottom level. Given any level of the hierarchy, L, since it is part of H (*qua* being a level of it) and H is part of L (*qua* being part of the lowest level and, given nesting, all higher levels than that), it follows that $H=L$. That is, it follows from the nesting theorem and the anti-symmetry of ‘@’ (i.e., “is at” or parthood) that the reductionist’s hierarchy has only one level.

This takes us to the contradiction. Since our “pairs” theorem (Theorem 1 in Appendix B to Chapter 3) requires that hierarchies have at least two levels, it then follows that there is some level L that is higher than B, since B has been designated to be the lowest level. But since $H=B$, it then follows from Leibniz’ Law that L is higher than H. But, from Theorem 3 of Appendix B to Chapter 3, H, being a compositional hierarchy,

must be a top level of itself: there is no level higher than H. Since L must be both higher than, and not higher than, H, we have an explicit contradiction.¹³⁰

Notice that we do not need to endorse any mereological supplementation principles to derive this contradiction: it is derived straightforwardly from the anti-symmetry of parthood and our Core Hierarchy Theory. We do not require an extensional mereology to acquire this result: the weakest (unsupplemented) mereology combined with the most basic doctrines of reductionism will suffice.

To summarize the argument: the versions of metaphysical reductionism considered above all imply a hierarchical ordering of the explanatory interests of science – specifically, a (upwards directed) compositional hierarchy – as well as the metaphysical view that everything is found at the bottom level of that hierarchy. Having shown in the previous chapter that, unless one places restrictions on the “is at” relation, Core Hierarchy Theory implies that all such compositional hierarchies are upwards nested, are the topmost levels of themselves, and must contain at least two distinct levels (Theorems 4, 3, and 1, respectively), we are able to derive a contradiction: since the Reductionist’s “basement” is in the hierarchy, and the hierarchy itself is in the basement (Core Hierarchy Theory tells us that the unrestricted hierarchy is the topmost level of itself and Reductionism tells us the basement contains all levels), it then follows that the Reductionist hierarchy is its own top and bottom level; but since every hierarchy must have two levels, that implies that there must be a level that is both higher and not higher than the lowest (also highest) level. A contradiction.

¹³⁰ Although the inconsistency appears earlier – having derived that the Reductionist hierarchy has only one level, it thus cannot then have two levels, as the pairs theorem requires – I draw out the explicit contradiction as a matter of methodological preference.

All of the forgoing has been derived using (A) and (B) and Core Hierarchy Theory. Since reductionists are committed to both (A) and (B), they must therefore reject Core Hierarchy Theory, or some relevant aspect of it. In what follows I offer some suggestions on behalf of the troubled Reductionist, but as we will see the matter is not at all straightforward.

6.1 *Option 1: There is no Hierarchy!*

The first reaction to the aforementioned problem is to concede that Reductionism is committed to a compositional ordering, but then respond that the “levels” implied by (A) are metaphysically vacuous. There simply *are* no real hierarchies. If there are no hierarchies, then it should not matter whether or not there is only one level. In fact, that there *really* is only one fundamental domain is exactly the sort of thing the Reductionist might want to claim. A contradiction cannot be generated, the Reductionist might say, because the “pairs” theorem has been misapplied.

Core Hierarchy Theory, however, tells us that there are hierarchies in the Reductionist’s domain of discourse. According to Core Hierarchy Theory, a hierarchy is any collection of levels (which are themselves collections of things), that are ordered by an asymmetric and transitive relation to which we assign the label, ‘higher than’. Reductionists accept that there are chemical, biological, and social phenomena that need explaining; they just say that these phenomena are, like everything else, in the *fundamental* domain by virtue of their decompositions into fundamental parts.

In order to explain higher level phenomenon, the Reductionist – as with the social scientist explaining to the authorities how the Volkswagen found its way onto the roof of

the campus chapel – must appeal to fundamental things working together (or in the case of the social scientist: a bunch of fraternity brothers playing a prank together). When individuals are taken together, however, we have all we need to count them as one. But, I shall now argue, this is all we need in order to apply our hierarchy theory, *pace* the Reductionist.

Core Hierarchy Theory, the simplest and most general account of hierarchies *qua* hierarchies presently available, is described using certain idioms: an individual is something that falls within the range of our existential quantifiers, and a collection is an individual with proper-parts (either other collections or atomic individuals). Since the Reductionist's higher level predicates are latching on to *collections* of fundamental things, those collections are captured by the idioms of Core Hierarchy Theory. And so, insofar as there is a collection within the Reductionist's domain of quantification, and the scientifically relevant predicates applied to that collection stand in a reductive relationship to the fundamental properties of the fundamental parts of that collection, it follows from the asymmetry and transitivity of (micro)reduction, and our Core Hierarchy Theory, that there *is* a Reductionist hierarchy. The Reductionist, then, cannot plausibly reject the Reduction hierarchy while also endorsing Core Hierarchy Theory.

6.2 *Option 2: Reject Core Hierarchy Theory.*

Since the existence of a hierarchy is implied by our Core Hierarchy Theory, one way to avoid all of the problems raised here would be to simply reject Core Hierarchy Theory. That theory, as described in the appendix to Chapter 3, consists in a definition and two axioms.

Rejecting the second axiom, the linkage axiom, is not apt to gain any ground for the Reductionist. And rejecting the first axiom, the existential axiom, has untoward consequences. To deny the existential axiom is to claim that nothing is a level of H, where H is some collection of individuals. But, given Definition 1 – something is a level of H iff it is a part of H and it stands in the higher than relation with something else in H –¹³¹ a denial of the existential axiom entails that either H does not have parts (which is impossible: everything is part of itself) or nothing in H stands *higher than* anything else in H. Since ‘higher than’ here means *reduces to*, it follows that a rejection of the existential axiom only places the reductionist’s view into deeper peril: if nothing in H stands *higher than*, then nothing reduces; and if nothing reduces, then (A), one of the two central reductionist doctrines listed above, is false. Rejecting the axioms of Core Hierarchy Theory, therefore, does nothing to help the Reductionist. The Reductionist must then either reject the definition of hierarchy, or the definition of being a level of a hierarchy.

The definition of hierarchy, again, is this: a collection of levels related and linked by some strict partial order, which we call “higher than.” This definition is something Reductionists (such as Oppenheim and Putnam) take for granted; it appears obvious.

The definition of *is a level of* offers a more promising option: to be a level of a hierarchy is to be a part of it that serves as a relatum of *higher than*. What justifies this definition, I maintain, is its simplicity and generality. The definition must be general in order to cover the vast number of things that a hierarchical analyst might want to treat as

¹³¹ Note that this is the definition of what it is to be a level *of a hierarchy*. To be a level *simpliciter* is even more straightforwardly defined: a level is nothing more than a relatum of *higher than*. Some might dismiss this definition as being cheap, but being cheap is not always bad: if a cheap suit fits the occasion, then wear it; if a cheap definition suits the model, then use it.

hierarchies, or as levels: sets, classes, objects, worlds, networks, or whatever else. Such generality is a virtue of the definition. If there is a problem with the definition, then, it must fall on the appeal to primitive-parthood. This brings us to our next option.

6.3 Option 3: Reinterpret '@' in Core Hierarchy Theory.

Reductionists who work through the proofs in the appendices to Chapter 3, will notice that they rely heavily on a particular feature of the definition of hierarchy used throughout this dissertation: that to be *at*, or *in*, a level, and *at* or *in* a hierarchy is to a *part* of that level or hierarchy. Most of the proofs, one might notice, rely on the anti-symmetry, reflexivity and transitivity of parthood. If it can be rationally argued that parthood is too general, or too specific, to count as the fundamental *association* relation between levels, their “contents”, and hierarchies, then Core Hierarchy Theory can and should be reinterpreted with a more appropriate association relation.

One might think that the primitive-parthood relation is too *specific* because there are other relations that might be taken as an association relation, but are not part-whole relationships. Hierarchies of sets or classes might be more appropriately specified by interpreting the association relation as the inclusion relation, or perhaps the elementhood relation, for example.

I maintain that parthood is *not* too specific. Recall from Chapter 2 that the relationship of parthood was specifically chosen for the analysis of hierarchies for its generality. Following David Lewis’s (1991) proposal that sub-classes are parts of classes, and the proposal given in Chapter 2 that elements (\in) are *special* parts of classes, the primitive-parthood relation, equipped with an extremely general understanding of

‘individual’ and ‘collection’ provide an analysis that is general enough to cover all of the important cases of hierarchies that are relevant to mathematical and empirical science, *including* those specified using the elementhood or inclusion relations. All one must do is place an appropriate restriction on parthood to yield a φ -@ hierarchy in the sense described in the previous chapter: x is \in -@ y iff x is part of y and $x \in y$.

Perhaps, for some reason, one might think that the parthood relation is *too* general. For instance, if a hierarchical analyst wants to analyze a collection of levels that happen, by coincidence, to be order-coextensive with a compositional hierarchy, Theorem 4 requires that one pins every individual at any level to all higher levels higher than the one containing it. For example, if we wanted to model the hierarchy {Piston \rightarrow Engine \rightarrow Car}, our theorems require us to consider each of the lower levels to be at all higher levels. That is, pistons are at the Piston, Engine, and Car levels; and engines are at the Engine and Car levels. But that means we cannot specify the Car level to consist in all and only cars and the Engine level to consist in all and only engines. It might then be objected that since Core Hierarchy Theory cannot give the hierarchical analyst the resources to model these restricted hierarchies because of its dictatorial interpretation of ‘@’, Core Hierarchy Theory’s ‘@’ must then be reinterpreted.

In response, let us note that Core Hierarchy Theory has the resources to model these hierarchies, and so ‘@’ does not *need* reinterpretation. Again, recall our earlier discussion of fiat-levels in the previous chapters. Using the existing interpretation of ‘@’, we can define a more complex φ -@ relation, which we can then use to represent the *restricted* hierarchy one wants. In order to represent the automotive components

hierarchy above we could say the following: where ‘ φ ’ is a variable ranging over kinds, $x\varphi\text{-}@y$ iff $x@y$ and x satisfies φ .

For example, we can say that x is *characteristically-at* y iff $x@y$ and x is of the type characteristic of y . When the analyst goes to model this hierarchy, she first specifies characteristic type for each of the levels, Car, Engine, Piston, and then represents the contents of the levels using “ $\varphi\text{-}@$ ” as defined above, rather than ‘ $@$ ’. By doing so, one yields a hierarchy in which all and only cars are (characteristically-)at the Car-level; all and only engines are (characteristically-)at the Engine-level; and all and only pistons are (characteristically-)at the Piston-level.¹³²

Core Hierarchy Theory’s parthood interpretation of ‘ $@$ ’ is therefore neither too general, nor too specific. The parthood interpretation of ‘ $@$ ’ is exactly as general as one should want: it allows one to define and represent very specific and restricted hierarchies, such as a set-theoretic hierarchy or the simple component hierarchy described here, using ‘ $@$ ’ as a basis, in the most unrestricted sense of ‘ $@$ ’. And as a final consideration, take note that the *parthood* interpretation of ‘is at’ resulted from our negotiation of the disagreement between Roberto Poli and Herbert Simon in Chapter 2. To reject the parthood interpretation of ‘ $@$ ’, one must say what has gone wrong in that negotiation. It is possible that something has indeed gone wrong, but it is now up to the Reductionist to tell us what that something is.

¹³² The hyphenations in parentheses are intended to be read as if they were implicit. Although $\varphi\text{-}@$ allows us to use ‘at’ idiosyncratically, there is no need to explicitly articulate the idiosyncrasy. As one sees if one reads the contents of the parentheses aloud, its success in disambiguating from the more general ‘at’ comes at the expense of making things unnecessarily confusing.

6.4 Option 4: The Reductionist's Hierarchy as a φ -@ Hierarchy

The response from the preceding section might spark an idea in the Reductionist. Reductionist metaphysicians might think to shift attention to a different hierarchy, one specified with a refined notion of @: some sort of " φ -@". If the Reductionist hierarchy is to be represented by a complex φ -@, rather than simple parthood, then upwards nesting would not apply to the restricted Reduction hierarchy. Understood in terms of φ -@ rather than @, the Reductionist's hierarchy would therefore enjoy consistency with Core Hierarchy Theory.

But alas, describing the Reductionist's hierarchy with a specialized " φ -@" relation is not going to make the basic, fundamental @-based hierarchy go away, no matter how clever the definition of " φ -@." The restrictions imposed by the φ -@ definitions strike me as little more than artificial representations of an analyst's conceptual scheme, designed for purposes known only to them. They therefore cannot plausibly be said to relate to the basic hierarchical structure of the world *independent* from our perspective and pre-theoretic biases; namely, to the sort of structure of interest to metaphysics. Hierarchical metaphysics is best conducted in the most general, and most foundational terms if one wants to ensure a reasonable distance from the idiosyncrasies of perspective.

In short: φ -@ implies @. So, the refined φ -@ hierarchy still implies the existence of a general @ hierarchy as described in the previous chapter. Since the @ hierarchy is more basic and more general in that the φ -@ hierarchy is defined on the @ hierarchy, the φ -@ hierarchy does not really get to the bottom of its own metaphysics; it is merely to

turn a blind eye to some of the contents of the levels, putting the spotlight on some of the others. The φ -levels are nothing other than a sort of epistemological favoritism. That is not a sound methodology for Reductionist metaphysics: a large part of the motivation for Reductionism is a devotion to simplicity, the very same devotion that motivates the unrestricted interpretation of @ over the φ -@ interpretation.

6.5 Option 5: 'Hierarchy' is Ambiguous

Without a good reason for rejecting the legitimacy of Core Hierarchy Theory as a theory of hierarchies, the Reductionist must find a way to live with it. One way of doing this might be to plead ignorance. Core Hierarchy Theory was not available to Reductionist's when laying down the groundwork for their metaphysics. Hierarchical structure was taken for granted and never given a rigorous analysis.

Now that a rigorous analysis has been given, the reductionist can nod in agreement: there are indeed hierarchies as described by Core Hierarchy Theory. But that, they might say, is irrelevant. This entire time Reductionists were objecting to the metaphysical robustness of a *restricted* φ -@ sort of hierarchy discussed above, rather than the unrestricted @-based (i.e., primitive-parthood based) hierarchy. They had no reason to think that there was a more fundamental hierarchy lurking in the background. But, now granting the meanings of 'individual' and 'collection' used to express the Core Hierarchy Theory, its level of generality might appear quite sensible to the reductionist. They might now agree that there is indeed this more fundamental hierarchy, having finally made clear what we mean by 'hierarchy'. But this is no slight against Reductionism; its DIP is only intended to apply to the *restricted* hierarchy, and not to the

fundamental hierarchy. Having picked its target-hierarchy carefully, there is nothing at all inconsistent with the Reductionist position.

Such a response should spark concerns that the Reductionist view collapses into Holism. Since the Holist's hierarchy is straightforwardly captured within Core Hierarchy Theory, and the Reductionist admits the existence of *that* hierarchy (in addition to, or rather, underwriting, their own special hierarchy), it follows that Reductionism with respect to one hierarchy is a sort of Holism with respect to another. For the Reductionist to base their defense on an ambiguity in the term 'hierarchy', the whole debate becomes a farce: the disagreement was no disagreement at all. Reductionists and Holists were simply talking about different hierarchies: the Holist is talking about the most general sort of hierarchy, while the Reductionist's view involves reference to a special and restricted sort of hierarchy. Since the disagreement certainly appears genuine, the ambiguity option is rather unsatisfying.

But perhaps the ambiguity option is unnecessary. Perhaps the Reductionist thinks that this concern about genuine disagreement warrants a reconsideration of the φ -@ option suggested in the previous section. That there is a genuine metaphysical disagreement between Holism and Reductionism might lead to an argument *against* the idea that the best way to investigate hierarchical metaphysics is to use the unrestricted @ interpretation in lieu of a φ -@ interpretation. The argument is this: if we investigate the debate using the unrestricted approach, then the debate disappears only because Reductionism has been unfairly or illegitimately suppressed on an *a priori* basis; there is a genuine debate that should not be unfairly or illegitimately suppressed, and so, we should not use the unrestricted interpretation of @.

Have we been unfair to Reductionists in adopting the unrestricted interpretation of ‘@’ to represent their (metaphysical) disagreement with Holists? Is that adoption illegitimate? While it is true that the dispute between Reductionism and Holism is not preserved by within unrestricted Core Hierarchy Theory, it takes a lot of hard work to see that. It is by no means obvious. If there is a dispute between two views V1 and V2, and an independently motivated (e.g., via simplicity and generality) metaphysical framework resolves the dispute in a way that favors V1 over V2, then to defend V2, one must challenge the motivation for the framework. That the framework favors V1 over V2 is no reason to reject that framework.

Even if one *insists* that there is a genuine metaphysical dispute between reductionists and holists that the unrestricted application of Core Hierarchy Theory fails to capture or illegitimately rules out, the Reductionist must argue for this insistence; they must offer an account of the restricted hierarchy that best frames the dispute, and then tell us why it is the best. Moreover, whatever restricted hierarchy the Reductionist offers must be one to which the Holist agrees, otherwise the dispute will not have been appropriately restored.

It is not obvious that there is a restriction that satisfies these constraints. The most obvious restriction involves a set-theoretic approach: to be “at” a level is to be an \in -part of it; that is, the idea that levels are then elements rather than primitive-parts. The Holist, however, will protest. The set-based approach is impure compared to the unrestricted approach, in the sense that set-based hierarchies do not consist entirely of their levels and the contents of those levels (and collections thereof): set-theoretic hierarchies also consist in the empty set, which cannot be mistaken for nothing, for the empty set is surely

something. The unrestricted mereological approach is then simpler (more parsimonious) by a mere hair. To bring the set-based approach into parsimonious parity with the unrestricted mereological approach, the set-based approach must construct their account of the layered worldview using a set theory that forgoes the empty set. I know of no such set theory. But, then again, such ignorance is hardly compelling. Nevertheless, the onus is on the Reductionist to defend a restricted approach.

As a final note, it should be emphasized that the nesting results obtained from the unrestricted Core Hierarchy Theory are by no means obvious, and are perhaps even astonishing. Non-compositional hierarchies, after all, are unaffected by the theorems; they behave exactly as one would expect under the unrestricted Core Hierarchy Theory, with no surprises. Only when we turn our attention to *compositional* hierarchies do we obtain surprising results. Before witnessing the theorems derived from the unrestricted interpretation of '@', I very much doubt that the Reductionist would object to that interpretation. It is not until the nesting results are obtained that the legitimacy of '@' is questioned. I thus see no reason to think that we have been *unfair* to Reductionism in adopting an unrestricted framework.

6.6 *Option 6: Challenging the Assumptions*

Are there any other options available to the Reductionist? Perhaps there are. One might, for example, deny the assumption of mereological monism stated in the introductory chapter, and then use an entirely different, non-classical, notion of parthood as the interpretation of '@'. One might *define* a hierarchy to be a set (instead of a general mereological collection) and then interpret '@' as the transitive closure of

elementhood.^{133, 134} One might take a hierarchy to be a plurality of levels, or one might take it to be a kind of relation. We have already considered these options, but there might be any number of alternatives to Core Hierarchy Theory that have not been considered here. Logical space has not been completely exhausted; the second chapter contains only a sampling of possible accounts of hierarchies.

Even so, the development of Core Hierarchy Theory has proceeded by argument. To adopt an alternative interpretation of '@', one must confront those arguments. To reject mereological monism, for example, one must reject the arguments given in its favor (see the introductory chapter). To reject the mereological approach in favor of a strictly set-theoretic approach, one must argue against the mereological approach, against the idea that the existence of a set hierarchy implies the existence of a more general mereological hierarchy (see Chapter 2), or at least against the assumption that 'hierarchy' is univocal. To challenge univocality one must also argue that the set-theoretic approach is not simply focusing on a sub-type of hierarchy that is already captured by a simple restriction on the mereological approach; otherwise, the set-based definition will suffer from a lack of generality. Not all hierarchies *need* be thought of as sets, unless a good reason can be given to the contrary.

As a final note, recall that two virtues of the mereological approach were listed in the second chapter: simplicity (specifically, parsimony) and generality. To offer an alternative interpretation of 'hierarchy' or '@' as a way to save Reductionism, those

¹³³ Recall, from Chapter 2, the argument that @ is a transitive relation: '@' means *in*; since knowing that x is in a level and L is in hierarchy H is enough to infer that x is in H, it follows that *in* is transitive, and thus so is @.

¹³⁴ One would also have to argue that the transitive closure of elementhood is not simply the parthood relation.

virtues must be retained, downplayed, or other virtues must be brought to bear. The Reductionist, after all, certainly does not want to abandon simplicity as a virtue; simplicity is among the main virtues of the Reductionist picture to begin with. We must then ask: aside from saving Reductionism, what benefit is gained from rejecting Core Hierarchy Theory as it has been developed here? Does the unrestricted Core Hierarchy Theory fail to capture some relevant feature of the Reductionist-Holist debate? If so, what is that feature and why are we unable to capture it?

There might be satisfying answers to these questions, but if so, they are not available to us at this stage of the dialectic. We have no choice but to proceed under the assumption that Core Hierarchy Theory is indeed the best available account of hierarchies.

7.0 Conclusion

Insofar as Core Hierarchy Theory is our best account of hierarchies *qua* hierarchies, if it can be said that the world is layered, it is *not* layered by virtue of the Reductionist's hierarchy. Although not much has been said about the epistemology of Reductionism, we have a complete critical analysis of its metaphysics using Core Hierarchy Theory as our central framework. The metaphysics of Reductionism is inconsistent with our simplest and most general understanding of hierarchies *qua* hierarchies: its two most basic doctrines are incompatible with Core Hierarchy Theory. To uphold their view, the Reductionist must develop reasons to reject or modify Core Hierarchy Theory.

In the next chapter, we will investigate the layered worldview as described by anti-reductionist philosophers. We will also pay special attention to the work of Jaegwon

Kim. Although Kim is a Reductionist, his work will help us to understanding the connection between the *realization* relation and the layered worldview.

