The Universe As We Find It

JOHN HEIL

For my Antipodean heroes, Jack Smart, David Armstrong, Keith Campbell

O World, thou choosest not the better part!
It is not wisdom to be only wise,
And on the inward vision close the eyes,
But it is wisdom to believe the heart.
Columbus found a world, and had no chart,
Save one that faith deciphered in the skies;
To trust the soul's invincible surmise
Was all his science and his only art.
Our knowledge is a torch of smoky pine
That lights the pathway but one step ahead
Across a void of mystery and dread.
Bid, then, the tender light of faith to shine
By which alone the mortal heart is lead
Unto the thinking of the thoughts divine.

G. Santayana

It's good to know
the earth is there,
compact below
the actual air,

its substance kept
immune, opaque,
when I have slept
as when I wake.

The clay commutes
its dark duress
to feed the roots
of consciousness,

and, thought or sensed,
the spirit's act
is shaped against
the stone of fact.

The levers set
by our purpose lock
with a purchase let
in the living rock.

The world's uncouth
old lengths decree
what chains of truth
shall make us free.

D. C. Williams

Contents

Preface

This volume incorporates my best efforts to articulate a metaphysical framework that illuminates the most general features of the universe as we find it: a fundamental ontology. I have tried to express myself clearly and in a way that presupposes no more than a passing acquaintance with topics central to contemporary metaphysics. In my judgement, there is little to be gained and much to be lost when we philosophers couch arguments in a style congenial only to other philosophers. My models in this endeavour are my Australasian heroes, Jack Smart, David Armstrong, and Keith Campbell, philosophers of unsurpassed ontological integrity to whom this volume is dedicated.

In writing on metaphysical topics, we are prone to take refuge in a technical vocabulary that obscures more than it reveals. Technical terms encode substantive doctrines that tend not to be scrutinized by those deploying the terms. The precision of arguments relying on such terms can disguise the fact that the real work is going on behind the scenes. Sharing a technical vocabulary is to share a tidy collection of assumptions. Reliance on that vocabulary serves to foreclose discussion of those assumptions. The goal should be to *arrive at* a stable of technical terms on the basis of reflection and investigation, rather than relying on them at the outset of investigation.

Today non-philosophers—academics, journalists, readers generally—are especially vocal about the 'irrelevance' of contemporary philosophy. Philosophers, it is said, discuss topics of interest only to themselves and in a way intelligible only to specialist fellow philosophers. To the extent that this is so, it is due partly to our chronic dependence on specialized terminology and our unwillingness, or maybe an inability, to express our thoughts plainly, to express them without reliance on in-group jargon.

Another, quite separate, factor that militates against the public embrace of philosophy is that the subject matter is often difficult for non-philosophers to appreciate or find compelling. This has always been so. Journalists, pundits, and cultural arbiters who lament what they regard as the decline of philosophy are apt to cite pressing moral

and political topics as examples of what philosophers ought to be addressing. But of course philosophers *do* engage with such topics and in a manner likely to be accessible to anyone with sufficient patience and interest. This is not *all* philosophers do, however, and it has never been all philosophers have done (see Saunders 2009).

When cultural arbiters bemoan the current state of philosophy, you might wonder whether they have ever themselves read Aristotle, or Ockham, or Aquinas, or Descartes, or Locke, or Spinoza, or Kant. The most important and influential works of these philosophers are difficult even for specialists. If the ideal is Ralph Waldo Emerson, you would need to toss out the towering figures of the past along with lesser lights of today. In my view, fundamental ontology is unavoidable once you begin thinking hard about the universe, but its unavoidability does not mean that the problems it addresses can be domesticated or that serious efforts to cope with these problems can be turned into best-sellers.

Although I very much hope that my discussion in the chapters that follow makes clear why metaphysics, and in particular, ontology, is worth pursuing, I am not naïve enough to imagine that very many non-specialists will care. Happily, as a trained and calloused philosopher, it is easy for me to be philosophical about such things. I shall rest content if what I have to say opens an occasional door and leads others more talented than I to improve on these uncertain steps.