Chapter 5

The Layered Worldview and the Metaphysics of Multi-level Explanation

1.0	INTRODUCTION	239
2.0	THE BASIC METAPHYSICAL FRAMEWORK	242
2.1	<i>THE HUMEAN VS. MECHANISTIC APPROACHES</i>	242
2.2	<i>JACKSON AND PETIT ON EFFICACY VS. RELEVANCE</i>	243
3.0	GLOBAL ORDERINGS: SUPERVENIENCE AND NESTING	245
3.1	<i>SUPERVENIENCE AND THE LAYERED WORLD</i>	245
3.2	<i>ARGUING AGAINST SUPERVENIENCE ORDERINGS</i>	247
3.3	<i>FORGING A GLOBAL HIERARCHY</i>	249
3.4	<i>ARGUING AGAINST MONOLITHIC GLOBALISM</i>	251
3.5	<i>AN UPWARDS-NESTED GLOBAL HIERARCHY</i>	252
4.0	LOCAL ORDERINGS: THE DUAL-HIERARCHY MODEL	255
4.1	<i>CAUSAL EXCLUSION AND FUNCTIONAL REDUCTION</i>	255
4.2	<i>ORDERS VS. LEVELS: THE DUAL HIERARCHY VIEW</i>	258
4.3	<i>THE DUAL-HIERARCHY VIEW, REDUCTION, AND CORE HIERARCHY THEORY</i>	260
5.0	THE MECHANISTIC HIERARCHY	261
5.1	<i>WHAT IS A MECHANISM?</i>	262
5.2	<i>GLENNAN ON THE DUAL-HIERARCHY VIEW</i>	263
5.3	<i>PRODUCTIVE ACTIVITIES VS. CAUSAL RELEVANCE</i>	265
5.4	<i>CRAVER AND THE LAYERED WORLDVIEW: LEVELS OF MECHANISMS</i>	266
5.4.1	<i>Craver, Mechanisms and Globalism</i>	268
5.4.2	<i>Craver on Levels-of-Mechanisms</i>	271
5.5	<i>MECHANISMS, CAUSATION AND MULTI-LEVEL EXPLANATION</i>	274
5.5.1	<i>Mechanistic Explanation and Programmatic Explanation</i>	276
5.5.2	<i>Mechanisms and Metaphysical-MLE</i>	277
5.5.3	<i>Nested Hierarchies and Metaphysical-MLE</i>	279
5.6	<i>THE CAUSAL EXCLUSION PROBLEM REVISITED</i>	283
5.6.1	<i>Mechanistic Hierarchies are Dimensioned Realization Hierarchies</i>	284
5.6.2	<i>Causally Excluded Activities and Mechanisms</i>	286
6.0	CONCLUSION	288

1.0 Introduction

Two thirds of our central thesis, Multi-Grade Holism, have now been defended. We have argued for Realism and for Holism. There is only one issue left to address: the issue of

so-called “multi-level” explanations in ecology (and perhaps other life sciences). At the outset of the dissertation, we inquired into the justification for the multi-level approach.

We listed two possibilities:

[Metaphysical-MLE] The hierarchical structure of the world justifies our describing, explaining, and analyzing things using hierarchical terms.

[Epistemic-MLE] The *usefulness* of hierarchical structure justifies our describing, explaining, and analyzing things using hierarchical terms.

We also noted some connections between these two justificatory approaches and the two theses defended thus far. For one, if Reductionism is true then neither Metaphysical-MLE nor Epistemic-MLE is true; if everything is at the bottommost level of the “layered” world, and all can be described, explained, and analyzed using resources contained within the bottommost level, then there is no need for multi-level explanation at all.

Reductionism is incompatible with the very notion of a multi-level explanation.

Nevertheless, we argued against Reductionism in the previous chapter in favor of Holism. Explanations of ecological phenomena that span over (by virtue of quantifying over individuals and instanceables at) multiple levels seem quite fitting under a Holist’s view. But now we must ask: what is it that makes these sorts of explanations (e.g., those that appeal to the nitrogen cycle) fitting: is it because the world is structured hierarchically, because it is *useful* to exploit hierarchical structures, or both?

In the introductory chapter, Realism was paired with Metaphysical-MLE. Having argued for Realism in the third chapter, one would expect that the metaphysical justification for multi-level explanation follows straightforwardly. It does not. In this chapter, I raise a problem for the pairing of Realism with Metaphysical-MLE with respect to a certain class of hierarchy: the upwards nested hierarchy. Inspired the Reductionist argument above, I offer some considerations which seem to *suggest* that Metaphysical-MLE is false, or at least dubious, given the sort of Realist-Holism defended so far. Near the end, I offer some suggestions on how the Realist-Holist might go about defending Metaphysical-MLE.

As it was characterized earlier the Holist thesis relates specifically to the philosophy of ecology. Unfortunately, most philosophical accounts of explanation in multi-level contexts fall outside of philosophical and theoretical ecology. So, to investigate Metaphysical-MLE within ecology, we must broaden our approach. Here I examine views outside of ecology that share the relevant features of Holism: views according to which the world is (non-reductively) layered in a way that is relevant to scientific understanding.

I begin with a close examination of Jaegwon Kim. Although Kim is a self-ascribed reductionist, his view nevertheless provides a useful starting point for our investigation. His approach to the layered world is unlike anything we have seen so far. After arguing against Kim's special approach to the reductionist layered worldview, I present the mechanistic views of Stuart Glennan and Carl Craver. Although I agree that *levels-of-mechanisms* can be used to characterize the layered worldview, by re-introducing Carl Gillett's work on realization, we will see that it provides a local layering

that, *pace* Craver, subsumes the mechanistic hierarchy. Before getting to the details of Kim's approach, it is prudent to distinguish the basic metaphysical commitments of Kim from those of the mechanists.

2.0 The Basic Metaphysical Framework

Taking Realism for granted, the world is layered in at least one of the two senses described above: either globalism or localism, or both, is true. We do not know which. Of the two, the local approach is far less contentious. Let us focus there, first by describing and contrasting the basic metaphysical commitments of the main proponents of a (local) layered worldview.

2.1 The Humean vs. Mechanistic Approaches

As Glennan (2010) notes, the metaphysics of "non-reductive" explanation comes in two general varieties. One variety is the standard Humean approach, which includes entities, the properties they exhibit, relations, and laws. The rival variety, the "mechanistic" approach, includes entities, the activities they engage in, the properties they exhibit and the relations they stand in. Proponents of either approach also tend to mention events in their discussions, most often without endorsing any specific account of events.

This notion of *activity* is the main point of disagreement between standard non-reductivists and mechanistic philosophers. Mechanists reify activity as a special ontic type (Glennan 2010, Machamer, Darden, and Craver 2000, Craver 2007), whereas supporters of the standard view typically understand *activity* simply as a series of property changes over some span of time.

The difference is crucial to their respective approaches to causal explanation. To causally explain a phenomenon, one gives some sort of account of how it came to be produced. For the mechanist, reified activities are what account for this productivity since activities, on the mechanist's view, are the ontic locus of productivity: they are understood simply as those bits of reality that are productive. Since activities are taken to be ontic primitives by mechanists, many think that no further account of productivity is then to be expected of them. Scientifically interesting phenomena occur because they are *produced* by the entities and activities organized within a mechanism. In short, the notion of causation is built into the mechanistic framework itself as an ontic primitive: activity.

2.2 Jackson and Petit on Efficacy vs. Relevance

For the standard non-reductivist, on the other hand, causation is far more mysterious. The relata of the causal relation are standardly taken to be events; and so productivity seems to be found either in the antecedent causal event itself, or some property that it instantiates.¹³⁵

Jackson and Petit (1990) attempt to capture the distinction between those properties (or events) that are productive and those that are not by drawing a conceptual distinction between *efficacy* and *relevance*. They justify this distinction using three assumptions, one of which is an account of *inefficacy*: property F is not efficacious in the production of an effect e if (i) there is a distinct property G such that F is efficacious in e's production only if G is; (ii) the F-instance and the G-instance are not sequential causal

¹³⁵ To say that the productivity is found within the causal relation itself suggests that the causal relation is an antecedent cause; but it does not seem right to say that the phenomenon of interest occurs because of some cause C as well as the causal relation R into which it enters. Causation does not cause the door to unlock when I turn the key. Rather, the turning of the key caused the door to unlock.

factors (they are not simply nodes in a single causal chain); and (iii) the F-instance and G-instance do not combine together to produce e, that is, they are not “coordinate causal factors.” (1990, p. 108) Given this, and their other two assumptions, they go on to argue that higher level properties, such as special science properties, are causally relevant but are not efficacious. Efficacy, and hence productivity, is found only at the level of fundamental physics according to Jackson and Petit.¹³⁶

I do not think we should endorse their causal metaphysics immediately. Their account of *inefficacy* (p. 108) seems to imply that in cases of genuine causal overdetermination (e.g., Bill and Suzy both throw rocks which simultaneously strike, and shatter, a window, yet either rock alone would accomplish the shattering) the overdetermining causal properties turn out to be inefficacious, since the overdeterminers meet the three aforementioned conditions for inefficacy: (i) qua simultaneous, if one rock-throw is efficacious, then so is the other; (ii) qua simultaneous the overdeterminers are not simply sequential instantiations in a single causal chain; (iii) qua individually sufficient for the effect, the overdeterminers are not coordinate causal factors (they do not work together to produce the effect). The problem is that genuine overdeterminers should be considered efficacious: what makes overdetermination cases puzzling for counterfactual accounts of causation is that in such cases there is more efficacy than is needed, but on Jackson and Petit’s account, there is no efficacy at all. The consequence that there is no efficacy at all is just as much a problem for Jackson and Petit’s argument as it is for naïve counterfactual accounts of causation.

¹³⁶ This metaphysical picture informs their special view of multi-level explanation, which they call *programmatically explanation*. We will contrast their view of multi-level explanation with that of Craver’s in a later section.

This is, of course, a minor complaint. It does not imply that there is no metaphysical distinction between efficacy and relevance; but rather, only that one should be careful to endorse Jackson and Petit's account of it. Since their key premise regarding inefficacy is faulty, we must leave their claim that only the lower level is efficacious as a metaphysical option for which we have insufficient evidence.

To summarize: the basic metaphysical picture, in general, includes entities, some locus of productivity, properties, relations, and often events. In the next section, I examine how these different approaches might facilitate a global hierarchical ordering.

3.0 Global Orderings: Supervenience and Nesting

Despite favoring a sort of reductionism, Kim's writings on the layered worldview are not entirely critical. Kim lays out a metaphysical view that utilizes not one, but *two* hierarchies, and the relations they bear to one another.

3.1 Supervenience and the Layered World

Although Kim (2002) expresses explicit interest in the *metaphysics* of the layered worldview, he is sure to note that it might in fact be incoherent to say that the world itself has a layered structure, and that levels might only apply to concepts, descriptions, or language. (2002, p. 4) He is also sure to note the importance of being clear about what is involved in the levels picture. He writes:

The idea of "higher" and "lower" is essential to the levels picture: the levels are ordered from lower to higher, and as noted it is often thought that there is a bottom level, a level than which there is no lower. Unless

one could meaningfully speak of “higher” and “lower”, there would be little point in talking about levels... those who use the levels language today are, we may assume, concerned exclusively with the physical domain, and the imagery evoked by levels talk is a picture, somewhat fuzzy and unarticulated, of the physical world neatly stratified into a structure of discrete levels, with a bottom level of basic particles... and the rest as forming a vertically ordered system of levels each resting on the one below and all ultimately resting on the base level of microparticles. (*ibid.*)

The view he describes here is not a view he intends to defend. What we see above is Kim’s description of the *globalist* approach to the layered worldview, which he later rejects in favor of a sort of localism.

This characterization of globalism is not sacrosanct. We have already argued against the idea that levels are necessarily discrete in the first and third chapters when describing the notion of a *nested* hierarchy. We might also question whether or not there must be a lowest level, as Kim suggests. And, following Wimsatt, we might question whether or not the layering must be non-branching, as Kim suggests with the expression “vertically ordered system.” Aside from these quibbles, Kim’s approach to the layered worldview roughly matches that which we have been discussing so far. He attempts to chip away at the globalist worldview by showing the discreteness and non-branching constraints to be objectionable. He is not chipping away at the more general understanding of globalism we had specified earlier.

When it comes to the layered worldview, Kim favors a supervenience approach (2002 p. 10).¹³⁷ Specifically, he appeals to the notion of *strong* supervenience between sets of properties U and L:

U strongly supervenes on L iff for any property u in U, if something x instantiates u, there is a property, l, in L such that x instantiates l, and, necessarily, if anything instantiates l, it instantiates u. (2002 p. 9)

To preserve the asymmetry of the *higher than* relation, Kim adds constraints to the strong supervenience ordering with the following definition: H asymmetrically supervenes on L just in case H strongly supervenes on L and L does not supervene on H. This gives us an analysis of *higher than* for the layered worldview.

(SV) H is higher than L iff H asymmetrically supervenes on L.

3.2 *Arguing Against Supervenience Orderings*

There are problems with the supervenience approach. Consider Angela Potochnik's (2010)¹³⁸ counterexample: the phenotypic properties of organisms are typically considered to be higher than genotypic properties, but phenotypic properties do not supervene on genotypic properties. An organism's camouflage, she argues, varies independently of any changes to the organism's genotype. There are thus cases in which x is taken to be higher than y, but x does not supervene on y. Kim's (SV) analysis seems to fail.

¹³⁷ It is unclear whether or not Kim is talking about C. Lloyd Morgan's hierarchy in particular or the layered world in general, when he describes supervenience orderings as the best option for order the levels of the layered worldview. I will presume the latter since it is congenial to his larger view.

¹³⁸ Also see Potochnik and McGill (forthcoming).

Potchnik's objection relies on an uncharitable reading of Kim. In earlier chapters, we noted that there is a difference between the *higher than* relation and the *is at a higher level than* relation; this distinction provides a way for Kim to avoid Potchnik's worry. If we allow that *higher than* relates the levels themselves whereas *is at a higher level than* relates the contents of those levels, we can take (SV) at face value, leaving *is at a higher level than* unanalyzed. In issuing her objection, Potchnik's must interpret (SV) to be an analysis of *is at a higher level than*: Kim is not committed to the idea that phenotypic and genotypic properties are themselves levels; he is only committed to the idea that these properties are *at* levels. If (SV) is not an analysis of *is at a higher level than*, then he is free to maintain that the level that contains phenotypic properties (i.e., the organism level) does supervene on the level that contains genotypic properties (i.e., the molecular level): the distribution of molecules in our world, by hypothesis, necessitates the actual distribution of phenotypic and other organismal properties. Higher *levels*, then, might indeed supervene on lower levels; Kim's account appears salvageable despite Potchnik's counterexample.

Having a response to Potchnik is not enough. There is a more serious problem for Kim's supervenience ordering. It involves the nature of the supervenience relation itself. If Kim's definitions are correct, then 'Q strongly supervenes on P' is about instantiation, elementhood, and a modal relation, viz., $\Box(P \rightarrow Q)$. Supervenience claims are not made true by virtue of the way *our* world is; they are made true by virtue of what our world *and other-worlds* are like.

That these claims are not made true by virtue of what our own world is like strongly suggests that supervenience relations do not exist in the same way as other,

similar, ordering relations: *reduces to* and *is realized by*, to use familiar examples. Claims about those relations, it seems, are made true by how things are in *our* world. Since it is difficult to see how an other-worldly (perhaps “many-worldly” is better) relation might give structure to *our* world, it is difficult to see supervenience as the relation to bring hierarchical order to our world. The existence of hierarchies, recall, is justified on the grounds that its defining-parts exist: that the collection of things being ordering hierarchically is a reified collection, and that the relation serving as *higher than* is also reified. Globalist hierarchies are no different. To put the point succinctly: since supervenience is not a reified (i.e., “real”) relation within the scientific domain, it does not qualify as a *higher than* relation for the layered world.

3.3 *Forging a Global Hierarchy*

So far, we have restricted the discussion to Kim’s remarks on a global hierarchy of properties. But there is more to the layered worldview than simply properties. There are individuals as well. To set up his argument against the globalist layered worldview, Kim chooses Oppenheim and Putnam’s hierarchy as his target (2002, p. 12). Starting with what he calls the universal domain of physics, U , Kim carves out a comprehensive, and familiar, account of a global layering:

First consider atoms: atoms and aggregates of atoms form a subdomain of U ; call it U_A ... Molecules and aggregates of molecules form a subdomain of U_A , and hence of U . Call it U_M . What are left out of U_M are elementary particles and atoms that are not part of molecules. In a similar fashion we can take cells and their aggregates as a subdomain, U_C , of U_M . We then

have the domain of multicellular living organisms, U_L , as a subdomain of U_C , and the domain of social groups U_S . Notice that U , U_A , U_M , U_C , U_L , and U_S form a nested series in which each domain that follows U is properly included in its predecessor. (Kim 2002, p. 17)

Following what Kim tells us here, we are presented with the following picture:




Figure 18: Kim's Nested Global Hierarchy

Figure 18 bears both a striking resemblance to the Oppenheim and Putnam downward-nested reductionist hierarchy discussed earlier and a stark contrast to the Allen and Starr upward-nested hierarchy from the first chapter. The only difference between Oppenheim and Putnam's hierarchy, and that of Kim, is that the former hierarchy is ordered by a *microreduction* relation, whereas Kim offers no explicit interpretation of 'higher than'. Implicitly, he seems to maintain that the asymmetric supervenience ordering is correct, at least for properties; and then adds a (de)compositional mereological ordering for *individuals* in the hierarchy. Having developed what he takes to be the strongest presentation of the globalist picture, Kim then aims to refute it.

3.4 *Arguing Against Monolithic Globalism*

Kim raises a number of problems for the approach to globalism set out in the previous section. His free-molecule problem, which we discussed in the first chapter, is a straightforward counterexample to the claim that the level of living organisms is properly included in the level of cells. Some of an organism's molecular parts might not be included in the level of cells (i.e., there might not be a cell for which the free molecule is a part), and so the level of organisms cannot be completely included within the cellular level.

While it might also be true that the organism level is properly included in the molecular level, Kim thinks it is unacceptable to simply crop out the cellular level since "cells form a significant nomic kind, and the biological functions and behaviors of organisms seem perspicaciously explainable in terms of the properties and behaviors of cells." (2002, p. 17) In short, Kim thinks cells must be included as a level in the layered worldview because of their significance for multi-level explanations of biological functions and behaviors.

Kim raises other concerns as well. He thinks that the possibility of artificial minds and the mere existence of computing machines "throws into doubt the idea of a single hierarchy of connected levels", (2002, p. 16) suggesting that there is little sense in asking whether such machines are higher or lower than sea slugs or birch trees. (*ibid.*)

To continue to endorse globalism, Kim tells us, there is little recourse but to say that such computers would reside at the molecular level, since they are not composed of biological cells, and are not living organisms. Otherwise, one would have to endorse what

we have been calling a branching-hierarchy, (2002, p. 17) in which the computational and cellular levels reside on separate branches of the hierarchy, and are not comparable in terms of “higher than”. Branching, however, degrades the neat and tidy globalist layered world as Kim conceives it, into a comparatively messy tree-like branching structure with overlapping levels. (*ibid.*) Although Kim has taken pains to provide the most charitable interpretation of monolithic globalism, his concerns ultimately lead him to reject it.

As we remarked earlier, however, Kim’s version of globalism is loaded with unnecessary details. A purer form of globalism claims only that the entire world has a hierarchical structure; not that it has a discrete, non-overlapping, and non-branching structure. Such qualities are not essential to hierarchies and so they are not essential to a globalist hierarchical worldview. Kim has not shown that a globalist approach to the layered worldview fails; he has only shown that the version he takes for granted fails. There are other available options aside from what has been offered by Oppenheim and Putnam.

3.5 *An Upwards-Nested Global Hierarchy*

Let us pause momentarily to take notice of how an upwards nested version of Kim’s global hierarchy can avoid the free-molecule and branching problems altogether. On such a view, social groups are not properly included in the level of organisms, but rather, vice versa. Starting at the top, the level of social groups (rather than physics) becomes the “universal” domain, organisms are properly included in that level (qua proper-parts of social groups), and the nesting of cells into the organism level, molecules into cells, and then atoms follows suit. This upwards-nested picture frees us from having to find a

reason for placing things in the level directly below; and at the same time, provides a reason to place things in the level directly above: the upwards-nestedness of (unrestricted) compositional hierarchies is a theorem of Core Hierarchy Theory, and so we do not have to account for *why* free-molecules are found on the cell level, they are there because we are using the simplest and most general interpretation of hierarchical structures Core Hierarchy Theory can offer.

One might be bothered by the fact that not everything belongs to a social group, leading one to think that it is inappropriate to use the level of social groups as the universal domain. I am sympathetic to such a concern: galaxies and planets are poorly suited for inclusion in the level of social groups. Nevertheless, this is a mere toy example (i.e., a local hierarchy) used to make a small point about how upwards nesting can resolve Kim's concern. In a truly *globalist* picture, the upwards nested hierarchy would have the universe itself as the topmost level.

Still, one might be bothered by the fact that not every organism, cell, molecule, and atom belongs to a social group. This might lead some to question the appropriateness of the nesting approach in general. To *this* concern, I am not as sympathetic. Antlers, to reuse an earlier example, are to be considered at the organism level, but they do not have to actually be attached to any deer for them to be appropriately positioned at that level.

Kim and his reductionists cohorts will complain that the upwards nested picture is at odds with the established view that physics is the only science with "full coverage" (Kim 2002, p. 16); i.e., the only one that seeks to provide "a comprehensive description and explanation of all phenomena of this world." (*ibid.*) Everyone agrees that this

ambition should be respected in a hierarchical representation of the world; but, they would argue, inverting Kim's picture to form an upwards nested hierarchy, in fact, inappropriately shifts these ambitions from physics to sociology, economics, or social science in general. That, we should all agree, seems incorrect.

Such complaints are of no serious concern. They are based upon a distortion of Oppenheim and Putnam's vision for scientific unity: the idea that the best or only way to represent the role of physics is to associate it with the lowest, i.e., the "basement", level of some grand hierarchy, and then apply a reductionist principle – such as the downward inclusion principle (DIP) from the previous chapter – in order to ensure that everything is included within the bottommost level. But there are other straightforward ways to represent the "full coverage" of physics aside from applying the DIP: physics, we can say, is the only science that seeks to explain the behaviors of things found at *any* level. Physics, therefore, has "full coverage" on the upwards nested view.

Some might dismiss this as "full coverage" in name alone; it was applied by stipulation. I disagree. Classical physics deals with the motion of macro-bodies, and astro-physics deals with the motions of galaxies and galaxy clusters. Moreover, on an upwards-nesting view, everything at any given level is also at all higher levels; atoms and other micro-particles studied by physicists, according to the nested view, are thus found at *every* level, from the bottom to the top. Physics should be labeled with "full coverage" since everything within the topmost level falls within its scope, and since everything falls within the topmost level on an upwards nested view, it straightforwardly follows that everything falls within the scope of physics. If micro-physics is associated with the bottommost level, physics, in general, should simply be associated with the universe

(topmost) level. The idea that everything falls within the scope of physics, but not everything falls within the scope of *micro*-physics is a perfectly sensible way to understand the structure of science and its connection to the structure of physical reality. Earlier we called this sort of view Holism.

To summarize: Kim's arguments against globalism – viz., the idea that the entire universe is structured hierarchically – are unconvincing. He has only raised problems for *certain* global hierarchies: such as those which expect to carve out a special place for everything of interest to scientists, and those for which the levels are taken to be discreet or downwards-nested. It was also argued that a supervenience ordering does not provide a satisfactory interpretation of 'higher than' for a layered worldview.

4.0 Local Orderings: The Dual-Hierarchy Model

Let us now examine Kim's localist position. Kim uses his locally layered worldview to frame and develop a novel sort of reductionism. This reductionism is a bit different from the microreductionist view we developed earlier; but, as we will see, it is no improvement.

4.1 Causal Exclusion and Functional Reduction

Kim (1993; 2002) endorses the standard metaphysics described in the previous section.¹³⁹ He takes the relations of reduction, supervenience, and causation to hold primarily between properties. Supervenience is characterized as a relation between *families* of properties; individual events stand in causal relations insofar as they are subsumed under

¹³⁹ Unlike most Humeans, his metaphysics involves some sort of notion of a *causal power*.

some law, where laws are understood as an objective connection between properties.

Kim's functional-model of reductionism is, in his view, the most plausible way to deal with the metaphysical problems that arise for the standard approach.

The problem Kim is concerned with, the *causal exclusion problem*, is well-known and has a substantial body of literature devoted to it.¹⁴⁰ The problem, roughly, is that there appears to be a tension between the common idea that higher level properties (having a desire, for example) are efficacious in bringing about certain higher level behaviors, (going for a run, for example), and the idea that those same behaviors are also brought about by lower level, that is, fundamental physical, properties. Suppose H and H* are higher level properties, and that H causes H* to be instantiated. Let L and L* be the respective lower level supervenience base for H and H*. Since H* supervenes on L*, the only way for H* to come about is for L* to come about. But if we then suppose that L causes L*, then we have a complete story for H*'s production without invoking H whatsoever; that is, there is no causal role for H to perform since L is causally sufficient on its own. Without a role to play, it appears that H is inefficacious and epiphenomenal.

Kim's point is that if we were to say that $H = L$, then H would play the appropriate causal role: the very same role as L. Epiphenomenalism can then be avoided via the reduction of H to L. To establish this sort of reduction, Kim introduces what might be called the *functional model* of reduction. His idea (1998, p. 98) is to construe H as a second-order property¹⁴¹ defined by its *causal role*, where a causal role might be

¹⁴⁰ I review some of this literature in the next chapter.

¹⁴¹ Kim characterizes second-order properties as follows: "F is a *second-order* property over set B of base (or first-order) properties iff F is the property of having some property P in B such that D(P), where D specifies a condition on members of B" (1998, p. 20).

thought of as a list of typical causes and typical effects.¹⁴² If we think of H as a second order functional property, that is, the property of having a property with causal role C, and this C is exactly the causal role that L plays (in all nomologically possible worlds), then Kim thinks we can say that $H = L$.¹⁴³

Consider the property of transparency, for example. Kim thinks that we should first construe transparency, standardly stated as *the capacity of a substance to transmit light-rays intact*, in a functional way: as the second-order property of having a first-order property that bestows the power to transmit light-rays intact to a substance that instantiates it. Next, we find an explanation, at the micro-level, for why the substance has this power. The relevant first-order property picked out in this explanation, call it L, then realizes the transparency. In this circumstance, the instance of transparency and the instance of its micro-realizer have the same causal powers. Kim thinks that the properties are then identical: transparency, it turns out, is nothing other than having L.

To summarize: Kim's solution to the causal exclusion problem is to preserve the efficacy of supervenient properties by first construing them in the right sort of way (according to their causal role), and then by identifying those properties instances with their supervenience-base by virtue of the fact that they have the same causal powers.

¹⁴² Cf. Shoemaker (2007) who characterizes such a property X in terms of its *causal profile*: a set forward-looking (typical effects of X) and a set of backward-looking causal features (typical causes of X).

¹⁴³ Alexander Rueger thinks that it would be better for Kim to draw the identification between the causal powers of H and the causal powers of L. (2006, p. 336)

4.2 *Orders vs. Levels: The Dual Hierarchy View*

The appeal to the “micro-level” must be qualified in Kim’s picture of functional reduction. He thinks that it is important to make a distinction between what he calls “orders” and “levels” if we hope to resolve his causal exclusion worries.

Kim’s distinction is not meant to suggest that orders and levels are different sorts of things altogether. The distinction is merely terminological. The terms ‘order’ and ‘level’ are both used to designate relata of, what appears to be, a *higher than* relation. If this is correct, then they qualify, in the mind of the core hierarchy theorist, as *levels*.

The purpose of the distinction is to specify the sort of hierarchy to which Kim thinks his reductionism applies. (1998, p.83) When Kim talks about higher or lower *orders*, he refers only to a hierarchy of functional properties: the property of having some other property. In other words, to instantiate a higher-order property is to instantiate the property *instantiates some lower-order property*; for example, the property of *having some property that plays such-and-such a causal role*.

Higher *levels*, in Kim’s terminology, are to be understood mereologically. Higher level properties are those properties belonging to wholes that do not belong to their parts.

Kim’s reductionism is intended to apply only to order-based hierarchies. He thinks that it is unproblematic to suggest that new causal powers appear on higher levels in the mereological hierarchy. Many, I suspect, will feel inclined to agree. A baseball has a certain mass, but that mass is not shared by any one of the baseball’s proper-parts. Conscious persons can have sophisticated philosophical thoughts, but their constituent atoms cannot.

For Kim, the realizers of high-level functional properties are not functionally reduced to lower *level* properties – i.e., micro-properties – but rather to properties found at that same higher level. In this respect, functional reduction is an intra-level relationship, rather than inter-level. The high-level, but *micro-based*, reducing properties (i.e., the realizers) are accounted for in the following way:

P is a micro-based property just in case P is the property of being completely decomposable into non-overlapping proper parts a_1, a_2, \dots, a_n such that $P_1(a_1), P_2(a_2), \dots, P_n(a_n)$, and $R(a_1, a_2, \dots, a_n)$. (Kim 1998, p. 84)




Figure 19: The Dual-Hierarchy Approach

Here P is a higher level property, a *micro-based* property, of a composite object, and the properties of that objects parts, $P_1 \dots P_n$ are the *micro-properties*. The situation is depicted in Figure 19. The idea is that a mereological aggregate instantiates both the functional property and P, the micro-based property. The parts into which the aggregate decomposes do not have P. They have other properties, viz. micro-properties, and they are related to one another in some way. Their relatedness can be complex.

4.3 *The Dual-Hierarchy View, Reduction, and Core Hierarchy Theory*

One might be tempted to think that the dual-hierarchy view can be used to relieve the tension between Reductionism and Core Hierarchy Theory brought out in the previous chapter. The mereological “levels” hierarchy, qua compositional, is perfectly consistent with Core Hierarchy Theory. The property-based “order” hierarchy, on the other hand, is not a compositional hierarchy of any sort. Decomposition does not figure into the property hierarchy since the realization relation, for Kim, holds between a functional property and a micro-based property that is instantiated in the *same individual* (Gillett 2002 calls this “flat” realization). Upwards nesting, then, is not a derivable feature of the functional-reductionist hierarchy¹⁴⁴ and so Kim can make the identification of functional and micro-based properties without conflicting with Core Hierarchy Theory’s theorems for unrestricted hierarchies.

Even so, Kim’s dual-hierarchy approach cannot be used to defend Reductionism, for his view is not Reductionist in the sense described at the outset. By separating the property-hierarchy from the mereological-hierarchy, the conflict with Core Hierarchy Theory is avoided entirely: by rejecting the slogan “everything’s in the basement!” with respect to the mereological hierarchy, Kim avoids the problem we raised for Reductionism; but at the same time, by rejecting the slogan Reductionism is abandoned. So, the dual-hierarchy approach cannot be used to defend Reductionism after all.

¹⁴⁴ But note that Kim’s “functional-reductionist hierarchy” turns out not to be a hierarchy at all, according to Core Hierarchy Theory. For Kim, realization is property-identity. As soon as it is claimed that the realizer is identical to the property it realizes, it follows from the symmetry of identity that Kim-style realization lacks the requisite asymmetric property that is necessary for being a higher than relation. Kim’s “reduction hierarchy” is therefore not a hierarchy, according to Core Hierarchy Theory.

Ironically, Kim turns out to be a kind of Holist. According to Kim, the *micro-based* properties are identical to the higher-order mental properties they realize, but are not identical to any collection of (or structuring of) *micro*-properties. This is perfectly consistent with Holism, as characterized in the introductory chapter: for Kim, there is a mereological hierarchy (the Holist agrees) and higher-level properties are not positioned at the bottommost level (the Holist agrees); higher-order properties are identical to lower-order properties, but are not identical to anything at the lower-level. Kim's functional reductionism, as it is described here, thus counts as a kind of Holism, as characterized in the introductory chapter.

This is a strange result. Kim thinks that there is a problem with non-reductivist metaphysics: the causal exclusion problem. Yet, Kim counts as a kind of Holist. Since non-reductive (i.e., Holist) *metaphysics* is the target of the causal exclusion problem, and this metaphysics appears to be the same as Kim's, we should expect that Kim's metaphysics fails to solve the causal exclusion problem (i.e., fails to avoid epiphenomenalism). Indeed, that is exactly what Stuart Glennan argues in order to advance his mechanistic approach to the layered worldview.

5.0 The Mechanistic Hierarchy

Glennan (2010) argues that Kim's reductionism fails to solve the causal exclusion problem. He accuses Kim of something called *property bias*, taking issue with the idea that higher level properties are determined *only* by base properties and appropriate nomological relations, which he says is implicit in Kim's view. Glennan thinks that this approach ignores that higher level properties are mechanistically explicable, and that

causal relationships are mediated by underlying mechanisms.¹⁴⁵ To ignore the role of mechanisms, Glennan thinks, is to invite causal exclusion worries back into the mereological hierarchy. According to him, taking special note of mechanisms, especially their *activities*, is the only way to put these causal worries to rest. I conclude, in part, by arguing otherwise.

5.1 *What is a Mechanism?*

Glennan's account of mechanism (c.f. 1996, 2002) is as follows:

A mechanism for a behavior is a complex system that produces that behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations. (2002, p. 344)

According to Glennan there is no such thing as “a mechanism simpliciter. One cannot even identify a mechanism without saying what it is that the mechanism does.” (Glennan 1996, p. 52).

Contrast this characterization with that of Machamer, Darden, and Craver (henceforth, MDC): “Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions.”

(MDC 2000, p. 3) MDC call their characterization “dualist” since on their account,

¹⁴⁵ The notion of ‘mechanistic mediation’ is also found in Craver and Bechtel’s (2007) account of so-called *top-down* causation. They argue that top-down causation refers to mechanistically mediated effects, which are understood partly as a constitutive claim and a causal claim: the causal relations, for Craver and Bechtel, are exclusively intra-level; whereas composition and constitution relations are inter-level. A similar account is provided by Max Kistler (2009), who adds a few refinements to Craver and Bechtel’s account.

entities and activities are both primitive ontological categories of the sort that can enter into a composition relation; that is, “mechanisms are composed of both entities (with their properties) and activities” (Machamer, Darden, & Craver 2000, p. 3).¹⁴⁶

In both characterizations, mechanisms are productive by virtue of their decomposition into parts that engage in productive activities. Glennan does not see the two characterizations to be at odds. In fact, when arguing against Kim, he adopts MDC’s idea that activities are the locus of productivity, rather than events, properties or laws. In this respect, he upholds what we earlier called the mechanistic, rather than Humean, metaphysical approach.

5.2 Glennan on The Dual-Hierarchy View

Glennan (2010) argues against Kim’s dual-hierarchy view by raising doubts as to whether it solves the causal exclusion problem. According to Glennan, to describe something in terms of a complex conjunction of properties and relations of its low-level parts implies that there is a configuration of those parts at the lower level. Those micro-parts, instantiating the very properties and relations mentioned in Kim’s high-level micro-based property, according to Glennan, also correspond to a complex lower level property.

I think Glennan’s point is a good one, albeit slightly misstated. The general idea is that we can concede to Kim that high-level (but low-order) micro-based properties – e.g., *being decomposable into the Ps standing R* – are instantiated by a high-level individual, expressed formally as Dx . This concession also requires that we admit the existence of an organized assortment of micro-entities, x_1, x_2, \dots, x_n , each instantiating micro-properties and

¹⁴⁶ How should we understand these ontic-primitive *activities*? According to Peter Machamer (2004, pp. 30-2), activities are what are referred to by verb expressions.

bearing the R relation to one another – e.g., *being the Ps standing R* – at the lower level. The “complex lower level property” to which Glennan refers is in fact a low-level relation:¹⁴⁷ expressed as Rx_1, x_2, \dots, x_n . If the relational-structuring of micro-entities and their properties, Rx_1, x_2, \dots, x_n , is sufficient for an effect, then the high-level micro-based property to which Kim attributes efficacy is made inefficacious. What further work does *being decomposable into the Ps standing R* do that is not already done by *being the Ps standing R*? None, it seems. Since what goes on at the micro-level of the mereological hierarchy is sufficient for what goes on everywhere else, the causal exclusion problems snaps back at Kim’s functional-reduction (i.e., intra-level identity theory) solution.

To respond, one might suggest that we should further reduce the properties of high-level composites to Glennan’s complex lower level relations, without reducing the composites themselves. That way the causal efficacy of higher level properties might be saved, via microreduction, without endorsing the fatal downward-inclusion principle. The problem, of course, is that Leibniz’ Law forbids relational properties, such as Kim’s micro-based properties, from being identified with lower level complex relations. Properties of high-level individuals are quite generally monadic, whereas the low-level complex relations are polyadic. Since one is not many, the difference implies that this further reductionism is untenable. Kim cannot avoid epiphenomenalism so easily.

¹⁴⁷ Otherwise, Glennan’s objection does not make sense. Since ‘level’ refers to mereological level in the current context, if he were to construe his complex structural property as a low-level property, then it would be a property possessed by something at the lower mereological level. But there is nothing at that level, no micro-part, that instantiates that sort of complex structure. So, Glennan’s complex structural property cannot be a lower-level property instantiated by a lower-level thing, but rather a relationship between lower-level things instantiating the properties that they do.

5.3 *Productive Activities vs. Causal Relevance*

Glennan (2010) thinks he has a way to resolve the issue. He thinks that properties, independent of the mechanisms that realize them, are quite generally inefficacious. The consequence we end off with in the previous sub-section is, in Glennan's mind, exactly correct.

Activities, for Glennan are where productivity, and hence efficacy, is found. Consider a simple example: a key opening a lock. Turning the key causes the lock to open by virtue of the key's shape and rigidity. Where Kim faces a struggle to preserve the efficacy of the key's macro-properties *shape* and *rigidity* in the presence of its microstructural properties, Glennan denies the efficacy of both. It is the *turning* of the key that is productive, and hence efficacious, in opening the lock. Properties, in fact, are only causally *relevant*, which is a notion he, following Jackson and Petit, takes to be distinct from productivity.¹⁴⁸ Since changing the micro-structure of the key, while holding its macro-structure fixed (e.g., expose the key to extreme heat and then let it cool), will make no difference to the opening of the lock, Glennan tells us that it is the macro-structure that is causally relevant rather than the micro-structure.

This is largely in agreement with Carl Craver's (2007, Ch. 6) mechanistic analysis of the causal exclusion worry. He denies that the exclusion problem applies to the mechanistic framework he and Glennan take for granted, since Kim's "argument simply does not apply to the central explanatory sense of levels in neuroscience." (Craver 2007, p. 198) Both Glennan and Craver deny the efficacy of properties altogether. Efficacy, for

¹⁴⁸ For more on this line of argument see Craver (2007, Ch. 6).

these mechanists, relates to *productivity*, and properties do not actively produce anything; mechanisms, on the other hand, are productive and hence efficacious.

Craver (2007, p. 165) thinks that the productivity of higher level mechanisms escapes causal exclusion worries because the causal exclusion argument affects only those accounts of local hierarchies that are based on a ranking of properties: Kim's realization hierarchy.¹⁴⁹ Craver, perhaps misleadingly, calls these "realization" hierarchies. Mechanistic hierarchies, he claims, are not realization hierarchies in this sense.

Here I hesitate. I worry that relying on a distinction between property hierarchies and mechanistic hierarchies bring us no closer towards solving the causal exclusion problem. The problem, I suspect, can be regenerated with respect to the mechanistic approach to the layered worldview as well. To understand why, let us first look at the mechanist's vision of the layered worldview.

5.4 *Craver and the Layered Worldview: Levels of Mechanisms*

The fifth chapter of Craver's (2007) book is devoted to an investigation of the usage of the word 'level' in the biological, especially neurological, sciences. His goal for the investigation is to shed light on the sort of "multi-level" explanation that is at work within neuroscience. He thinks that *levels-of-mechanisms* are most relevant to

¹⁴⁹ As far as I am aware, Craver accepts Kim's levels/orders distinction, as well as Kim's claim that the causal exclusion worries apply only to the intra-level (i.e., order-based, rather than mereologically-based) property hierarchy. Craver thinks (2007, Ch. 5 and Ch. 6) that realization is, as Kim claims, an intra-level relationship.

explanations surrounding spatial memory,¹⁵⁰ which he takes as a paradigmatic example of a multi-level explanation in neuroscience.¹⁵¹

He understands ‘level’ to be “multiply ambiguous” (2007, p. 163); there are levels of abstraction, analysis, behavior, complexity, description, explanation, function, generality, organization, science, theory, and mechanism. (2007, p. 164) He thinks that the application of term “requires only a set of items and a way of ordering them as higher or lower” (*ibid.*), and uses this general understanding to carve out a taxonomy (2007, p. 170) of the different senses of *level*. He then uses this taxonomy to narrow in on the sense of ‘level’ used in neuroscience. I have two minor quibbles.

First: ‘Level’ is not multiply ambiguous. Craver defines it well enough as “a set of items and a way of ordering them as higher or lower.” This is simply a set-theoretic variant of the Core Hierarchy Theory formulation: *levels are the relata of higher than*. The main source of ambiguity, it strikes me, is not found in the term ‘level’; but rather, in the *way of ordering them*. ‘Higher than’, we have seen, courts many suitors.

Second: there is a problem with his taxonomy. It is missing a rather important sense of ‘level’ in both neuroscience and philosophical metaphysics: *levels-of-realization*. Craver hastily dismisses levels-of-realization on the grounds that realization is a relation between properties possessed by the same individual. He agrees with Kim that realization

¹⁵⁰ There are four levels in Craver’s main example: at the top, you have a mouse navigating a water maze (2007, p. 166-7); lower, you have the hippocampus generating a spatial map (p. 167); lower still are neurons inducing long-term potentiation (p. 169); and at the bottom, the molecular level, you have NMDA receptors activating (p. 169). He notes that the choice of four levels is an over-simplification. But I do not dispute that it is a useful one.

¹⁵¹ There are many other examples that fit with Craver’s picture. Bechtel, for example, thinks the human eye is an example of a multi-level mechanism that corresponds to a multi-level explanation of human perception.

as an intra-, rather than inter-, level relation. But there is a different way to think of realization. Carl Gillett (2002) argues in favor of a *dimensioned* view of realization. He gives the following analysis:

(REALIZATION) Property instances $F_1—F_n$, in individuals $s_1—s_n$ (or individual s^*), realize a property instance G , in individual s^* under background conditions $\$,$ *if and only if* the powers contributed by $F_1—F_n$ to $s_1—s_n$ (or s^*), which are constituents/parts of s^* , together comprise the powers individuating G , in s^* under $\$,$ but not vice versa. (2007, p. 202)

This sort of analysis incorporates two kinds of realization: property, or “flat”, realization in which both the realized and realizing property are instantiated in the same individual, and dimensioned realization in which the property of the whole is realized by the properties of (a structuring of) its parts. Kim’s approach to realization is flat, whereas (our interpretation of) Glennan’s argument – that the causal exclusion argument generalizes from order-based hierarchies to mereological hierarchies – exploits dimensioned realization. By overlooking the hierarchy of *dimensioned* realization, Craver ignores a viable contender for levels of explanation in neuroscience.

5.4.1 Craver, Mechanisms and Globalism

Although Kim and Craver disagree about the efficacy of higher-order properties, they are in agreement with respect to the tenability of globalism: they are skeptical. Presuming that the Oppenheim and Putnam approach to be best representative of the globalist

ambition, and that the metaphysics of a global ordering is grounded in composition, Craver argues that there is no neat correspondence between levels of objects and levels of sciences (physics, chemistry, etc).

First, the atoms-to-societies ranking is incomplete. It is missing solar systems, galaxies, and other important scientific features. Societies are at the top-level of Oppenheim and Putnam's compositional global hierarchy, but solar systems are not societies. Nor, says Craver, will it do to simply add levels to the hierarchy. Solar systems, he claims, do not have societies as parts, so they cannot be higher than societies in the global compositional hierarchy.

Craver's claim is puzzling. Surely *some* solar systems have societies as parts. Solar systems have biospheric and non-biospheric planets as parts; biospheres have ecosystems as parts; ecosystems have communities as parts; and societies are a sort of community.¹⁵² From the transitivity of primitive-parthood it follows that societies are parts of some solar systems. They are not parts that cosmologists care about in their professional work, but the interests of cosmologists do not determine what counts as a part of what, in the most basic sense. Societies, therefore, fit quite nicely in a *global* compositional hierarchy stretching from galaxies down to atoms. We should only agree that the society level is ill-placed in *local* hierarchies involving distant solar systems that contain no life.

Nevertheless, Craver raises other objections to the tidy correspondence between the compositional hierarchy and the hierarchy of scientific fields. He uses neuroscience

¹⁵² Note that I use 'biosphere', 'ecosystem', and 'community' in the sense discussed in Chapter 1.

as one example. “Theories in neuroscience, in short, do not correspond to tidy levels of nature.” (2007, p. 176) Moreover, the mere presentation of Wimsatt’s wildly branching global “hierarchy”¹⁵³, Craver thinks (2007, p. 190), is evidence enough.

But note, globalism is not the view that there is a perfect, or near-perfect, fit between a global hierarchy of scientific fields and a global hierarchy of composition. That is only *one version* of globalism. Globalism, as it was characterized earlier, is simply the idea that the entire scientific domain exhibits hierarchical structure. On such a characterization, there might be *many* ways in which the entire domain can be ordered into a hierarchy; it imposes no requirement that any global hierarchy whatsoever must “match up” with existing scientific fields. Craver’s objection is restricted to something that might be called *integrative globalism*: the idea that all or many of the global hierarchies, can be integrated into one singular “master” hierarchy. I agree that integrative globalism is implausible. Globalism, the idea that the entire scientific domain exhibits a hierarchical structure as a whole, in one or more ways, remains a live option.

Even so, the existence of a globalist mechanistic or compositional hierarchy does not imply that Craver’s *localist* account of levels-of-mechanisms is problematic. Craver only means to assert that hierarchical structures should be investigated on a case-by-case basis, depending on what is “explanatorily relevant for a given phenomenon.” (2007, p. 191) Let us now examine Craver’s localist approach in greater detail.

¹⁵³ Quotations are used because, as it was noted in the previous chapter, Wimsatt’s hierarchy is not genuinely hierarchical: there are failures of the asymmetry constraint on *higher than*.

5.4.2 Craver on Levels-of-Mechanisms

Craver's view of mechanistic hierarchies begins from two assumptions, which he makes explicit:

(LM1) Levels of mechanisms are ordered compositionally. (2007, p. 184)

(LM2) Levels of mechanisms are *exclusive*, in the sense that each item appears at only one level in a given hierarchy. (*ibid.* p. 185)

When characterizing levels-of-mechanisms as a sort of compositional ordering, Craver warns against the presumption that levels-of-mechanisms are levels-of-aggregates, where "aggregates" are understood as mereological sums, irrespective of their structure and organization. Neuroscience properties, he argues, are not always aggregate-like, they depend on the spatial location of the parts. Aggregation does not depend on location. So, neuroscience levels are not levels of aggregation. (2007, p. 187)

Craver also warns against the presumption that mechanistic levels correspond to levels of spatial containment. Such a characterization, he argues, fails to differentiate between *components* and *pieces*. Pieces of mechanisms are arbitrarily specified hunks of it; components are "pieces that make identifiable contributions to the behavior of a mechanism." (2007, p. 188) He understands 'identifiable contribution' in terms of causal relevance: components make a difference to the behavior of a mechanism as a whole.

Following Stuart Kauffman (1971), Craver thinks that decomposition into components is always relative to some behavior of the system. He summarizes his positive account of levels-of-mechanisms in the following way:

Levels of mechanisms are levels of composition... the relata are behaving mechanisms at higher-levels and their components at lower levels. These relata are properly conceived neither as entities nor as activities; rather, they should be understood as acting entities... X's ϕ -ing is at a lower mechanistic level than S's ψ -ing if and only if X's ϕ -ing is a component in the mechanism for S's ψ -ing. Lower-level components are *organized together* to form higher-level components... By organization, I mean that the parts have spatial (location, size, shape, and motion), temporal (order, rate and duration), and active (for example, feedback) relations with one another by which they work together to do something. (2007, p. 189)

His diagram (Figure 20) is most helpful:




Figure 20: Craver's Depiction of a Mechanism With Three Levels (2007, p. 189)

At the top-level, we have the mechanistic behavior to be explained: S's ψ -ing. The level below it is the mechanism that explains S's ψ -ing; at that level, we see x's ϕ -ing. The bottom level is the mechanism for x's ϕ -ing.¹⁵⁴

Craver thinks his view corresponds nicely to Wimsatt's general remarks on levels: components are smaller than the mechanisms that contain them; levels-of-mechanisms are "loci of stable generalizations", and can be seen as "local maxima of regularity and predictability." (Craver 2007, p. 190)¹⁵⁵

His presentation differs from Wimsatt's in an important way. The difference, in fact, might be taken to raise trouble for Core Hierarchy Theory. The difference is that levels-of-mechanisms are not levels of entities since mechanistic components, Craver argues, are not always entities. The synapse, he gives as an example, is part axon-terminal, part dendrite, and the gap between them. (2007, p. 190)

If our Core Hierarchy Theory were expressed in terms of 'entity' rather than 'individual', then this subtle point would indeed place Craver's levels-of-mechanisms out of its reach. An individual, recall, is anything that falls within the scope of our singular quantifiers. All entities are then individuals, but not all individuals are entities. The synapse, for an example, is an *individual* in the idiomatic sense of Core Hierarchy Theory, as evidenced by our quantifying over it in this very sentence.

¹⁵⁴ Note, Craver does not mean that the contents of levels (the "relata") are not literally entities; acting entities are still entities, after all. He is merely suggesting that all mechanisms are active. For the mechanism to exist, it must be an organization of entities, and the activities they engage in. No activity, no mechanism.

¹⁵⁵ Note that Craver's approach to mechanistic levels has nothing to do with Wimsatt's waveform diagram (i.e., the scalar hierarchy) from the previous chapter.

So, Craver has no reason to complain that Core Hierarchy Theory's range to be too narrow. To do so, he must show that there are components of mechanisms that are not individuals in the idiomatic sense of Core Hierarchy Theory, which is implausible. To give the counterexample, he must mention *something* outside the range of our singular quantifiers. But by mentioning *it*, he thereby shows that *it* is, in fact, within the range of our singular quantifiers. Craver's account of levels-of-mechanisms, therefore, falls within the scope of our Core Hierarchy Theory.

5.5 *Mechanisms, Causation and Multi-level Explanation*

To give a multi-level *explanation* of some phenomena is to describe the mechanism for that phenomenon by describing (which involves quantifying over) the behaviors of its relevant components. What makes the explanation *multi-level*, rather than single-leveled, is that these behaviors require quantification into different levels of composition. To explain a phenomenon mechanistically is to give the entities and activities that make a contribution to the production of that phenomenon. If this list of entities and activities includes components of the mechanism, then giving the explanation requires quantification into different levels (e.g., there exists, at compositional level L, a mechanism m such that...).

For Craver's take on both the layered worldview and multi-level explanation, *higher than* and *is-at-the-same-level-as* must be understood as relations to mechanisms. (2007, p. 192) The question, "what's higher than what?" is only sensible if the subjects of inquiry are both components of the same mechanism. (2007, p. 191)

He thinks that this insight helps to resolve confusions regarding inter-level causation, claiming that “one ought not to say that things at different levels causally interact with one another.” (*ibid.* p. 195)¹⁵⁶ According to Craver (2007) and Bechtel (with Craver, 2007), there is no genuine *inter*-level causation: seemingly inter-level causal claims are, they say, partially a causal claim and partially a constitutive claim. This suggests, first, that downward causal-mechanistic claims are in fact, claims about mechanisms being composed of parts, such that *those parts* stand in lower level causal relations, and second, that upwards causal-mechanical claims are, in fact, about components in a mechanisms for some high-level phenomena, such that the operation of those components are relevant to the production of that phenomena. Inter-level causal claims, they claim in contrast to popular opinion, have little to do with scale.

Craver (2007, p. 195) offers an example: an elephant crushing a flea. To think of interactions at different scales as inter-level causation, he thinks, is to confuse levels-of-mechanisms with scalar hierarchies: a different sort altogether. According to Craver, the crushing causal claim is not inter-level, but rather, intra-level. What “places two items at the same mechanistic level is that they are in the same mechanism, and neither is a component of the other. Two items at very different size scales can satisfy this relationship.” (*ibid.*) If the crushing of the flea is a component of the mechanism for flea extinction, then the elephant and its victim must be at the same level of that mechanism.

In the first chapter, we said that the development of a practical hierarchical theory could be given by specifying some levels, an account of what it is to be *at* each of these levels, and a way of ordering the levels as higher and lower. We can see from the above

¹⁵⁶ Also see Craver and Bechtel (2006), and Craver (2001; 2002).

that Craver's theory of mechanistic hierarchies fits the bill. His characterization suggests the following analysis:

(@-LEVEL_M) x is at level L of mechanism M iff L is a mechanism, L is a component of M and x is a component of L.

That is, if x is a component of sub-mechanism L of M. Craver can also account for the distinction between the *higher than* relation, which is a relation between levels, and the *is at a higher level than* relation, which is a relation between the contents of levels:

(HIGHER-LEVEL_M) H is higher than L iff L is a sub-mechanism of H and H is not a sub-mechanism of L.

(@-HIGHER-LEVEL_M) For all x and y, x is at a higher level than y iff the lowest level x is at is higher than the lowest level y is at.

5.5.1 *Mechanistic Explanation and Programmatic Explanation*

Craver's account of mechanistic levels and multi-level explanation is largely coincident with other approaches to multi-level explanation. But there are differences. Consider Jackson and Petits *programmatic explanation*. Their idea is this: all the causal work is actually found at the level of fundamental physics, just as the microreductionist claims, but we can nevertheless explain higher level phenomena by appeal to *relevant* (but non-efficacious) properties.

Let us borrow Glennan's lock-and-key example to help clarify the difference between programmatic and mechanistic explanation. To explain why the lock opened, on Jackson and Petit's view, you need only mention the relevant properties, that is, the

properties which *program* for the effect. The idea is this: a property, P, programs for an effect whenever the instantiation of P implies the co-instantiation of some lower level property, P₁, which *is* efficacious to the event. In the lock-and-key case, the lower level micro-structural properties of the key do all the causal work, but the higher level macro-structure is no less relevant.

The programmatic view differs from that of Glennan, Craver, and Bechtel in the following way: these mechanists think that there is genuine causation on many levels. Mechanisms are found on many levels, and mechanisms are productive. To produce is to cause. So there must be genuine causation at many levels, on the mechanist's view *pace* Jackson and Petit. Does either approach have the upper hand? We have reason to think so. Having argued against Jackson and Petit's account of efficacy, we have every reason to reject Jackson and Petits restriction of efficacy to the bottommost level.

5.5.2 *Mechanisms and Metaphysical-MLE*

Recall that Craver's development of the local mechanistic hierarchy is founded on two assumptions: (LM1) and (LM2), we stated them earlier. Recalling the theorems of our Core Hierarchy Theory, as well as our discovery that the mechanistic hierarchy falls within their scope, it might seem that (LM1) and (LM2) are inconsistent with Core Hierarchy Theory.

Why might they seem incompatible? (LM1) implies that the mechanistic hierarchy is a compositional hierarchy. We learned earlier that *unrestricted* compositional hierarchies are necessarily nested. Nested hierarchies are, by definition, non-exclusive: items appear at multiple-levels in nested hierarchies; for example, in ECOHIERARCHY

lower levels are contained within higher levels. (LM2) requires that the levels are exclusive. (LM1), therefore, *appears* to imply the negation of (LM2).

In fact, they are compatible. In defense of Craver, his hierarchy is not an unrestricted compositional hierarchy. It is a *restricted* compositional hierarchy: for x to be at a level of Craver's hierarchy is for x to be both a primitive-part of that level and a *mechanism*; that is, for Craver, to be at a level is to be a ϕ -part (viz., a mechanistic-part) of that level. Since restricted hierarchies are not (necessarily) nested hierarchies, Craver's mechanistic hierarchy fits nicely within Core Hierarchy Theory.

Such an insight might seem to lead to yet another apparent incompatibility, this time between Craver's mechanistic hierarchies and Metaphysical-MLE. Recall the distinction between restricted and unrestricted hierarchies according to Core Hierarchy Theory. A restricted hierarchy is simply an unrestricted hierarchy in which one simply ignores (by fiat) the parts of levels that are irrelevant to the interest of the hierarchical analyst. According to Core Hierarchy Theory, higher levels include everything found at lower levels in the most basic sense of "is at". Assuming that the most basic and unrestricted sense of "is at" is what is relevant to metaphysics (*qua* its independence from human interests), it then appears that Craver's mechanistic hierarchy, the hierarchy that conforms to (LM1) and (LM2), is *not* a purely metaphysical hierarchy. Metaphysical-MLE then appears *false* under Craver's view: what justifies multi-level explanation does not appear to be the hierarchical structure of the *world*, but rather, Craver's epistemic and explanatory interests. That is, it is only because Craver's restricted mechanistic hierarchies are *useful* in achieving our epistemic and explanatory aims that multi-level explanations are justified.

I think Craver can respond to this worry. Craver's restricted hierarchy is still getting at the hierarchical structure of the world, despite the fiat restrictions that have been placed on *being at a level*. What is most relevant here is that the fiat restrictions involve *mechanisms*, which are not individuals-by-fiat; they are *real* things according to Craver and most other mechanists. Thus, while it might be true that Craver's hierarchy is a fiat-refinement of the basic nested hierarchy whose existence is implied by Craver's approach, this refined hierarchy seems no less *metaphysical* than the basic nested hierarchy which underlies it. Metaphysical-MLE is therefore preserved under the mechanistic approach to the layered worldview.

5.5.3 *Nested Hierarchies and Metaphysical-MLE*

Let us agree that Craver's mechanistic vision of the layered worldview supports Metaphysical-MLE. We can also agree that Craver's approach is only one of many. ECOHIERARCHY, recall from the first chapter, is an upwards nested hierarchy (e.g., the organism level contains cells, molecules, and atoms), and as such fails to meet Craver's exclusivity constraint expressed in (LM2) above. Although the coherence and motivation for ECOHIERARCHY is vindicated by Core Hierarchy Theory and its theorems, I will now argue that this same vindication appears to undermine Metaphysical-MLE with respect to that sort of hierarchy.

Multi-level explanation, generally speaking, is a special sort of explanation that requires quantification into multiple levels to describe all of the features relevant to the production of the phenomenon to be explained; to restrict quantification to a level is to

leave something out of the story. In the introductory chapter, a standard explanation for lake eutrophication was given as an example.

Reductionist metaphysics, recall from the outset of this chapter, is incompatible with genuinely multi-level explanation for the simple reason that all of the resources required to explain the phenomena are found at the bottommost level. Scientific explanations, according to the Reductionist, *really* only quantify into one level.

If such an argument works with respect to the Reductionist's downward nested hierarchy, then there is a worry that the same sort of argument can be given with respect to upwards nested hierarchies. Since everything within the range of the ecologist's quantifiers is found at the topmost level of the explanans hierarchy, *qua* upwards nested, it follows that ecological explanations too only quantify into one level: either the topmost level of the entire hierarchy (*viz.*, the biosphere level) or the lowest level at which the phenomenon to be explained is found (e.g., in the case of lake-eutrophication, the ecosystem level).

While it is true that in upwards nested hierarchies, some of those higher level things are *also* lower level things, this does not salvage the metaphysically robust conception of multi-level explanation enjoyed by the mechanists. There is no need to dive to the bottom of the ocean when the treasures one is looking for have floated to the top. By the same token, there is no need to bother with ecological levels at all, when everything needed to explain (for example) lake eutrophication has "floated" to the

top?¹⁵⁷ If we do bother with levels in an upwards nested ecological context, it seems, it is not because ecology is structured in a way that requires reference to levels, but rather, because it is *useful* for us to refer to them. It appears, then, that insofar as ECOHIERARCHY is an explanatory framework, Metaphysical-MLE is false; at best, the justification for “multi-level” explanation conforms to Epistemic-MLE.

Alas, Epistemic-MLE faces a similar problem. Notice that in our original example of lake eutrophication, the explanation was given without using, or mentioning, the word ‘level’ at all. This leads us to wonder: why is a levels framework *useful* at all within ecology? The world is indeed structured exactly as described by ECOHIERARCHY, but the worry is, that structure might be entirely irrelevant to our understanding of ecological phenomena. Both Metaphysical-MLE and Epistemic-MLE appear false within the context of ecological explanations. It seems that, insofar as ECOHIERARCHY is the basis of such explanations, phenomena can be described, explained, and analyzed *without* using any hierarchical terms.

To reply to this sort of worry, one might think to simply reject the nested conception of the ecological hierarchy in favor of a non-nested one, such as that presupposed by Potochnik and McGill. But, as we saw in the first chapter, Potochnik and McGill are quite skeptical about the justification and relevance of compositional hierarchies, and raise a number of problems for them in ecology. The nested conception of ECOHIERARCHY, recall, was used as a *defense* against some of their criticisms.

¹⁵⁷ In the metaphorical sense of “floated”, of course; I do not mean to suggest that green sludge explains everything.

Another reply to the worry might be to draw out another common feature of nesting: that upwards nested hierarchies are top levels of themselves. The idea is this: if the ecosystem level *is* a hierarchy consisting of all levels below it (not to forget the ecosystems composed of those things found at lower levels), then we cannot avoid an implicit appeal to hierarchical levels when offering ecosystemic explanations that utilize, for example, the nitrogen cycle. Since the molecular level is included in the ecosystem level on the nested approach, any reference to molecules within an ecosystemic explanation (e.g., for lake eutrophication) involves an *implicit* reference to lower levels, by virtue of those levels being part of the higher level as well. So it might appear perhaps, that Metaphysical-MLE is true with respect to ECOHIERARCHY after all.

We should remain skeptical. Even with the concession that the ecosystem level contains all lower levels as well as the contents of those levels, that does not imply (or even suggest) that Metaphysical-MLE is true. If the ecological hierarchy is upwards nested, as in ECOHIERARCHY, the topmost level of the hierarchy must be such that it has all the resources needed to completely explain the phenomenon of interest without an appeal to “levels” or any other hierarchical term: when we restrict our quantifiers to that which is contained within the topmost level, the contents of the entire ECOHIERARCHY remain within reach; not so with Craver’s mechanistic levels – by restricting one’s quantifiers to the topmost level of a mechanism, one obtains nothing but a big black box of mystery (see Figure 20 above). What you obtain is a puzzle: a phenomenon that needs explaining.

Although the matter is hardly settled, it appears that Core Hierarchy Theory supports skepticism regarding Metaphysical-MLE with respect to ECOHIERARCHY.

The argument for Multi-Grade Holism, the central thesis, of this dissertation is then complete. By assuming that Core Hierarchy Theory is the best (i.e., simplest and most general) account of hierarchies, it provides strong support for Realism, Holism, and now we see that it appears that Metaphysical-MLE is false with respect to the ecological hierarchy. With respect to the mechanistic hierarchy, on the other hand, Metaphysical-MLE appears justified by virtue of the fact that Craver's restriction on *being at a level* is metaphysical in nature, and ensures that the mechanistic hierarchy is non-nested. Since the problem for Metaphysical-MLE is restricted to nested hierarchies, which I maintain are the most basic sort of hierarchical structure, it might follow that the *basic* hierarchical structure of the world fails to justify multi-level explanations, but it nevertheless holds true that *restricted* mechanistic hierarchies succeed in justifying multi-level explanations on a metaphysical basis.

5.6 *The Causal Exclusion Problem Revisited*

So far, it appears that the mechanistic approach to the layered worldview is superior to Kim's. Kim's dual-hierarchy view, recall, succumbs to a slightly recast causal exclusion problem. Mechanists, we noted earlier, generally take their view to be beyond the reach of causal exclusion arguments by virtue of the fact that their view of efficacy is grounded in productive activities. Activities are obviously productive, they say, no matter what compositional level they are on; the causal exclusion problem, they think, applies only to realization hierarchies.

In "flat" realization hierarchies, functional properties appear inefficacious since their *micro-based* realizers are sufficient to accomplish whatever it is that functional

properties accomplish. Kim's solution to the problem is to identify functional properties with their micro-based realizers. As Glennan argues, Kim's solution is not enough: an analogous exclusion problem applies to Kim's mereological hierarchy. Taking for granted Gillett's notion of *dimensioned* realization, what Kim calls higher level micro-based properties – e.g., *being decomposable into the xs instantiating $P_1 \dots P_n$ and standing R* – are in fact dimensionally realized by a complex relation involving lower level micro-properties – e.g., *being the xs instantiating $P_1 \dots P_n$ and standing R* . Since the complex of low level micro-properties are sufficient to accomplish whatever it is that higher level micro-based properties are said to accomplish, the higher level micro-based properties thus appear excluded from being efficacy.

The mechanist feels safe from these worries: elephants crush, planets rotate; crushing and rotation are productive. Such feelings might be hasty. Their view, despite what they say, might in fact succumb to (a version of) the causal exclusion problem.

5.6.1 *Mechanistic Hierarchies are Dimensioned Realization Hierarchies*

We have what we need to plausibly argue that a mechanistic hierarchy is a sub-species of realization hierarchy, taking for granted Craver's account of the mechanistic hierarchy, as well as Gillett's aforementioned account of dimensioned realization.

Mechanisms are organized entities and activities that are productive to some phenomenon of interest. The mechanism for that phenomenon is often multi-level in the sense that the mechanistic behavior, S's ψ -ing, decomposes into sub-mechanisms, the xs's ϕ -ing, at lower levels. If S is ψ -ing, then S has the property *is ψ -ing*, and if the xs are ϕ -ing, then the xs have the property *is ϕ -ing*.

The mechanist's causally efficacious mechanisms are then subject to Gillett's analysis. Insofar as S's ψ -ing has powers that are individuating of ψ -ing instances *in virtue of* the powers contributed by the property instances belonging to its constituents (e.g., the xs instantiating *is φ -ing*), but not vice versa, it Gillett's dimensioned account, REALIZATION, implies that the mechanistic hierarchy is a sort of realization hierarchy. In short, mechanistic decomposition meets the conditions for dimensioned realization.

The mechanistic and realization hierarchies are not coextensive hierarchies since there are realization hierarchies that are non-mechanistic. However, since every mechanistic hierarchy accounts for higher level causal efficacy in terms of the efficacy of its components, the inter-level mechanistic relationship implies a realization relationship as well. So, every mechanistic hierarchy is a realization hierarchy, but not vice versa.

Mechanists might scoff at the mentioning of *powers* in the above argument, and might try to motivate a rejection of it on those grounds. Mechanists talk of activity, not power.

Such a response seems promising, but must be developed carefully. Although mechanists make no *explicit* mention of powers, they do talk about properties. One of the more popular views of properties found within the debates over realization is the power-based theory of properties, which can be characterized by the claim that properties are individuated by the causal powers they bestow.¹⁵⁸ If mechanists have a different account of property, then we must hear it. Only then can we evaluate whether avoiding the use of 'power' can help the mechanist distance herself from the realization hierarchy. Under the

¹⁵⁸ See Shoemaker (2007) for a very useful discussion.

scope of our present assumptions, however, the conclusion that S's ψ -ing is (perhaps, mechanistically) realized by its sub-mechanisms has not been made any less compelling.

And so, the mechanist approach to multi-level causation succumbs to the causal exclusion argument. By Craver's admission, the causal exclusion argument applies to the realization hierarchy; and as we have now discovered, the mechanistic hierarchy is a sort of realization hierarchy.¹⁵⁹

5.6.2 *Causally Excluded Activities and Mechanisms*

There is another way to argue that levels-of-mechanisms succumb to the causal exclusion problem. We simply adjust its initial formulation. Recall that Glennan diagnoses the problem with causal exclusion as a category error: properties are not the sorts of things that can be efficacious; productive activities are. Properties, the mechanist says, are only *relevant* to the production of the phenomenon.

The idea is simple, but it is not obviously correct. On the contrary, shifting efficacy from properties to activities strikes me as an ineffective solution to causal exclusion worries. Let a be an activity of some mechanism, m , let $R[x$'s a_1 -ing... x 's a_n -ing] be the acting entities of its lowest-level sub-mechanisms with a complex relation R holding between them, and let e be some phenomenon that is produced by m .

Now consider the following: if m 's a -ing produces e , and m 's a -ing mechanistically-decomposes into $R[x$'s a_1 -ing... x 's a_n -ing], then e is also produced by $R[x$'s a_1 -ing... x 's a_n -ing]. Why? Since m 's a -ing decomposes into x 's a_1 -ing... x 's a_n -ing,

¹⁵⁹ An interesting corollary, Realization hierarchies, pace Craver, might then best capture the sense of 'level' we see throughout the life sciences and special sciences.

if m 's a -ing suffices for the production of e , then $R[x$'s a_1 -ing... x 's a_n -ing] is presumably sufficient for the production of e : presumably, any world where you have $R[x$'s a_1 -ing... x 's a_n -ing] is a world where m is a -ing. Following Kim, we do not want to say that e is *overproduced*, and so the vexing question resurfaces: why posit a higher level mechanism when a complex of inter-related lower level mechanisms will do? Since $R[x$'s a_1 -ing... x 's a_n -ing] is sufficient for e , it appears that higher level mechanisms are superfluous; they appear inefficacious and thus unproductive. This is a terrible result for any mechanistic philosopher of science who favors a metaphysically robust multi-level approach to explanation.

Consider a simple example: a bicycle. Mechanists should want to say that the chain's *pulling* on the gear is productive of the rear wheel's turning. They should also want to say that the chain and gear mechanistically decompose into sub-mechanisms: each link of the chain *grabs* a tooth of the gear. The mechanism decomposes even further than that, but let us bottom out here for sake of simplicity. The point is, the plurality of chain-links grabbing teeth is a plurality of inter-related sub-mechanisms, and this plurality of sub-mechanisms, related as they are, is productive of the rear wheel's turning.¹⁶⁰ Since the plurality of sub-mechanisms (each link grabbing a tooth), related as they are, is sufficient to make the wheel turn, the higher level mechanism (*viz.*, the chain's pulling on the teeth) appears superfluous. To avoid the *overproduction* of the effect, it seems, we must give up the efficacy of the higher (or perhaps lower?) level. The causal exclusion problem is not as easily pacified as Glennan and Craver assume.

¹⁶⁰ That Craver draws a "production arrow" at the bottom level of Figure 20(above) is certainly suggestive of this. It suggests that, for each mechanism at the second level, there are similar production arrows exiting the mechanism boundary at the bottom level. What, then, is the higher level production arrow accomplishing that the lower level arrows are not? Nothing, perhaps.

One way to solve the problem might be to issue identity statements relating the higher level mechanism with the complex of inter-related lower level sub-mechanisms. But there is a familiar problem with such a suggestion. Identity is a one-to-one relationship: one thing is not many things. The relationship between the singular a , and the plural, $a_1 \dots a_n$, is therefore not identity. Perhaps, then, it is something like identity. In the next chapter, the causal exclusion argument is the main focus.

6.0 Conclusion

Our inquiry into the metaphysics of multi-level explanation has led us to a problem with causation in multi-level contexts. We began by laying out our metaphysical options: what we called the Humean approach and the mechanistic approach. We then undertook a rigorous investigation of Kim's take, a mostly Humean take, on the layered worldview. We saw both that his supervenience view fails, and that his functionalist *dual*-hierarchy view succumbs to the very problem it was intended to solve: the causal exclusion problem.

The mechanistic approach to the layered worldview is, by contrast, highly-developed and well-supported by actual scientific explanation and practice. One commonly mentioned feature of mechanistic explanation, and the mechanistic view of the world, is that it facilitates a view of multi-level explanation. The mechanistic approach to the layered worldview, we argued, supports the metaphysical justification for multi-level explanation: it satisfies Metaphysical-MLE.

Not so for the ecological hierarchy. Since the compositional hierarchy, ECOHIERARCHY, is upwards nested, everything at any level is contained within all

higher levels; the explanatory resources one needs to explain a phenomenon at level L is, thus by default, contained in L. Since we can restrict our quantifiers to a single level without losing anything relevant to the explanation of the phenomenon, it appears that the metaphysics of upwards nested hierarchies (much like Reductionist hierarchies) does not justify our describing, explaining, and analyzing things using hierarchical terms. Metaphysical-MLE, therefore, appears false with respect to ECOHIERARCHY and, taking Core Hierarchy Theory for granted, all other unrestricted compositional hierarchies as well.

To end, we raised a familiar problem for the mechanistic approach to the layered worldview. What purportedly makes mechanistic philosophy different from standard Humeanism is its reification of *productivity*. That difference, it seems, is unhelpful. There are compelling reasons for thinking that the mechanistic notion of higher level productivity succumbs to (a version of) the causal exclusion argument. If mechanists hope to advance their metaphysics on the grounds that it solves causal worries in multi-level contexts, they must tell us what is wrong with the causal exclusion argument. Insofar as Kim's dual-hierarchy succumbs to the problem, and mechanistic philosophy hopes to do better than Kim, mechanists owe us a story. But what could that be? In the next chapter, our attempt at answering this question reveals deep and profound features of the nature of causation that have hitherto been unexplored.

Chapter 6

Problems in the Causal Basement

1.0	INTRODUCTION.....	291
2.0	THERE IS NO HIGHER LEVEL: THE ELEATIC OBJECTION	292
2.1	<i>THE CAUSAL EXCLUSION ARGUMENT: SIMPLY STATED.....</i>	<i>292</i>
2.1.1	<i>The Problem of Causal Drainage.....</i>	<i>295</i>
2.1.2	<i>The Diffusion of Exclusion: Causal Selfishness</i>	<i>298</i>
2.1.3	<i>The Diffusion of Exclusion: The Triviality Response.....</i>	<i>301</i>
2.1.4	<i>The Diffusion of Exclusion: The Near-Identity Response</i>	<i>303</i>
2.2	<i>THE GENERAL ELEATIC ARGUMENT.....</i>	<i>306</i>
2.2.1	<i>The Eleatic Principle</i>	<i>307</i>
2.2.2	<i>The Higher level Does Not Make a Difference.....</i>	<i>308</i>
3.0	THERE IS NO REALIZATION: THE ELEATIC PROBLEM ONCE AGAIN	310
3.1	<i>FOUR TESTS FOR DIFFERENCE-MAKING</i>	<i>310</i>
3.1.1	<i>The Quinean Test.....</i>	<i>311</i>
3.1.2	<i>The Modal Juxtaposition Test</i>	<i>312</i>
3.1.3	<i>The Contrastivist Test.....</i>	<i>314</i>
3.1.4	<i>The Refined Contrastivist Test.....</i>	<i>317</i>
4.0	RESURRECTING THE REALIZATION HIERARCHY	318
4.1	<i>THE REIFICATION OF REALIZATION: IT MAKES A DIFFERENCE.....</i>	<i>318</i>
4.1.1	<i>The (Causal) Disconnection Between Levels</i>	<i>319</i>
4.1.2	<i>Co-opting the Contrastivist Test.....</i>	<i>321</i>
4.2	<i>THE REIFICATION OF THE HIGHER LEVEL: IT MAKES A DIFFERENCE.....</i>	<i>323</i>
4.2.1	<i>The Contrastivist Strategy Revisted.....</i>	<i>323</i>
4.2.2	<i>Implementing the Contrastivist Strategy.....</i>	<i>324</i>
4.2.3	<i>The Problem for Property Realization.....</i>	<i>326</i>
4.3	<i>ELIMINATING THE ELEATIC PRINCIPLE.....</i>	<i>328</i>
5.0	IMPLICATIONS FOR THE CAUSAL EXCLUSION ARGUMENT	331
5.1	<i>WEIGHING THE OPTIONS.....</i>	<i>332</i>
5.2	<i>INTERPRETING THE RESULT.....</i>	<i>335</i>
6.0	CONCLUSION	337

1.0 Introduction

The defense of Realism with respect to the ecological hierarchy depends on whether or not there really are ecosystems, communities, organisms, and the like. If there are, then there are ecological hierarchies. If there are not, then Core Hierarchy Theory tells us that there are no such hierarchies.

One might then present two metaphysical worries: that higher level things do not exist; and that *higher than* relations do not exist. In this chapter, I defend Realism against two arguments that lend credence to these sorts of worries. The first is quite famous. It is a simplified variant of Jaegwon Kim's (1993, 1998, 2005) causal exclusion argument. The second is a more general concern that lends credence to *both* worries: the general Eleatic argument. In short, the argument is as follows: those things that do not make a difference, as a matter of metaphysical principle, do not exist. Higher level things (Beckermann, 1997) and the inter-level relations that hold between them (Wrenn 2010) do not make a difference. Therefore, there are no higher level things and there are no inter-level relations.

In this chapter, my primary purpose is to solve these causal worries. I do three things. I explore some resolutions to the causal exclusion argument. I argue that realization, and higher level things *do*, in fact, make a difference to the world. And I argue that, even if they do not, we have good reason to reject the metaphysical principle that calls for their elimination.

2.0 There is No Higher level: The Eleatic Objection

In what follows, I will examine the arguments against the reification of higher level property instances. I begin with Kim's infamous causal exclusion argument, followed by Beckermann's general Eleatic argument.

2.1 *The Causal Exclusion Argument: Simply Stated*

The causal exclusion argument can be stated most simply with two premises. The first premise can be called *the principle of causal sharing*.

[CSHARE] If x is realized by y, then everything brought about by x is also brought about by y, its realizer.

$$\forall x \forall y (x/y \rightarrow \forall e (x \gg e \rightarrow y \gg e))^{161}$$

CSHARE derives its plausibility from some of the standard tenants of physicalism: the causal closure thesis. Since the physicalist thinks that everything that occurs in the macro-world occurs by virtue of what goes on in the micro-world, all macro-events are, in some sense, determined by lower-level events. Many physicalists take this to imply that every event has a completely sufficient micro-physical cause. Causation at higher levels, these physicalists think, must be underwritten by causation occurring at lower-levels; that is, presuming that there is such a thing as higher level causation. Given this commonly presumed connection between higher and lower-level causation, many conclude that lower-level realizers must then be causally sufficient for anything its higher level realization might cause. Such is the nature of causal underwriting, at least according to common physicalist intuitions.

¹⁶¹ Where 'x/y' is 'x is realized by y', and 'x >> e' is 'x brings about e'.

The second premise of the exclusion argument, call it the *principle of no-overdetermination*, is a rather strong claim.

[NOD] If x is realized by y, then everything e that is brought about by y is not also brought about by x. $\forall x \forall y (x/y \rightarrow \forall e (y \gg e \rightarrow \neg x \gg e))$

What NOD tells us is that lower-level events *exclude* their higher level realizations from being causally efficacious. As noted above, the claim seems a bit strong: if water is realized by a structuring of H₂O molecules, NOD implies that drinking water does not bring about the quenching of thirst, because the H₂O molecules are sufficient for the quenching. The same thing can be said with respect to many everyday activities.

To motivate NOD, its defenders offer the following sort of consideration: if lower-level L is sufficient for whatever its realization H is sufficient for, as CHSARE implies, then it seems intuitive to say that the *real* work is being done at the lower level. Drinking water, for example, does not appear to do anything that is not accomplished by its realizer, say, H₂O molecules entering the body. But if all the work is being done at the lower-level, then it appears that there is no causal work left for H to accomplish. H, it seems, makes no difference to the world that is not already made by L. To maintain that H makes a difference, while accepting that L makes *all* the difference, is to suggest that there is more causation in the world than necessary. Insofar as one desires a parsimonious metaphysic, one should not posit more causation than necessary. Thus, insofar as we feel sympathy for parsimony, NOD appears warranted.

This poses a serious problem for non-reductive physicalist: CHSARE and NOD jointly entail that there is no causation at higher levels. If H is realized by anything at all, then H is inefficacious. It is inert. Higher level things thus cannot stand in causal relations

or enter into genuine laws. Drinking water, then, quenches no thirst! We can state this conclusion as follows:

[Exclusion Conclusion] If x is realized by y then x is causally inert (there is no e such that x brings about e).

$$\forall x \forall y (x/y \rightarrow \neg \exists e (x \gg e))$$

Moreover, if higher level events cannot stand in causal relations, then it is hard to see how higher level properties and events can figure into genuine laws. Given the exclusion conclusion, higher level things are made to seem causally and scientifically irrelevant. Without causal and scientific relevance, however, the very existence of higher level things is called into question: intuition tells us that we should never posit irrelevant things. Given the conditions for hierarchy reification, realization hierarchies (e.g., the mechanistic hierarchy, as it was argued in the previous chapter) fail to provide a robust account of the layered worldview unless one can justify a rejection of the exclusion argument.

Obviously, non-reductivist response options are few: they must reject CSHARE, NOD or both. All three options carry unfortunate consequences, however. On the one hand, rejecting CSHARE and accepting NOD entails that realization sometimes results in the production of novel causal powers and causal relations not found at lower levels. Realization seems to give rise to ontological emergence in such cases. In fact, the non-reductivist would be well-advised to hold realization to be emergent in all cases. The typical, non-emergent, higher level products of realization are still subject to the principle of NOD, effectively losing their efficacy; and with that gone, their very right to exist.

If the non-reductivist were to reject NOD and accept CSHARE, on the other hand, they would be forced to accept what many take to be an offensive sort of systematic overdetermination. Genuine overdetermination is supposed to be rare; it is not often that two bullets fired from different pistols simultaneously penetrate a victim's heart, competing equally for being *the* cause. Such rarity is inconsistent with a wholesale rejection of NOD, however; for then it will turn out that overdetermination is *everywhere*.

The third option is, of course, to reject both. Rejecting both principles with respect to every instance of realization might be overkill when rejecting one will do; but then again, choosing which principle to reject on a case-by-case basis hardly seems any better. Nevertheless, whichever route is ultimately chosen, one thing is certain: if you accept CSHARE, you must reject NOD.

Given the mystery of emergence and the promiscuousness of overdetermination, all three of these response options seem unpalatably sour. Even so, the non-reductivist had better get used to the taste. An argument advanced by Ned Block (2003)¹⁶² shows that an acceptance of both CSHARE and NOD lead to an absurdity.

2.1.1 *The Problem of Causal Drainage*

Block points out that, as far as we know, there might not be a fundamental level at all; or more specifically, there might be infinite realization. This possibility, however, is incompatible with the conclusion of the causal exclusion argument. If nothing is fundamental (i.e., there is an infinite chain of realization¹⁶³), then the exclusion

¹⁶² The argument can be found within Schaffer (2003) as well.

¹⁶³ Stated more formally, $\forall x \exists y (x/y)$.

conclusion entails that there is no causation anywhere.¹⁶⁴ Since it is an open question whether or not there is a fundamental level, but, according to Block, it is not an open question as to whether there is causation, there must be something wrong with Kim's argument: either CSHARE or NOD (or both) must go.

In response, Kim (2003, p.168, p. 175) suggests that there is a way to block the drainage. Even if you think that mental properties do not reduce to neurological, biological, or physical properties, recent successes in science appear to support reduction: we have discovered, for example, that water = H₂O and (perhaps) that thermodynamics reduces to statistical mechanics.¹⁶⁵ This suggests that the properties and events competing for causal status attain a reductive nature as soon as we get down deep enough; at some lower-level, they can be identified with one another. And as soon as we reach a level at which the properties or events in competition for being cause can be identical, the drainage will bottom out at that level. Such a bottoming-out will then allow causation to take root in the world, even if the entities instantiating the identified properties can be infinitely decomposed. Absurdity averted, or so Kim thinks.

In fact, the response is weak. There are multiple ways of formulating Block's worry and Kim's reply is only effective against one such formulation. In his original formulation, Block understands the possibility of non-fundamentality as nothing more than the possibility of *infinite supervenience*. Since identity is a supervenient relation Kim expects to make his *a priori* claim about event and attribute identities without facing too much controversy; the matter, he thinks, will be left in the hands of empirical science.

¹⁶⁴ More formally, $\neg\exists x\exists y(x \gg y)$.

¹⁶⁵ The case of thermodynamic reduction is, in fact, not so straightforward. It is a contentious issue.

Notice, however, that CSHARE and NOD contain no reference to supervenience. They contain only the notions of *realization* and *brings about*; and the notion of non-fundamentality is taken to be infinite realization, as opposed to infinite supervenience. Realization, we learned in the previous chapter, has a dimensioned and flat interpretation. We are using the dimensioned interpretation for the specification of the non-reductivist global and local hierarchy. This subtle difference in set-up, however, is enough to block Kim's response; as we have seen, if some H is (dimensioned) realized by some L, then the asymmetry (and irreflexivity) of realization implies that H is not identical to L, pace Kim. Kim's reductionist move is unavailable unless he stands firm against infinite (dimensioned) realization; that is, accept finite (dimensioned) realization as a brute fact about the world.

But even this is not enough; he must also state, even more boldly, that realization *necessarily* bottoms-out. Any world in which there is causation as well as infinite realization is *ipso facto* a world where either CSHARE or NOD is false. But why should the truth of these two principles turn only on whether or not realization is infinite: what explains the "*ipso facto*"? Does infinite realization preclude micro-determination? It should not. But if it does not, then why does it preclude CSHARE? I see no reason why it would. Does infinite realization suddenly pave the way for massive overdetermination, making it less offensive? Again, I cannot see why it would. Until someone convincingly *argues* that realization does bottom out, I think we have good reason (though we do not yet know what it is) to cast general suspicion on the causal exclusion argument: either CSHARE or NOD must be false. The trouble is simply deciding which, and why.

2.1.2 *The Diffusion of Exclusion: Causal Selfishness*

Like most physicalists, realization physicalists think that everything happens *by virtue* of the micro-physical. However, by rejecting CSHARE, the non-reductivist will admit a case where a causal relation's obtaining at a higher level in fact *excludes* its holding at the lower-level: there is an effect that is not caused by virtue of micro-physical events. The problem with this, as noted earlier, is that it conflicts with the causal closure thesis, one of the central theses of non-reductive physicalism. Preserving higher level causal efficacy at the expense of standard physicalist views hardly seems like a profitable trade for the realizationist.

But is such a trade really necessary? There are multiple ways the physicalist might characterize the thesis of causal closure, and some interpretations might suggest a way to reject CSHARE while upholding closure. One way to characterize closure, albeit weakly, is to say that everything that occurs, occurs *by virtue* of micro-physical occurrences. A slightly stronger characterization might tell us that causation, everywhere, occurs *by virtue* of microphysical causal occurrences. The strongest interpretation, of course, would convey the thought that every event is ultimately *caused* by micro-physical events; that every event has a completely sufficient physical cause.

Notice that the last characterization is the only one that is explicitly causal. While it cannot be denied that 'by virtue of' could be interpreted causally in the previous characterizations, we need not commit ourselves to such an interpretation. The phrase is clearly quite vague. Instead we might choose a more general interpretation, where 'by virtue of' only picks out the *determination* relation rather than causation. By making a distinction between determination and causation, we might choose to recast the causal

closure thesis, not as the claim that all effects are *caused* by micro-physical states, but merely that all effects are *determined* by micro-physical states. Determination can be causal. But there might be other sorts of determination as well. Take realization, for instance. It is not at all implausible that realization is a sort of determination.

Such a reformulation of the closure thesis is nothing new. Philosophers of science who are skeptical about the role of causation within physics, among many others, should be happy to interpret closure as expressing a fact about determination (insofar as they think anything is determined and not mere probability) rather than causation. Skepticism about causation in general, after all, does not entail skepticism about determination.

The determination interpretation of the closure thesis remains silent on the existence of inter-level causal relations between micro-physical causes and macro-physical effects. Yet, it seems to perform just as well when it comes to capturing the general point about the micro-physical world that physicalists really care about: that the micro-physical world is, in some sense, responsible for every event that occurs.

By understanding this responsibility in terms of determination rather than causation, the non-reductivist then has a way to respond to the exclusion argument: simply reject CSHARE. By doing so, the physicalist will admit that there are events that are brought about by higher level occurrences and *not* brought about by micro-physical occurrences. But this is perfectly consistent with the idea that micro-physical states nevertheless *determine* what is brought about by those higher level occurrences. By recasting causal closure as a thesis about determination rather than causation, the non-reductivist can avoid the troublesome inconsistency.

Even so, there are still major obstacles preventing the outright rejection of CSHARE in this sort of way. For one, when we replace the occurrence of ‘brings about’ in CSHARE with ‘determines’, the principle clearly cannot be rejected, given the new understanding of causal closure. And if we make the same terminological replacement with respect to NOD, the exclusion argument will then entail that there is no higher level determination, which is just as bad as saying that there is no higher level causation. Block’s argument will still entail something the physicalist considers an absurdity: that there is no determination anywhere. The modified principle of NOD, therefore, must also be rejected if one takes this tack in rejecting CSHARE.

Rejecting the modified principle of NOD might be less offensive than rejecting the original ‘brings about’ version, however. When two assassins simultaneously pierce the heart of the same target, the death of that target is not just overdetermined, it is *over-caused*. The intuition supporting the original causal version of NOD seems to rely more on worries regarding massive *over-causation* than overdetermination.

Overdetermination, when divorced from over-causation, does not seem as problematic. While over-causation is indeed supposed to be rare, *over-determination* need not be. In fact, if we accept deterministic causation, then over-determination will turn out to be trivially widespread: any time the first node of a linear causal chain is deterministically sufficient for each subsequent event in the same chain, each of those subsequent events are equally guilty of overdetermining the final event in the chain. Strictly speaking, in a linear chain, the first event was sufficient for, and hence determines, the last; but the same can be said for the intermediating nodes as well. Causal chains are not the only example of vacuous overdetermination, however; realization

chains are just as effective in demonstrating the same point. Non-reductivists are generally in agreement that the instantiation, H, of a higher level property is determined by its lower-level realizer, each step in the realization chain connecting H to its most basic realizer R must then overdetermine the realization of H. Inter-level overdetermination is vacuous.

But there is another, more serious, obstacle preventing a rejection of CSHARE. The idea that there are effects without a micro-physical cause, for some reason, strikes many as being wildly implausible. And even when we entertain the possibility, and recast the closure thesis as suggested, we are still left with the question of how higher level causation can operate without micro-physical causation doing the obvious work; or more plainly, how it is that the micro-physical occurrences can determine without also causing. Until this mystery is solved, a rejection of NOD will no doubt be the more popular option when responding to Kim's exclusion worries.

2.1.3 The Diffusion of Exclusion: The Triviality Response

Assuming that non-reductivists would endorse a widespread rejection of the principle of NOD, how one should go about motivating this rejection becomes the important question. One common strategy is call into question the *genuineness* of the purported overdetermination.

One way to argue that causal exclusion-related overdetermination is ingenuine is to take a standard example of overdetermination, use it to create some sort of general analysis of genuine overdetermination, then use that analysis to show that causal-exclusion related overdetermination fails to meet some necessary condition for

genuineness. Jesper Kallestrup (2006, p. 472) has explored this sort of strategy using the aforementioned example of the overdetermined assassination. In the standard case, when the two bullets simultaneously pierce the heart of the unfortunate victim, we typically acknowledge that either bullet is sufficient to bring about the effect in the absence of the other. This example accords with the following notion of overdetermination:

[OD] e is overdetermined by H and L iff (i) H is sufficient for e, (ii) L is sufficient for e, (iii) if H occurred without L, e would have occurred, and (iv) if L occurred without H, e would have occurred.

In Kallestrup's view, the non-reductivist should accept that inter-level overdetermination meets (i)-(iv), but nevertheless fails a condition for genuineness: non-vacuity. By assuming that realization is supervenient, he argues that one of these conditions is trivialized: if H supervenes on L, for example, (iv) becomes vacuous. Since the condition is vacuously satisfied, the overdetermination is deemed inoffensive.

I hesitate to endorse this response, however. It is not clear why a vacuous satisfaction of one of (i)-(iv) makes the overdetermination to be less problematic. If the overdetermination were completely vacuous, then all of (i)-(iv) should be vacuous. But in the exclusion case, only one is vacuous. The overdetermination found in the exclusion argument, then, is not *entirely* vacuous.

Moreover, the *partial* vacuity does not change the fact that L is independently sufficient for e but not vice versa: H is not part of the causal chain that connects L to e, and since the direction of determination progresses from lower to higher, H's causing e depends on L's causing e. That is, everything is being determined by the lower-level,

independently of any higher level causation. So then what's the point of higher level causation? It doesn't seem to make a difference to anything, and that is the only motivation NOD seems to need. It does not matter whether or not the overdetermination is partially vacuous in order for overdetermination to seem problematic.

2.1.4 The Diffusion of Exclusion: The Near-Identity Response

A simpler way to spin-doctor the overdetermination result in the realizationist's favor is to maintain that realizations are not identical to their realizers, but are nevertheless "nothing over and above" those realizers. In such a scenario, it becomes natural to think that whatever causal relations higher level events happen to stand in, those relations are likewise nothing "over and above" the causal relations at lower-levels.

The problem with this response is that it appears to unnecessarily obfuscate the issue. The idiom "nothing over and above" is notoriously vague and quite well-despised by philosophers who cannot make sense of it. Clearly it does not mean 'nothing other than'. If it did then realization would be nothing other than identity. But 'nothing over and above' is *not* an expression of identity. If it were, then there would not be any need for scare quotes when putting the expression to use.

That the realization hierarchy is upwards-nested might help us finally make sense of it. In a nested hierarchy, such as ECOHIERARCHY, the parts are at the same level as the whole. The whole is, quite plainly, *not* over and above the parts in the realization hierarchy. For every level at which the realized thing resides, the realizer is there along with it.

Nested hierarchies aside, we do not need to put ‘nothing over and above’ into precise terms to reject the principle of NOD. We already know NOD must be false; we merely need to give a sketch of an explanation as to why.

To fill in the sketch, we might compare other examples of relations associated with the phrase ‘nothing over and above’, looking for common patterns. Finding examples is the easy part: the notoriety of the expression is well-founded. Everyday objects, it is often said, are “nothing over and above” their parts or constituents. Again we are faced with this mysterious notion.

What constitution, composition, realization, and the correspondence relation between higher and lower-level causal relations (whatever that relation might be) seem to have in common is the property of *identity-likeness* (Baxter 1988, Lewis 1991). Within these sorts of cases, x’s being nothing “over and above” y seems to be something quite *like* identity, even if it falls a bit short of what we standardly call identity.

Despite its shortcomings, most notably its vagueness, I think identity-likeness is enough to defuse the causal exclusion argument. Although identity-likeness is by no means a precise notion, it nevertheless explains why the sort of overdetermination cited in NOD is not *genuine*.

In classic cases of genuine overdetermination, the overdetermining causes are in no way identity-like. If the two bullets that simultaneously penetrate a victim’s heart were to have stood in an identity-like relation, then the death would not seem to have been

genuinely overdetermined.¹⁶⁶ That is, if the two bullets are *almost* the same bullet – e.g. bullet₁= the bullet you see flying from the gun-barrel and bullet₂= the sum of undetached bullet-parts – then the two causal relations holding between each bullet and the death-event should likewise be seen as virtually identical. Although, strictly speaking, there are two causal relations, those relations are virtually one. That there are two (barely) distinct relations turns out to be nothing more than a technicality bearing very little ontic significance.

The ontic insignificance of this technicality can be argued for, given the shared identity-likeness of composition, constitution, and realization. The argument is this: worrying about higher level events entering into “overdetermining” causal relations on the grounds that their lower-level realizers are getting the job done on their own is analogous to worrying about overfeeding your cat on the grounds that the sum of your parts will have fed it as well. But that is silly. Your feeding the cat (nearly) *is* your parts feeding the cat. Since such overfeeding is not the sort of thing pet-owners need to worry about, the same holds for the physicalist. The principle of NOD, being akin to something very silly, should therefore be rejected. Being akin to silliness is a weaker form of *reductio*, but it is a sort of *reductio* nonetheless.¹⁶⁷

Although there is more to be said for the causal exclusion argument, I think we have said enough to move forward. The problem with the argument is the principle of

¹⁶⁶ Where identity-likeness, as a minimal requirement must involve spatiotemporal co-location; in this sense it is much more than a sort of similarity.

¹⁶⁷ Again, I must emphasize, this solution is available to the non-reductivist even if the notion of *identity-likeness* turns out to be incoherent. All that matters is that realization has the same ontic innocence as composition. Since overdetermination is innocent in the case of composition, it should be equally inoffensive in the case of realization.

NOD. Although it seems intuitive at first blush, it must be false. Given Block's insights combined with the principle of causal sharing, it leads to an absurdity: the end of all causation everywhere. Moreover, the intuitiveness of NOD depends entirely on our emphasizing the *distinctness* of higher and lower-level events. But we can argue that this emphasis should be downplayed. The connection between higher and lower-levels (e.g., realization) might not be one of identity; but, as is the case with composition, the connection does seem to bear an important resemblance to identity. Borrowing from Lewis (1991), the non-reductivist can maintain that identity-likeness is sufficient for the ontological innocence of higher level determination.

2.2 *The General Eleatic Argument*

It is a mistake to think that embracing overdetermination will be enough to save the realization hierarchy, however. Part of the motivation driving the causal exclusion argument is the intuition that causally relevant things should, at least in most circumstances, make some sort of difference to their effects; and by accepting massive overdetermination one appears to concede that higher level property instances make no difference to the world, *qua* redundant.¹⁶⁸ Even if higher level determination is ontologically innocent, we are still left to wonder why these purportedly innocent things should be posited in the first place.

¹⁶⁸ For another thing, if all higher level causation is guilty of overdetermination, then the micro-physical level still dominates the causal and the value of multi-level explanation is called into question.

2.2.1 *The Eleatic Principle*

The above considerations point to a well-known method for unmasking false innocence. To ensure that one treads cautiously when issuing new posits, austere realists commonly appeal to this intuitive connection between ontology and causal difference-making. They reason that since existent things should be relevant to the world, they should therefore make a causal difference to it. This sort of reasoning is often presented in the form of a general naturalistic principle: the *Eleatic principle* – existent things make a causal difference.

Despite the austere realist's noble intentions, the principle nevertheless faces challenges (see Colyvan 1998, for example). After all, anyone committed to the existence of abstracta is already accustomed to rejecting Eleatic intuitions. Nevertheless, the principle is not entirely useless. As Wrenn (2010) notes, the principle can be sufficiently weakened so as to remove mathematical abstracta from its scope.

(NEP) For all x , if x (physically) exists, then it makes a causal difference.

Such a weakening should restore the intuitive merit of the Eleatic principle. If the only complaint is that it rules out mathematical abstracta, then the parenthetical qualification within NEP should quell the complainer. Indeed, the intuition favoring NEP appears widespread even amongst non-reductive physicalists.

In support of this intuition is the widely held assumption that an adherence to something like NEP is part and parcel of the contemporary approach to science. When constructing theories, scientists are pressured to drop certain bits of their theory whenever

they are discovered to be irrelevant to the explanandum; and when a piece is dropped from the theory, it is also dropped from our scientific ontology.

To deny NEP leaves one to wonder why: to accept that there might be something that fails to make a difference but nevertheless physically exists, then streamlining our theories according to custom would disallow the reading of ontology off the quantifiers of our best theories; for that custom would mistakenly *entail* the non-existence of certain things that might very well exist. For those who wish to trim the excess bits from their scientific and philosophical theories without disrespecting ontology in the process, the principle appears quite helpful.

But good help comes at a price. If the characteristic posits of a theory fail to make a causal difference, then sympathy for NEP will bring that theory to bankruptcy. The theorist must then provide a reason to reject it.

2.2.2 The Higher level Does Not Make a Difference

To use the Eleatic principle to deny reification to the higher level, we must first say a bit more about the key feature of the principle: making a causal difference. To be charitable to the non-reductive physicalist, Chase Wrenn (2010) characterizes ‘making a causal difference’ in the following way: a property or relation makes a causal difference if it makes a “difference in the causal powers had or manifested by particulars.” (Wrenn 2010, p. 6) In other words, the satisfaction of a predicate must contribute in some way to what things do or can do, if that predicate is to correspond to a reified property or relation.

The general Eleatic argument can be stated in different ways. Beckermann's (1997) version is straightforward. He writes: What would be the difference between the case where an object *a* possesses both F and its realizer

...and the case in which *a* possesses both the micro-structure [$C_1, \dots, C_n; R$] and the property F? Obviously there would not be any difference. For, if *a* has the micro-structure [$C_1, \dots, C_n; R$], then, according to [realization theorists], *a* has all features which are characteristic of F, and therefore everything in the world will just go as if *a* had F. (p. 313)

Given the NEP interpretation of the Eleatic principle, these irrelevant higher level property instances, therefore do not exist. But if there are no such property instances as *being an ecosystem, being a community, etc.*, then ecosystems, communities, organisms, etc., simply do not exist. Evidently, there is no ecological hierarchy, nor any other realization hierarchy.

Although I am unconvinced by Beckermann's *ex impossibile* reasoning – you simply cannot conclude anything interesting from contrasting impossible cases¹⁶⁹ – his general point remains: the non-reductivist typically concedes that low-level micro-physical properties are metaphysically sufficient for, and hence completely determine, the actual distribution of causal powers throughout the actual world. But, if that is the case, then higher level properties do not appear to make a difference to what things do or can do: they are causally redundant. NEP then entails their non-existence. The

¹⁶⁹ Or at least it is very rare that such reasoning gives legitimate conclusions. With respect to Beckermann's case: if *a* has all features which are characteristic of F, then *a* cannot fail to have F; and therefore it is trivial everything in the world will just go as if *a* had F, because it does have it. Assuming otherwise gets you nowhere.

layered worldview appears deflated; realization looks to be a relation without relata.

Before I respond to Beckermann's argument, let us first examine a recent extension of it.

3.0 There is No Realization: The Eleatic Problem Once Again

Chase Wrenn (2010) extends this general eliminativist worry to the non-reductivist's *realization* relation itself. Since micro-physical properties are metaphysically sufficient for the actual distribution of causal powers in the actual world, the realization relation does not make a causal difference to that distribution. Therefore, given NEP, realization does not exist.

The non-existence of realization is problematic for the non-reductivist for the same reason that the elimination of higher level things is a problem: realization plays an explanatory role within their theory. Even when describing the most superficial features of realization, LePore and Loewer (1989) characterize it as a "necessary connection which is explanatory." (p. 179) Being realized by a particular molecular structure, for example, explains why a particular substance possesses the characteristic of solubility. This explanation is cast into doubt, however, if it turns out that there is no such thing as realization, since it does not make a difference.

3.1 Four Tests for Difference-Making

To buttress his eliminativist argument, Wrenn (2010) devises some tests for causal difference-making and argues that the realization relation fails all of them: the Quinean test, the modal juxtaposition test, and the refined juxtaposition test. All such arguments

can be generalized, presenting us with three separate, but equally general, tests for reification.

3.1.1 *The Quinean Test*

The Quinean test (Wrenn 2010, p. 310) interprets difference-making in terms of differences made to the distribution of energy across space-time. Since the conjunction of micro-physical facts is metaphysically sufficient for all the facts about energy distribution, Wrenn concludes that there is no room for realization. Realization does not make a difference to energy distributions; therefore, realization does not exist. Moreover, given micro-physical sufficiency, it appears that there is no room for higher level properties as well. Higher level properties do not make a difference to energy distributions since all the difference is already made by the micro-physical.

The Quinean test is unconvincing. Those who adopt Quine's picture, and also take NEP seriously, have gotten themselves into a mess. If we buy into the picture of an expansive four-dimensional manifold of distributed mass-energy, then it is hard to see how anything is making a *causal* difference to anything. The distribution is just *there*. The fundamental scientist is interested in describing general characteristics of this mass-energy distribution. They look for patterns, and then they look for ways to describe them in a way that facilitates prediction. There is nothing *causal* here. Many philosophers have observed this. It gives rise to a view I call *causal deflationism*, and it is gaining in popularity. The general idea is that *causation*, has no real "oomph"; the term is used to track patterns of *what-increases-the-probability-of-what* within this mass-energy distribution ranging across a four-dimensional manifold. To understand difference-

making in terms of “oomph”, it turns out that NEP eliminates the entire manifold from existence. To understand difference-making as the causal deflationist does has the consequence of narrowing the applicability of the Quinean test to a very specific conception of the world. It is no longer acceptable as a general test.

Moreover, it is doubtful that the non-reductivist would agree that causal difference-making requires that one makes a difference to energy distributions; any non-reductive physicalist who sees fit to analyze objects, properties, or relations as being something *other* than a distribution of energy is going to reject this presumption. For the non-reductivist, if X makes a difference to Y, X is justifiably reified whether or not Y is a distribution of energy. All that matters is that Y exists, and X makes a difference to it. The eliminativist needs a more neutral test for difference-making.

3.1.2 The Modal Juxtaposition Test

The modal juxtaposition test (Wrenn 2010 p. 308) also makes use of the metaphysical sufficiency of the micro-physical. According to Wrenn, if realization (*viz.*, L realizes H) makes a causal difference to the world, then there must be a causal difference between cases where L realizes H and where L does not realize H. However, finding a causal difference between these cases requires that both such cases are possible. But, Wrenn argues, such cases are not possible: given the aforementioned facts about the metaphysical sufficiency of the micro-physical, ‘L realizes H’, whenever true, is metaphysically necessary. Thus, the realization relation does not exist.

Wrenn argues for the strong modal profile of realization in two related ways: as a purely logical exercise and by exploiting its explanatory role. The logical argument appears in a footnote. According to Wrenn,

Every account of realization either analyses the claim that L realizes H as the claim that something is necessarily so (e.g., as the claim that, necessarily, if x has L then x has H) or as a claim that happens to be necessary... (2010, p. 314 n. 13)

In either case, Wrenn derives ‘Necessarily, L realizes H’ with a quick appeal to S5. If realization claims are analyzed as claims that happen to be necessary, then realization claims themselves must also be necessary.¹⁷⁰ Wrenn, however, is making too much of the supervenient nature of realization. Even if Wrenn is correct that realization should be analyzed using the supervenience claim ‘necessarily, if x has L then x has H’, that claim cannot be all there is to the analysis: realization claims can only turn out true if it is also *actually* the case that x has L. The analysans of ‘L realizes H’ must therefore also affirm the antecedent of Wrenn’s supervenience claim. But such an affirmation is typically contingent. The analysans, therefore, must also be contingent. S5 thus has no bearing on the present circumstances.

Wrenn makes a similar point regarding realization’s explanatory role. On all accounts of realization, he claims, for L to realize H is for L to explain H in the right sort of way. On “all these accounts, though, it is not contingent whether L explains H; if L explains H in the requisite way, it is necessary that L explains H in that way.” (2010, p.

¹⁷⁰ S5, recall, tells us that $\Box(R \text{ iff } P)$ and $\Box P$, together, entails $\Box R$.

314) Wrenn infers from this that true realization claims are metaphysically necessary. But again this cannot be correct: L cannot explain H in world in which L does not exist; and so even if L actually explains H, it cannot *necessarily* be what explains H. H might be explained by something else, L*, in worlds where L does not exist. Indeed, that some higher level properties are multiply realizable is a core feature of non-reductive physicalism. The modal juxtaposition test is therefore no more convincing than its Quinean cousin, at least in its current formulation.

3.1.3 *The Contrastivist Test*

The strongest version of Wrenn's eliminativist argument (Wrenn 2010, p.315 *fn.* 14), call it the *contrastivist test*, again begins with a juxtaposition:

(CT) X makes a causal difference to Y (atop the causal contribution made by Z) if and only if there is a causal difference in Y between cases where X and not-X (holding Z fixed in both cases).¹⁷¹

What this tells us is that L's realizing H makes a difference to some particular (atop the contribution made by the micro-physical instantiation of L) only if there is a causal difference in that particular between cases where L obtains and H is realized and cases where L obtains, but H is not realized. But, given that L is supposed to be metaphysically sufficient for H's being realized, the second case is impossible. Thus, realization makes no causal difference: there is no causal difference between L-obtaining cases in which 'L realizes H' holds and L-obtaining cases in which it does not, since the latter cases are impossible.

¹⁷¹ Wrenn ultimately wants to convince us that there is no contribution made by realization that is not already made by some micro-physical property instantiation on its own.

But now the refined test is no more convincing than Beckermann's original *ex impossibile* argument. The test might be convincing if realization failed on the grounds that both cases were possible, and still no difference could be found by their juxtaposition. But that's not the result here. Here the failure hinges only on the fact that the juxtaposition itself is impossible.

The non-reductivist can simply argue, however, that if juxtaposition is impossible, then the juxtaposition test becomes useless. Eliminativists are pushing it beyond its limit. Just because there is no juxtaposition does not entail that there is no causal difference. It seems more appropriate to say that if we *cannot* compare cases, then we simply cannot use juxtaposition to test the causal contribution of realization atop the contribution made by L's being instantiated.

This problem echoes one of the more serious limitations of interventionist accounts of causation (Woodward 2003): just because it is metaphysically impossible to intervene on, or "wiggle", an element in a causal system does not immediately entail that element's inefficacy, it entails that the metaphysical possibility of an intervention on X is not necessary for X to be a cause.¹⁷²

Consider a magic pipe. Suppose you have a tobacco-pipe that is metaphysically impossible to alter in any sort of way.¹⁷³ This means, for any particular feature of the

¹⁷² Woodward instead relies on the logical possibility of interventions. In cases where an intervention is not physically possible, we can instead look to an appropriate theory to justify counterfactual claims involving purely hypothetical interventions (Woodward 2003 p. 128-30).

¹⁷³ I do not mean to make any substantive claim about the world by talking about magic pipes. Rather, I'm questioning the coherence of the testing method, and so hypothetical examples should carry weight, however simplified. If hypothetical examples tell us that the test can fail in circumstances where a juxtaposition or intervention is impossible, then we absolutely cannot use the test to conclude anything substantive when such circumstances in fact obtain.

pipe, it is impossible to “wiggle” that feature independently of the others in order to test its relevance. Is the shape of the handle relevant to the causing of its owner’s lung-cancer? Perhaps not; but even so, the causal vacuity of the handle’s shape does not seem to have anything to do with the fact that the pipe is impossible to alter. That one cannot compare the original case to a juxtaposed case in which the pipe has been altered has nothing to do with the causal relevance, or irrelevance, of any given feature of the pipe. As with all cancer-causing pipes, some feature of the pipe is partially responsible for the cancer it helps to bring about.

This is not to say that intervention tests are bad tests, only that they have their limit. The same can be said with the contrastivist test for difference-making. If the results of the test are not in realization’s favor, then so be it; but if the nature of the thing being tested makes it impossible to even make a juxtaposition or intervention, then the test simply cannot be performed. No results whatsoever can be obtained.

We can see this point again from a more abstract perspective. Suppose you have a causal chain, *A causes B causes C*. If we assume deterministic causation, then A is metaphysically sufficient for B, and B is sufficient for C. If we apply the contrastivist test, however, it will follow that B makes no difference to C atop the difference already made by A: *qua* deterministic, there can be no case where A and not-B; and thus there is no difference between an A and B case, and an A and not-B case. Given NEP, B, therefore does not exist and as such is not part of any causal chains. But that is clearly false: B does exist as part of a causal chain, by stipulation. And besides, it also seems clear that B, being part of the causal chain, is ipso fact *part* of the difference made by its predecessor, A. The contrastivist test, in its current formulation, thus fails to recognize

the difference made by intervening nodes in a deterministic causal chain; and therefore, the test simply fails.

The third, and final problem with the test is that it illegitimately entails that there are no metaphysically necessary events. If an event *e* (e.g., some great explosion or expansion near the beginning of time) is metaphysically necessary, then you cannot compare a case in which *e* obtains with a case where it does not, since there are no cases of the latter sort. And so, given the contrastivist test and NEP, no such *e* makes a difference and therefore metaphysically necessary events do not exist.

This line of reasoning is quite suspicious. In general, the discovery that *X* makes a difference should not permit us to infer that *X* is not necessary. The contrastivist test simply does not seem like an appropriate way to settle disputes about the modal status of events.

The main reason the contrastivist test is unconvincing is not because it illegitimately settles modal disputes, however; it is because, on the face of it, big bangs, necessary or not, make a big difference.¹⁷⁴ Ultimately, any test that tells us that modal status matters for its difference-maker status is not a test anyone should accept.

3.1.4 The Refined Contrastivist Test

Wrenn's contrastivist argument is therefore unsound. Because the modal status of an event is irrelevant to its ability to make a difference, we must reject the key premise: *X* makes a causal difference (to *Y*) at *t* if and only if there is a difference made by *Z* if and only if there is a

¹⁷⁴ As an anonymous referee has pointed out, one might add deities to the list: God, should He exist, is generally taken to be both a difference-maker and a necessary being.

difference (in Y) between cases where X and Z hold and cases where not-X and Z hold. To repair the test, that biconditional premise needs to be stated as part of a greater biconditional: it is to be affirmed if and only if a difference between such cases is possible. Only with such a modification does the test gain credence. Modifying the test in such a way, of course, makes it impossible for Wrenn to draw the conclusion he wants.

Somewhat ironically, with this more plausible test in hand, non-reductive physicalists find themselves in a position to respond to the general eliminativist argument. In what follows, I show how they can use the repaired contrastivist test to restore the realization hierarchy, by arguing that the realization relation, and the higher levels realizations do in fact make a difference to the world, at least within the framework provided by non-reductive physicalism.

4.0 Resurrecting the Realization Hierarchy

To defend the realization hierarchy from Eleatic concerns, I argue for two claims. First, I argue that the realization relation should be considered to be reified within the framework of science, since it makes a difference. Second, I argue that the higher levels of the realization hierarchy are similarly reified, since *they* make a difference as well.

4.1 The Reification of Realization: It Makes a Difference

Even if each of Wrenn's specific Eleatic tests can be challenged, non-reductivists still need to respond to the general eliminativist argument. They still must say how it is that the properties and relation targeted for elimination do in fact make a difference to the causal powers of particulars, even though micro-physical property instantiations are

metaphysically sufficient for the total distribution of causal powers throughout the universe.

One way to do this is to show that micro-physical property instantiations are not, *by themselves*, metaphysically sufficient to account for the causal powers of particular higher level property instances. One might argue that while micro-physical property instantiations are metaphysically sufficient for the distribution of causal powers found at the higher level, this fact holds only *because* of these important higher level properties and connecting relations.

4.1.1 The (Causal) Disconnection Between Levels

Eliminativists think that the set of microphysical events is sufficient for all other events; and non-reductivists typically agree. Eliminativists also tends to think that micro-physical facts are also sufficient to explain all other facts; but here the non-reductivist disagrees. For the non-reductivist, the modal properties of micro-physical property instances explain little (LePore and Loewer 1989, p. 178-9). Indeed, the logical relationship of metaphysical sufficiency that holds between micro-physical L-instances and higher level H-instances – e.g. $\text{Nec}(L \rightarrow H)$ – is itself something that needs to be accounted for. Such modal facts are supposed to be accounted for by the fact that L realizes H; they are too important to be left to brute supervenience.

Given the analytical impotence of mysterious modal facts, the non-reductivist can argue that micro-physical property instantiations alone cannot account for the causal powers of higher level property instances. While satisfying the predicate *instantiates-L* (where ‘L’ denotes some micro-physical property) can account for a particular L-

instance, x, having the powers it does, higher level property instances (H-instances) do not happen to instantiate L. If the H-instance, h, were identical to the L-instance, then the causal power h obtains by instantiating L would be accountable simply by the fact that h instantiates L, and the strong modal connection between the L-instance and the H-instance would be likewise accountable by virtue of their being identical.

However, if the H-instance is *not* identical to the L-instance, the L-instance's instantiation of L on its own does not account for the H-instance's having the causal powers it does; nor does it account for the aforementioned strong modal connection: that the L-instance is metaphysically sufficient for the H-instance. The non-reductivist will ardently maintain that there must be *some* explanatory relationship between the L-instance and the H-instance if the L-instance is to account for the powers possessed by the H-instance. Presumably the L-instance gets its powers by instantiating L, but how then does the corresponding (but non-identical) H-instance get its powers? Assuming that the instances are not identical, the instantiation of L on its own does not account for much at all.

The layered worldview under discussion, recall, is one of *dimensioned* realization, which is irreflexive: nothing realizes itself. The H-instance, h, then cannot be identical to the L-instance that realizes it. Indeed, if they were identical it would be hard to see how non-reductive physicalism can genuinely be considered non-reductive.

4.1.2 *Co-opting the Contrastivist Test*

Given that the mere fact that L is instantiated fails to account for the causal powers possessed by higher level property instances, the non-reductivist is presented with a way to show that realization makes a causal difference.

Consider the claim ‘H is realized by L’, where L is an instance of a micro-physical property and H is an instance of a higher level property. As noted earlier, such a claim must be considered contingent in order to accommodate the non-reductivist concept of multiple-realizability; there are some worlds where H instantiates ‘realized by L’ and there are some worlds where it does not. Moreover, in any such world where H exists but is not realized by L, H must nevertheless be realized by something else, some L*.

Focusing on these two worlds, we now have two cases to feed into the contrastivist test. Is there a difference between cases where H instantiates *realized-by-L* and cases where H instantiates *realized-by-L**? Although the repaired test should deliver an answer (since the difference between cases is possible), the answer it delivers is going to depend on the nuances of the particular form of non-reductive physicalism one favors. Specifically, the answer will depend on one’s reaction to the following principle:

(Power-essentialism) For any two property instances, h and h*, and any two micro-physical properties p and p* if h and h* are instances of the same property and p and p* are instances of different properties, and h is realized by p and h* is realized by p*, then h and h* have all the same causal powers.

Otherwise put: higher level property instances have the same causal powers no matter how they happen to be realized. Non-reductivists do not endorse this principle. Some powers, it is standardly thought, are conditional: their acquisition depends on how the higher level instance is realized.

Given the denial of power-essentialism, the non-reductivist is free to assert both (1) and (2) below:

(1) By instantiating *realized-by-L*, the H-instance obtains a certain set of causal powers.

(2) By instantiating *realized-by-L**, an H-instance obtains a different set of causal powers.

By affirming (1) and (2) and Wrenn's contrastivist test, it follows that a difference has been made to the causal powers of the H-instance. Specifically, we see that satisfying the predicate *realized-by-L** makes a causal difference to the H-instance atop the contribution made by the H-instance itself.

Does this mean that *realization* makes a difference? It does. Since this complex predicate can be broken down into two components – the realization relation and the L* instantiation – given that the L*-instance cannot account for the causal powers of the H-instance on its own (since it does not instantiate L), it will then follow that the *realization* relation does in fact make a causal difference. Although I have not said exactly what that difference is, the argument strongly suggests that a difference is made by realization.

On the other hand, if one accepts this principle of power-essentialism, this response option quickly evaporates. Given power-essentialism, juxtaposing a case where

H is realized by L with a case where H is realized by L* yields a different answer: it follows that no difference is made to the partitioning of causal powers to the H instance by virtue of being realized by L*. Luckily for the non reductivist, there is no good reason to accept power-essentialism, although some philosophers might find it tempting.

4.2 The Reification of the Higher level: It Makes a Difference

In the previous section, we used the contrastivist test to describe one sort of difference that realization makes to the causal powers of particulars. As the test was originally designed to eliminate realization, the result is rather ironic. In this section, I will press forward with the irony, employing the same strategy with respect to higher level property instances. I argue that higher level property instantiations do make a difference, pace Beckermann's argument from section two.

4.2.1 The Contrastivist Strategy Revisted

Rather than juxtaposing the micro-physical property instantiations against the realization relation, here we must instead juxtapose the realizing micro-physical property instantiation against the higher level property instantiation that is realized. Recall the contrastivist aspect of Wrenn's argument. Assuming that both claims *X and Z* and *not-X and Z* are possible,

(CT) X makes a difference atop the difference made by Z if and only if there is a difference between cases where X and Z and cases where not-X and Z.

To use this principle to positively affirm that H makes a difference atop the difference made by L seems to require contrasting a case where both H and L obtain with a case

where L obtains but H does not. However, given that L is metaphysically sufficient for H, such a contrast turns out to be impossible; indeed, this is the very feature that Beckermann and Wrenn attempted to exploit when arguing against realization physicalism. As we learned from our study of their arguments, however, impossible contrasts do not imply Eleatic inefficacy; they only imply that Eleatic efficacy cannot be determined from that particular test.

Still, one might be able to find a way to use the contrastivist test to *indirectly* demonstrate that H does make a difference within non-reductive physicalism. Given that H and L are in competition for the status of difference-maker, if we can use the contrastivist test to show that L does *not* make a difference atop the difference made by H, it will then follow that H does make a difference; for if it did not, then it would follow that the micro-physical property does not make any sort of difference at all. Anything that fails to make a difference atop something that makes *no* difference must also be something that makes no difference. But since it is unquestionable that micro-physical properties do make some sort of difference (given the non-reductivist's commitment to the causal closure of the micro-physical), it then follows that higher level properties must also make some sort of difference within non-reductive physicalism. Saying exactly what that difference is becomes the only remaining challenge.

This is only the sketch of the argument. We have yet to perform the test.

4.2.2 *Implementing the Contrastivist Strategy*

To set things up, let us consider the following scenario:




Figure 21: Hs realized by Ls

Let us imagine that the H's are not reducible to the L's, but rather, are realized by them: let ${}^R H_1$ denote the set of possible lower-level realizers of H_1 and ${}^R H_2$ denote the set of possible lower-level realizers of H_2 . Finally, let us imagine that any world where an element of ${}^R H_1$ obtains is a world where an element of ${}^R H_2$ obtains.

Given this scenario, let us run the contrastivist test. Since any possible realizer of H_1 is sufficient to bring about a realizer of H_2 , it follows that if it were the case that H_1 obtains but L_1 does not, H_2 would nevertheless obtain. In other words, H_2 is insensitive to changes in L_1 when we hold fixed the contrast class H_1 (cf. Van Frassen 1982; Schaffer 2005).

Since H_2 is insensitive to changes in L_1 , holding fixed H_1 , we have thus failed to satisfy the right-hand side of the contrastivist test. By holding H_1 fixed in both cases, no difference is made to H_2 by altering L_1 . Therefore, it follows that L_1 makes no difference to H_2 atop the difference made by H_1 . Maintaining the assumption that L_1 makes *some* difference to H_2 on account of its bringing about H_2 's actual realizer, L_2 , it therefore follows that H_1 also makes *some* difference. Again, what exactly that difference is remains a mystery.

Note that this result is consistent with what we see in actual scientific practice. Particle physicists cannot predict the precise behavior of individual particles. They are

just too erratic. Instead, statistical mechanics is called in to make predictions from an initial set of possible particle configurations. Those predictions nevertheless turn out to be quite accurate on the macro-level. This accuracy suggests that the particular configuration in which those particles actually find themselves is quite irrelevant to the macro-level behavior of the system. It makes no difference to the prediction.

Now that it has been argued (from the eliminativist-endorsed contrastivist test) that both higher level property instances and the realization relation make some sort of causal difference within non-reductive physicalism, we are warranted to think that there is a sort of non-reductive physicalism is compatible with both the NEP interpretation of the Eleatic principle and the SEP interpretation. But there are other versions of the realization hierarchy.

4.2.3 The Problem for Property Realization

Some non-reductive physicalists will be wary to endorse the above solution. They will be bothered by the presumption that the relata of realization are property instances rather than the properties themselves. Kim and Shoemaker, for example, maintain that realization holds between properties.¹⁷⁵

How does the eliminativist's argument hold up against this alternate conception of realization's relata? If we take the view that higher level properties and their realizers are instantiated by the same particulars, then realization will end up succumbing to the eliminativist argument: that one property realizes another does not seem to have any bearing on the powers of particulars. If a particular already instantiates the lower-level

¹⁷⁵ In fact, Shoemaker (2007) holds that realization holds between property instances *and* it holds between properties.

realizer, it automatically acquires all of the powers bestowed by that realizer. There is nothing to be gained by also instantiating a higher level property or by introducing a realization relation connecting the higher level property to the lower. Neither realization nor higher level property instantiations make a difference, and are therefore eliminated by NEP.

If, on the other hand, the non-reductive physicalist maintains that nothing instantiates both a higher level property and its realizer, we can run the same argument as we did for instances. Since higher level particulars do not instantiate the lower-level realizer property, the powers they acquire are gained solely by their instantiation of higher level properties (presuming that powers are acquired by particulars via the instantiation relation). With such restrictions in place, higher level properties do make a difference to the powers of particulars that instantiate them.

The realization relation is a different story. For those who think that realization relates properties as opposed to property instances, one is left to wonder what realization adds to the causal story. Since higher level particulars gain all the powers they do by instantiating a property, the fact that that property is realized by another contributes nothing. Since particulars fall outside the domain of the realization relation, there are no grounds to claim that realization contributes anything.

In this case, however, defending the existence of realization does not require arguing that it is a relation that makes a difference. One can instead argue that it is, strictly speaking, a mathematical relation. Take Shoemaker's (2007) "subset" view of realization as an example. For Shoemaker, X is realized by Y iff the set of X's forward

looking causal features is a subset of Y's forward looking causal features, and Y's backward looking causal features are a subset of the backward looking causal features of X. The subset relation, however, is a mathematical relation. Having been analyzed as a mathematical relation holding between sets of powers, the realization relation fails to satisfy the "physical" antecedent of NEP, and so it is exempt from having to make a difference to the world, passing through on its set-theoretic credentials. No conflict between subset-realization and NEP is generated. The non-reductivist who *disagrees* with Shoemaker about the mathematical nature of the realization (subset) relation, however, will have to fight back.

4.3 *Eliminating the Eleatic Principle*

Those non-reductivists who rail against a mathematical analysis of realization, but who also fail to account for the difference it makes will have to reject NEP. A passive dismissal does them no favors, however. It would be in their interest to mount a case *against* the principle. If it can be shown that NEP fails to hold as a metaphysical principle, then the NEP connection between existence and difference-making can be severed; and the non-reductive physicalist has nothing to fear from the Eleatic eliminativist.

Metaphysical principles are typically those which bear a strong modality: metaphysical necessity. NEP is no exception; it tells us that difference-making is metaphysically necessary for existence. The entire point of the principle, after all, is to draw a metaphysical connection between causality and existence. If there are beings in possible worlds where difference-maker status is *not* required for non-mathematical

existence, then that connection is broken. Even worse, if that connection breaks down in a model of world like ours, then the “intuition” that ours is one of those lucky worlds where NEP holds becomes nothing more than bad faith.

The following is a model of a world like ours. Although it makes use of the notion of a conditional power we introduced earlier, it is nothing special: it merely represents a familiar sort of overdetermination. The philosophical importance of overdetermination cases, of course, is well-known. Traditionally offered to quiet the claim that difference-making is necessary for causation, such cases are what keep counterfactual accounts of causation shrouded in skepticism.¹⁷⁶ But as we will see, those problems extend to the eliminativist interpretation of the Eleatic principle, NEP, as well.

Consider a causal model consisting of four events D , C_1 , C_2 and E , held together by the following chain of relationships: (a) D , given background conditions α , causes both C_1 and C_2 ; (b) C_1 is a completely sufficient cause of E ; and (c) C_2 is too. (b) and (c), of course, tell us that E is overdetermined by C_1 and C_2 .

This scenario poses a problem for NEP. If making a causal difference were metaphysically necessary for existence, then it would follow that C_1 and C_2 do not exist: as we noted earlier, if an effect is genuinely overdetermined then each of its overdeterminers is causally redundant. Individually, they make no difference to the

¹⁷⁶ The literature on this issue is enormous. Shaffer (2007) is as good a place to start as any.

effect.¹⁷⁷ But their non-existence contradicts the parameters of the model: recall that each of these four things had been stipulated to exist at the outset.

Taking C_1 and C_2 to compose a single event C_3 is of little help to NEP. By invoking the notion of a conditional power, we can extend the model to entail that C_3 does not make a difference. If, for example, we suppose that D has the power to bring about both C_1 and C_2 *only* under the α conditions, we can then add the stipulation that D has the power to bring about some C^* given any other background condition, β . We can then further stipulate that this C^* is itself metaphysically sufficient for E and derive via the contrastivist test that C_3 does not make a difference: there is no difference between cases where D and C_3 hold and cases where D and not- C_3 since not- C_3 and D together imply C^* ,¹⁷⁸ which is also guaranteed to bring about E . Represented diagrammatically:



Figure 22: Two Causal Situations

Making a causal difference, therefore, cannot be a metaphysical requirement for reification. If it were, then the above causal model would involve overdeterminers that both do and do not exist, which is nonsense.

The eliminativist's interpretation of the Eleatic principle (NEP) therefore does not accurately portray the metaphysical connection between difference making and ontology,

¹⁷⁷ As can be confirmed by the revised contrastivist test. It is still possible, however, that C_1 and C_2 are nevertheless causes of the effect despite not making any causal difference.

¹⁷⁸ Not- C_3 and D implies C^* since not- C_3 and D implies that the α -condition do not obtain, which implies that the β -conditions do obtain, which, when combined with D , implies that C^* obtains.

suggesting an agnostic approach instead: do not eliminate the irrelevant,¹⁷⁹ but do not reify it either.

5.0 Implications for the Causal Exclusion Argument

Before closing, it is worth mentioning the implications that the rejection of NEP has for the causal exclusion argument: none whatsoever. The fact that some things might fail to make a difference to the world does not make systematic overdetermination suddenly more plausible, nor does it make it any more likely that the higher level brings about phenomena that the lower-level does not. Epiphenomenalism still looms. The rejection of NEP does, however, prevent the eliminativist from using Eleatic arguments as a way of turning that epiphenomenalism into stark eliminativism.

Our contrastivist test, on the other hand, does seem to have implications for the exclusion argument. The causal sharing principle (i.e., CSHARE) tells us that anything a higher level instance brings about is also brought about by lower-level instances. As we have seen here, there could be cases where the lower-level instance makes no difference to the effect atop the difference made by the higher level instance. One might take this to suggest that the causal work is being done at the higher rather than the lower level; but really, we cannot say for certain: the contrastivist test cannot confirm that the higher level makes a difference atop the difference made by the lower. If causal work sometimes turns out to be higher level work, then we would expect that the higher level would indeed make a difference atop the difference made by the lower. But there is no way to confirm

¹⁷⁹ That is, unless you have an independent (non-Eleatic) reason for doing so.

that with our analysis: we would need a test for contrastive difference-making that does not require us to reason *ex impossibile*.

5.1 Weighing the Options

We might still try an alternative way to reason through the issue. We can describe the space of options regarding inter-level difference-making in the following set of diagrams (a) through (h). Where A designates the higher level and B designates the lower, and the diagonal shading designates that a difference was made:¹⁸⁰

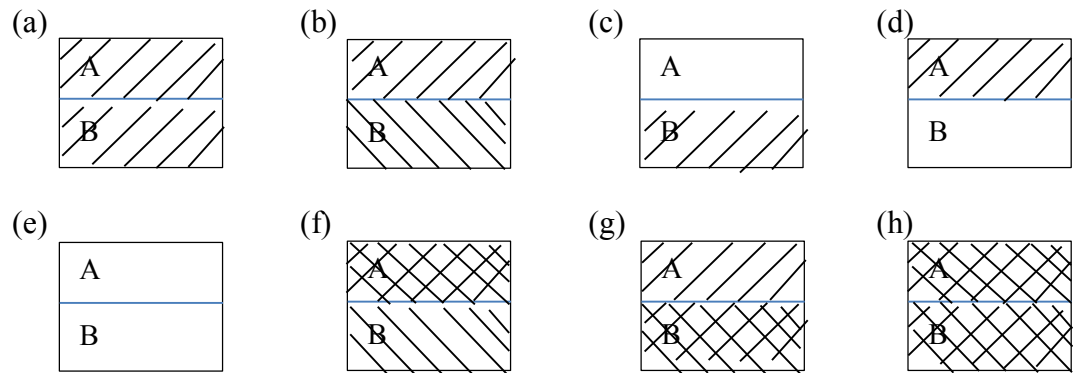


Figure 23: Eight Difference-making Options

(a) tells us that the difference made by A is the same difference as that made by B; (b) tells us that A and B both make a difference, but they are entirely separate; (c) tells us that all the difference is made by B and A makes no difference; (d) tells us that all the difference is made by A and that B makes no difference; (e) tells us that no difference is made whatsoever; (f) tells us that A and B both made a difference, but the difference B makes is part of the difference made by A; (g) tells us that A and B both make a

¹⁸⁰ We could also represent the space symbolically: (a) $0 \neq A=B$; (b) $0 \neq A \neq B \neq 0$; (c) $0 = A \neq B \neq 0$; (d) $0 \neq A \neq B = 0$; (e) $0 = A=B$; (f) $A = B + N$, where $N > 0$; (g) $B = A + N$, where $N > 0$; (h) $0 \neq A + B = B + A$. I see no advantage in doing so, however.

difference, but the difference A makes is part of the difference made by B; (h) A and B both make a difference, and they both contribute to the difference each makes.

Many of these options seem to suggest trouble for CSHARE. If we accept (b), (d), or (f), then the higher level makes a difference that the lower level does not. Since this suggests that the higher level is contributing something that the lower-level is not, it appears that CSHARE can be rejected.

Most of the above options can be eliminated immediately, however. In those special cases in which the lower level does not make a difference atop the difference made by the higher, we can eliminate (b), (c) and (g). In all these cases we have B making a difference not already made by A.

We can also eliminate (d) and (e). In those special cases where the lower-level makes no more difference to the effect than the higher level, the contrastivist test should nevertheless confirm that the lower-level instances make a difference atop the difference made by, say, some distant planet. Even in those special cases, then, the lower-level instances still make some sort of causal difference. (d) and (e) must therefore be eliminated.

This leaves us with a trilemma: either the difference made by the two levels is the same as in (a), the difference made by B is part of the difference made by A as in (f), or the difference made by each is partially made by each other as in (h).

It is hard to get a grasp of the latter horn, I admit. But there is a reason for that: the trilemma is a false one. (h) and (a) amount to the same thing. The reason (h) does not seem to be ruled out by the special case under consideration is because, even though B is

making its own difference, it is still not making a difference atop the difference made by A, since the difference made by A *includes* the difference made by B. But the same can be said for A, and this is why the dilemma is a false one. It follows from (h) that the difference made by A is the difference made by both A and B and the difference made by B is the difference made by both A and B. The difference made, therefore, is the same across levels: (h) collapses into (a).

And so we actually face a dilemma: either (a) the difference made is the same, or (f) the lower-level difference is part of the higher-difference. If we accept the causal closure of the micro-physical, as nearly every non-reductive physicalist does, then (f) can also be eliminated. (f) gives us a case where the higher level brings new differences into the picture, differences not found at the micro level. If one sees fit to reject (f), then, one must also reject the causal closure of the micro-level; and with that, CSHARE can be rejected.

Note also that if (f) is true, then it becomes necessary to reject the principle of no-overdetermination (i.e., NOD); (f) is a picture of a case where the lower-level makes enough of a causal difference to bring about the effect (since it brings about the realizer of the effect), and the higher level nevertheless makes an extra difference on top of that. The effect is not then just overdetermined; even worse, it is over-overdetermined. We have no choice but to reject NOD.

This leaves option (a). If true, (a) also seems to suggest that overdetermination is unavoidable, again mandating a rejection of NOD.

5.2 Interpreting the Result

There are only two possible ways to account for the possibility of sameness of the differences made by the higher and lower-levels. First, one might take a metaphysically reductionistic, identity-theoretic stance towards difference-making, independent of any claims reaching towards inter-theoretic reduction that motivated the variety of reductionism favored in Putnam's and Nagel's day – namely, the reduction of scientific predicates or posited properties.

Such an account calls for a change in the common pictorial representation of inter-level causation; where the vertical bar | denote realization:



Figure 24: Inter-level Causation

This picture suggests that the entire difference made is made by the hierarchical structure that includes both the H₁ instance and the L₂ instance. One might then interpret the situation as depicting a different sort of scientific reduction. If non-reductive physicalism holds for scientific predicates, perhaps those involved in the reductionist project should take to the business of looking for ways to reduce the *differences* made by the predicate-instances to the difference made at the lower level: reduce the differences made rather than the properties. Although I acknowledge the (epistemic) possibility of such a reductionist position, I have no reason to favor it. I simply have no reason to reject it.

The second option a quasi-reductionistic stance towards difference making, whereby the difference-making relation does not reduce, but *nearly*-reduces in the same sort of way we said that higher level property instances are *nearly*-identical to the instances that realize them. That is, we might say that the difference made by the higher level is *realized* by the difference made by the lower. Again, this calls for a difference in the standard picture:



Figure 25: Realizing Causation

Here we see that the higher level difference is, in fact, realized by the lower-level. Again, I have no argument for this position.¹⁸¹ Either option is sufficient to mandate a rejection of the no-overdetermination principle (i.e., NOD). Both epiphenomenalism and Eleatic eliminativism can be avoided.

This defense need not reset on a defense of (a), however: if either (a) or (f) holds, then we have what we need to reject the causal exclusion argument. If (a) holds, then NOD must be rejected: you have a case where both the lower and higher levels make a sufficient difference to the effect; if (f) holds, then both NOD and CSHARE must be rejected: not only do there appear to be novel causal powers at higher levels within (f),

¹⁸¹ I am tempted to appeal to relational properties here. If relational properties, in general, can be realized at the lower-level, then H₁'s instantiating *makes-a-difference-to-H₂* might be said to be realized by L₁'s instantiating *makes-a-difference-to-H₂*. There are too many issues at work here to be able to say anything with any warrant. Besides, pursuing such a line is tangential to my main goal: defending against hierarchical deflationism.

mandating a rejection of CSHARE, there is also more causal power in the picture than is needed to bring about the effect, requiring us to also abandon NOD. Either way, the contrastivist test, therefore, requires that NOD must be rejected.

Higher level things, therefore, make a difference. The causal exclusion argument is unsound. Block's drainage argument tells us that we have to reject at least one of the premises, and the contrastivist test then confirms that it is NOD that it has to go, supporting our original intuition.

6.0 Conclusion

I should note before closing that the eliminativist positions discussed here were never offered as arguments against the appeal to hierarchies in scientific endeavors. I have simply redeployed them, exploring the concern that our talk of levels in everyday, analytical, and scientific discourse is just that: talk.

As noted in the outset, eliminativism about higher level things or the ordering relation that holds between them entails levels-deflationism with respect to that hierarchy. The causal exclusion argument purports to show that higher level property instances are inefficacious. *Qua* inefficacious, some austere realists then feel compelled to eliminate higher level properties by invoking the NEP interpretation of the Eleatic principle: if something exists, it must make a causal difference.

There are two ways non-reductive physicalists can respond: they can reject the causal exclusion argument or they can reject the Eleatic principle. If one does the former, one must give reason to think that higher level things make a difference. If one does the

latter, one must give an example of something that does not make a difference but nevertheless exists.

Here I have pursued both avenues. In response to the causal exclusion argument, I considered the rejection of both premises. Ultimately, they amounted to the same thing: NOD must be rejected. In the end, it appears that the sort of overdetermination involved in the causal exclusion argument is ingenuine. Either way, epiphenomenalism is averted.

But the eliminativist has other resources at his disposal. Eliminativists commonly dismiss non-reductive physicalism more directly, on Eleatic grounds. The Eleatic principle, they say, ought to be part of our approach to science. The construction of theories is an ongoing process involving frequent refinements. Part of what drives such changes is an implicit adherence to the Eleatic principle: drop whatever is causally irrelevant. It simply does not exist. Because non-reductive physicalism is committed to higher level properties and a realization relation that do not make a difference, it appears that it is incompatible with the Eleatic principle. Non-reductive physicalism, they conclude, provides a poor metaphysic for science.

In the second half of the chapter, however, I went on to describe a variety of ways eliminativists attempt to show that non-reductivist elements do not make a causal difference: the Quinean test, the juxtaposition test, the contrastivist test, and the refined contrastivist test. It was shown that the only acceptable test was the refined test, and that the requisite refinement effectively removed the bite from the eliminativist's contrastive argument; the refined test ousts neither the realization relation nor higher level property

instantiations from the non-reductivist's ontology. Ironically, the same refinement allows the non-reductivist to respond to the general eliminativist argument.

I then went on to show that there is no incompatibility between non-reductivism and the Eleatic principle. Higher level properties and the inter-level realization relation do in fact make a difference, by the eliminativist's own standard. When an effect is tolerant to changes at the lower-level, holding fixed the higher level, the contrastivist test implies that higher level properties make a difference. What that difference is, exactly, remains a mystery.

But there is a deeper lesson to be learned from our investigation. Applying the contrastivist test to a simple causal model shows that making a difference is not necessary for physical existence. That something makes a difference might be sufficient for it to exist, but it is not necessary.

Conclusion and Directions for Future Research

Progress is made in the metaphysics of science whenever it can be shown that the assumptions of an explanatory framework have consequences about the nature, contents, and structure of the world. In the preceding examination of the use of hierarchical structures within ecology, we have done just that.

The thesis defended consists in three claims: the claim that the world is hierarchically structured (i.e., Realism); the claim that the world is hierarchically structured as ecologists claim, but it is not the case that everything in the hierarchy is, in fact, contained within the bottommost level (i.e., Holism); and the claim that Metaphysical-MLE – the idea that the hierarchical structure of the world is what justifies the practice of describing, explaining, and analyzing things using hierarchical terms – appears false with respect to some hierarchies.

I began by providing a clear and concise statement that describes the ecological hierarchy, as it is conceived by Allen and Starr (1982), and others (e.g., Webster 1979, and Salthe 1985). The ecological hierarchy defended here is a sort of *compositional* hierarchy. I then defended this description, ECOHIERARCHY, from some recent objections.

In the second chapter I reviewed the most prominent general theories of hierarchies *qua* hierarchies: theories of hierarchies that are detached (or abstracted from) any concrete example of a hierarchy; theories that capture the notion of ‘hierarchy’, as it is employed within science, *however* the term is employed. After arguing that levels are

parts of hierarchies, it was then argued that the theoretical considerations of parsimony and generality favor an unrestricted mereological approach over a set-based approach for understanding what hierarchies *are*. To be “at” a level, it was argued, is to be a part of that level, since such a conception provides a general foundation for the construction of *any* hierarchy. What is it for one level to be higher than another? Two answers were considered: the partial order view and the *strict* partial order view. We came to favor the strict view.

Taking these considerations for granted, in the third Chapter a robust general theory of hierarchies *qua* hierarchies was developed: Core Hierarchy Theory. Its formalization has interesting consequences: it vindicates the conception of ECOHIERARCHY from the first chapter by giving us a way to argue that all unrestricted compositional hierarchies are nested hierarchies. That is, Allen and Starr’s (as with Webster and Simon’s) idea that compositional hierarchies are *upwards nested* hierarchies turn out to be a theorem of the formalized Core Hierarchy Theory when the “is at” relation is left unrestricted.

These results were then applied to well-known issues in the metaphysics of ecology: debates between Reductionism and Holists, and issues surrounding mechanistic vs. Humean approaches to scientific explanation. Having shown that Reductionism is inconsistent with Core Hierarchy Theory when the “is at” relation is left unrestricted, I argue that Reductionism should be rejected: unless it can be shown that the most basic unrestricted sort of hierarchical structure fails to capture some important aspect of the debate, Holism is to be preferred.

Causal explanation, mechanistically-driven or otherwise, is commonly associated with compositional hierarchies. Since Craver's *levels-of-mechanisms* approach to multi-level explanation utilizes restrictions on the "is at" relation, it does not follow that mechanistic hierarchies are upwards nested hierarchies. Unrestricted compositional hierarchies, such as Allen and Starr's, Salthe's, and ECOHIERARCHY, on the other hand, are upwards nested hierarchies. Metaphysical-MLE (the metaphysical justification for multi-level approaches to explanation) is problematic on such a conception: since, in an upward nested hierarchy, everything at any level is found at all higher levels, it follows that the top-most level (e.g., the ecosystem level) of such a hierarchy contains all the resources needed to provide a complete explanation of any (e.g., ecosystemic) phenomena of interest. Core Hierarchy Theory then suggests that, with respect to certain hierarchies, Metaphysical-MLE is false, despite that those hierarchies are "real".

There are those who challenge the reality of hierarchies. They think that there is an internal conflict within hierarchical metaphysics, arguing that since higher level things, and inter-level realization relations, do not make a difference, they do not exist. In response to these worries, I indirectly show, from our starting assumptions, along with a well-known scientific heuristic, that these higher level things and inter-level realization relations do indeed make a difference. So, we should think that they exist after all.

Directions for Future Research:

There are five general areas of inquiry for which this dissertation might serve as a starting-point: philosophy of ecology, philosophy of mind, metaphysics, philosophy of mathematics, and formal ontology and bioinformatics.

Within the discussion of the ecological hierarchy, there were issues surrounding specific levels of the ecological hierarchy: challenges to the *community* concept, for example. Similar challenges have been raised for the ecosystem level as well. It would be valuable to contrast ECOHIERARCHY with different ecological level-schemes, to see whether or not the way we structure ecology has any influence on progress within the discipline. For example, a more detailed examination of the use of compositional hierarchies vis-à-vis the scalar alternative.

An examination into to the application of diversity metrics across hierarchical levels within ecology, I think, provides another avenue for future research. Along with an investigation into the potential mechanistic modeling has towards understanding ecological phenomenon.

This dissertation also opens up further avenues in the philosophy and metaphysics of mind. Here I have largely presupposed Carl Gillett's work on realization. Seeing how Gillett's account fares against criticisms is the most direct way to contribute to the general project I began here. There are other approaches to realization that I have not discussed. Seeing how well these alternative conceptions work within a layered worldview is an interesting and worthwhile project. It is also worth investigating the extent to which the results of this dissertation apply to hierarchical modeling within cognitive science and neuroscience.

Third, this dissertation project serves as a foundation for further investigations in metaphysics. Questions surrounding the notion of *fundamentality* and *grounding* have been gaining interest amongst contemporary metaphysicians. It is common to associate

fundamentality with hierarchies: to be fundamental is to be at the bottom level. An investigation into the extent to which this association holds is a potential research project.

A more detailed investigation into the metaphysics of multi-level causation is also warranted by the results of this work: the dissertation ends with an argument that higher level things can make a causal difference to the world, but arguments do not pass for understanding. For a robust understanding of causation in multi-level contexts we must do more than wrestle with the causal exclusion argument; a coherent *account* of causation in multi-level contexts is needed.

The theory of hierarchies developed and defended in this work, Core Hierarchy Theory, is a mereological approach to understanding hierarchical structures. The important theorems of Core Hierarchy Theory were derived by exploiting a connection between hierarchies and mereology: classical mereology informs this understanding of hierarchies. Perhaps the informativity holds in the other direction as well: a more detailed investigation into the connection between hierarchies and mereology might shed light on, or help to reconceive, mereology itself. For example, Core Hierarchy Theory might lead to insights on the status of supplementation and extensionality principles in classical mereology.

This third area of future research suggests a fourth: philosophy of mathematics. In the dissertation I argued against a set-based characterization of hierarchies. It is then worth investigating whether or not the mereological approach to hierarchies leads to any insights about set theory, category theory, or the relationship between mathematical theories.

Formal ontology and bioinformatics might also benefit from my research.

Presenting data in a layered way is common in medical information systems, geographic information systems, and knowledge management information systems. The formal system presented in Appendix A to Chapter 3, I think, might serve as a foundational and universal hierarchical language for information systems. In any case, its potential must be more thoroughly examined.

References

- Allen T. F. H. and Starr, T. B. (1982), *Hierarchy: Perspectives for Ecological Complexity*, Chicago: University of Chicago Press.
- Allen, T. F. H. and Ahl, V. (1996), *Hierarchy Theory: A Vision, Vocabulary, and Epistemology*, New York: Columbia University Press.
- Anderson, P. W. (1972), "More is Different," *Science*, 177: 393-396.
- Baxtor, D. (1988), "Many-One Identity," *Philosophical Papers*, 17: 193-216.
- Beckermann, A. (1997), "Property Physicalism, Reduction and Realization," in M. Carrier, and P. Machamer, *Mindscapes: Philosophy, Science, and the Mind*, Pittsburgh: Pittsburgh University Press, pp. 303-321.
- Block, N. (2003), "Do Causal Powers Drain Away?" *Philosophy and Phenomenological Research*, 67 (1): 133-150.
- Bunge, M. (1969), "The Metaphysics, Epistemology, and Methodology of Levels," in L.L. Whyte, A.G. Wilson and D. Wilson (eds.), *Hierarchical Structures*, New York: Elsevier, pp. 215-228.
- Causey, R. (1972), "Attribute-Identities in Microreductions," *Journal of Philosophy*, 69 (14): 407-422.
- Causey, R. (1976), "Identities and Reduction: A Reply." *Nous*, 10 (3): 333-7.
- Colyvan, M. (1998), "Can the Eleatic Principle Be Justified?" *Canadian Journal of Philosophy*, 28 (3): 313-335.
- Conger, G. (1925), "On the Doctrine of Levels," *Journal of Philosophy*, 22 (12): 309-321.
- Craver, C. (2001), "Role-functions, Mechanism, and Hierarchy," *Philosophy of Science* 68: 53-74.
- Craver, C. (2002), "Interlevel Experiments and Multi-level Mechanisms in the Neuroscience of Memory," *Philosophy of Science Supplemental Volume*, 69: S83-S97.
- Craver, C. (2007), *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*, Oxford: Oxford University Press.
- Craver, C. and Bechtel, W. (2007), "Top-Down Causation Without Top-Down Causes," *Biology and Philosophy*, 22 (4):547-563.
- Dawkins, R. (1976), "Hierarchical Organization, a Candidate Principle for Ethology," in

- P. P. G. Bateson and R. A. Hinde (eds.), *Growing Points in Ethology*, Cambridge: Cambridge University Press, pp. 7-54.
- Deneault, J. and Ricard, M. (2006), "The Assessment of Children's Understanding of Containment Relations: Transitivity, Asymmetry, and Quantification," *Journal of Cognition and Development*, 7: 551-70.
- Enc, B. (1976), "Identity Statements and Microreductions," *Journal of Philosophy*, 73 (11): 285-306.
- Epstein, B. (MS), "Science Without Levels," Accessed (04/23): http://epstein.org/brian/PhilPapers/sci_without_levels.pdf.
- Feibleman, J. (1954), "Theory of Integrative Levels," *The British Journal for the Philosophy of Science*, 5: 59-66.
- Fodor, J. (1974), "Special Sciences: Or, the Disunity of Science as a Working Hypothesis," *Synthese*, 28: 97-115.
- Gibson, D. (2009), *Grasses and Grassland Ecology*, Oxford: Oxford University Press.
- Gillett, C. (2002), "The Dimensions of Realization: A Critique of the Standard View," *Analysis*, 62 (4): 316-22.
- Gillett, C. (2007), "Understanding the New Reductionism: The Metaphysics of Science and Compositional Reduction," *Journal of Philosophy*, 104 (4): 193-216.
- Gleason, H. A. (1926), "The Individualistic Concept of the Plant Association," *Bulletin of the Torrey Botanical Club*, 53: 7-26.
- Glennan, S. (1996), "Mechanisms and the Nature of Causation," *Erkenntnis*, 44: 49-71.
- Glennan, S. (2002), "Rethinking Mechanistic Explanation," *Philosophy of Science*, 69: 342-353.
- Glennan, S. (2010), "Mechanisms, Causes and the Layered Model of the World," *Philosophy and Phenomenological Research*, 81 (2): 362-381.
- Greene, T.R. (1989), "Children's Understanding of Class Inclusion Hierarchies: The Relationship Between External Representation and Task Performance," *Journal of Experimental Child Psychology*, 48: 62-89.
- Greene, T.R. (1994), "What Kindergartners Know About Class Containment Hierarchies," *Journal of Experimental Child Psychology*, 57: 72-88.
- Guttman, B. (1976), "Is 'Levels of Organization' a Useful Biological Concept?" *Bioscience*, 26: 112-113.

- Hempel, C. and Oppenheim, P. (1948), "Studies in the Logic of Explanation," *Philosophy of Science*, 15 (2): 135-75.
- Hull, David (1980), "Individuality and Selection," *Annual Review of Ecology and Systematics*, 11: 311-332.
- Jackson, F. and Petit, P. (1990), "Program Explanation: A General Perspective," *Analysis*, 50 (2): 107-17.
- Jagers op Akkerhuis, G. (2008), "Analysing Hierarchy in the Organization of Biological and Physical Systems," *Biological Reviews*, 83: 1-12.
- Kallestrup, J. (2006), "The Causal Exclusion Argument," *Philosophical Studies*, 131: 459-85.
- Kauffman, S. A. (1971), "Articulation of Parts Explanation in Biology and the Rational Search for Them," in R.C. Buck and R.S. Cohen (eds.), *PSA 1970*, Dordrecht: Reidel.
- Keller, D. and Golley, F. (2000), *The Philosophy of Ecology: From Science to Synthesis*. Athens: University of Georgia Press.
- Kemeny, J. and Oppenheim, P. (1956), "On Reduction," *Philosophical Studies*, 7: 6-19.
- Kim, J (1993), "Causation, Nomic Subsumption, and the Concept of Event," in his *Supervenience and Mind: Selected Philosophical Essays*. New York: Cambridge University Press.
- Kim, J. (1998), *Mind in a Physical World: An Essay on the Mind-Body Problem*, Cambridge: MIT Press.
- Kim, J. (1999), "Making Sense of Emergence," *Philosophical Studies*, 95: 3-36.
- Kim, J. (2002), "The Layered Model: Metaphysical Considerations," *Philosophical Explorations*, 5: 2-20.
- Kistler, M. (2006), "Mechanisms and Downward Causation," *Philosophical Psychology*, 22 (5): 595-609.
- Kitcher, P. (1984), "1953 And All That," *Philosophical Review*, 93 (3): 335-73.
- Koestler, A. (1967), *The Ghost in the Machine*, Oxford: Macmillan.
- LePore, E., & Loewer, B. (1989), "More on Making Mind Matter," *Philosophical Topics*, 17: 175-191.
- Lewis, D. (1991), *Parts of Classes*, Oxford: Blackwell.

- Lycan, W. (1995), *Consciousness*, Cambridge: MIT Press.
- Maddy, P. (1990), *Realism in Mathematics*, Oxford: Oxford University Press.
- Machamer, P. (2004), "Activities and Causation: The Metaphysics and Epistemology of Mechanisms," *International Studies in the Philosophy of Science*, 18: 27-39.
- Machamer, P., Darden, L., & Craver, C. (2000), "Thinking About Mechanisms," *Philosophy of Science*, 67: 1-25.
- Melnyk, A. (2003), *A Physicalist Manifesto: A Thoroughly Modern Materialism*, New York: Cambridge University Press.
- Murphy, G. (2002), *The Big Book of Concepts*, Cambridge: MIT Press.
- Nagel, E. (1961), *The Structure of Science: Problems in the Logic of Scientific Explanation*, New York: Harcourt, Brace, and World.
- O'Neill, R., DeAngelis, D., Waide, J., and Allen, T. (1986), *A Hierarchical Concept of Ecosystems*, Princeton: Princeton University Press.
- Oppenheim, P. and Putnam, H. (1958), "Unity of Science as a Working Hypothesis," *Minnesota Studies in the Philosophy of Science* 2, Minneapolis: University of Minnesota Press.
- Pickett, S. T. A., Jurek K., and Jones, C. (1994), *Ecological Understanding: The Nature of Theory and the Theory of Nature*, San Diego: Academic Press.
- Poli, R. (2001), "The Basic Problems of the Theory of Levels of Reality," *Axiomathes*, 12: 261-283.
- Poli, R. (2006), "The Theory of Levels of Reality and the Difference between Simple and Tangled Hierarchies," in G. Minati, E. Pessa, M. Abram (eds.), *Systemics of Emergence: Research and Development*, Berlin: Springer, pp. 715-22.
- Poli, R. and Gnoli, C. (2004), "Levels of Reality and Levels of Representation," *Knowledge Organization*, 21: 151-60.
- Potochnik, A. (2010), "Levels of Explanation Reconceived," *Philosophy of Science*, 77: 59-72.
- Potochnik, A. and McGill, B. (forthcoming), "The Limitations of Hierarchical Organization," *Philosophy of Science*.
- Ricklefs, R. (2008a), *The Economy of Nature, Sixth Edition*, New York: W. H. Freeman.

- Ricklefs, R. (2008b), "Disintegration of the Ecological Community," *The American Naturalist*, 172 (6), pp. 741-50.
- Rueger, A. (2006), "Functional Reduction and Emergence in the Physical Sciences," *Synthese*, 151 (3): 335-46.
- Rueger, A., and McGivern, P. (2010), "Hierarchies and Levels of Reality," *Synthese*, 176: 379-397.
- Sadava, D., Heller, C., Orians, G., Purves, W., Hillis, D. (2008), *Life: The Science of Biology, Eighth Edition*, New York: W. H. Freeman.
- Salthe, S. (1985), *Evolving Hierarchical Systems: Their Structure and Representation*, New York: Columbia University Press.
- Salthe, S. (2009), "A Hierarchical Framework for Levels of Reality: Understanding Through Representation," *Axiomathes*, 19: 87-99.
- Schaffer, J. (2003), "Is there a Fundamental Level?" *Nous*, 37 (3): 498-517.
- Schaffer, J. (2005), "Contrastive Causation," *The Philosophical Review*, 114: 297-328.
- Shapiro, L. (2008), "Understanding the Dimensions of Realization," *Journal of Philosophy*, 105 (4): 213-222.
- Shoemaker, S. (2007), *Physical Realization*, Oxford: Oxford University Press.
- Shoemaker, S. (2011), "Realization, Powers, and Property Identity," *The Monist*, 94: 3-18.
- Sider, T. (2007), "Parthood," *Philosophical Review*, 116: 51-91.
- Simon, H. (1962), "The Architecture of Complexity," *Proceedings of the American Philosophical Society*, 106 (6): 467-82.
- Simon, H. A. (1973), "The Organization of Complex Systems," in H. H. Pattee (ed.), *Hierarchy Theory*, New York: George Braziller, pp. 3-27.
- Simons, P. (1987), *Parts: A Study in Ontology*, Oxford: Oxford University Press.
- Slattery, B., Stuart, I., and O'Hara, D. (2011), "Testing for Transitive Class Containment as a Feature of Hierarchical Classification," *Journal of the Experimental Analysis of Behavior*, 96: 243-60.
- Steel, D. (2004), "Can a Reductionist Be a Pluralist?" *Biology and Philosophy*, 19: 55-73.

- Tansley, A. (2000), "The Use and Abuse of Vegetational Concepts and Terms," in D. Keller and F. Golley (eds.), *The Philosophy of Ecology: From Science to Synthesis*, pp. 55-70.
- Tuomisto, H. (2010) "A Diversity of Beta Diversities: Straightening Up a Concept Gone Awry. Part 1. Defining Beta Diversity as a Function of Alpha and Gamma Diversity", *Ecography*, 33: 2-22.
- Turney, P. (1989), "The Architecture of Complexity: A New Blueprint," *Synthese*, 79: 515-42.
- Van Fraassen, B. (1982), *The Scientific Image*, Oxford: Oxford University Press.
- Varzi, A. (1999), *Parts and Places*, Cambridge: MIT Press.
- Voorhees, B. (1983), "Axiomatic Theory of Hierarchical Systems," *Behavioral Science*, 28: 24-34.
- Webster, J. (1979), "Hierarchical Organization of Ecosystems," in Halfon (ed.), *Theoretical Systems Ecology*, New York: Academic Press, pp. 119-29.
- Whittaker, R. J., Willis, K. J., Field, R. (2001), "Scale and Species Richness: Towards a General, Hierarchical Theory of Species Diversity," *Journal of Biogeography*, 28: 453-470.
- Wimsatt, W. (1976), "Reductionism, Levels of Organization, and the Mind-Body Problem," in Globus, Maxwell and Savodnik (eds.), *Consciousness and the Brain*, New York: Plenum Press, pp. 199-267.
- Wimsatt, W. (1994), "The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets," *Canadian Journal of Philosophy Supplement*, 20: 207-74.
- Wimsatt, W. (2007), *Re-engineering Philosophy For Limited Beings*, Cambridge: Harvard University Press.
- Winer, G. A. (1980), "Class Inclusion Reasoning In Children: A Review of the Empirical Literature," *Child Development*, 51: 309-28.
- Woodward, J. (2003), *Making Things Happen: A Theory of Causal Explanation*, Oxford: Oxford University Press.
- Wrenn, C. (2010), "The Unreality of Realization," *Australasian Journal of Philosophy*, 88 (2): 305-22.