Philosophers steeped in contemporary analytical metaphysics reading these words are hereby warned that I shall not be engaging much with their work. My focus differs from theirs. My style of argument and the positions I favour are closer to Locke's and Spinoza's than they are to those of philosophers currently publishing papers on topics in metaphysics in mainstream journals. Exceptions again include Armstrong, Campbell, Smart, and C. B. Martin. I see these philosophers as rooted in a tradition that includes Descartes, Locke, and Spinoza, among many others, a tradition largely supplanted by the linguisticization of philosophy in the twentieth century. I am content to let the linguisticizers pursue their interests. My aims are different, or apparently so. I lack the heart, the patience, and certainly the wits to engage in debates that currently occupy centre stage.

Some readers will be put off by my characterization of contemporary analytical metaphysics as 'linguisticized'. However, the ongoing

tendency to conflate predicates and properties, truth conditions and truthmakers, philosophy of language and metaphysics, makes my point. I see this as the legacy of the programme inspired by Wittgenstein, Carnap, and their successors to reduce philosophical, and, in particular, metaphysical, questions to questions about language and linguistic practice. Philosophy might have left the programme behind, but its influence lingers, all the more insidious by virtue of being unacknowledged, implicit—or even explicitly disavowed.

In an effort to improve readability, I have tried to avoid unnecessary formalisms and worked to keep footnotes to a minimum. I regard the ever increasing use of footnotes to evince a kind of self-indulgence that characterizes too much of today's philosophical writing. We philosophers cannot resist offering our opinions on every detail of every aspect of whatever we choose to discuss. That this might be distracting for a reader, that it might be annoying or off-putting, seems of little concern.[1] What I have to say will, on occasion, doubtless annoy you the reader, but I shall at least spare you the kinds of mental fatigue brought on by footnote frippery.

A final disclaimer. I profess no special originality for what I have written here. The book reflects influences of a number of historical figures, in particular my early modern and Enlightenment paragons—Locke, Descartes, Spinoza—and more recent philosophers with whom I have crossed paths, including Smart, Armstrong, Campbell, and especially Charlie Martin, the philosopher's philosopher. Virtually every idea on these pages is traceable in one way or another to these figures—and to D. C. Williams, whom I never met, but whose philosophical instincts mark my own.

I am indebted to Donald Davidson, E. J. Lowe, John Bigelow, Jaegwon Kim, Jonathan Bennett, Wallace Matson, and my erstwhile colleague, David Robb. These philosophers have shaped my thoughts on all the topics addressed here. Their philosophical influence is of a piece with their influence on me as human beings. Philosophers are not invariably the best people, but the best philosophers I know are the best people I know. Other philosophers with whom I have talked and corresponded led me to think more clearly and distinctly about particular topics. These include Elizabeth Prior, David Sanford, Galen Strawson, Michael Esfeld, Anna Marmodoro, Dennis des Chene, Laura Franklin-Hall, Robert Garcia, Phil Dowe,

---

[1] See what I mean?

CHAPTER I

Introduction

When the objects of an inquiry, in any department, have principles, conditions, or elements, it is through acquaintance with these that knowledge . . . is attained. . . . The natural way of doing this is to start with the things that are more knowable and obvious to us and proceed toward those that are more knowable and clearer by nature; for the same things are not 'knowable relative to us' and 'knowable' without qualification. In the present inquiry we must follow this method and advance from what is more obscure by nature, but clearer to us, towards what is more clear and more knowable by nature.

(Aristotle, *Physics* 184a 1–20)

1.1 Ground Rules

In the pages that follow, I advance an ontological picture of the universe as we find it, a picture meant to be realist (the universe is as it is independently of our thoughts about it), particularist (existing things are particular; generality is a feature of our ways of representing the universe), naturalistic (the universe is all there is), and from the gut (it stems, not from a nuanced analysis of talk about the universe, but from repeated head-on confrontations with the universe). Minds and their contents, no less than billiard balls, trees, and planets, are what result from arrangements of the fundamental things. The mental–physical distinction is one of conception only, not what would traditionally have been called a *real* distinction. This is not dualism, not physicalism, not materialism; but it is not reductionism either. If the mental–physical distinction is not ontologically deep, there is no question of reducing the one to the other.

A picture of this kind has its roots in age-old reflections on the nature of being, the nature of what there is. What there is, however,

is not something that could be ascertained by looking at what we 'quantify over' in our best confirmed theories. It is one thing to know what theories you accept, to know what you take to be true, quite another matter to know what the world is or must be like if those theories *are* true. This is so for philosophy no less than it is for the sciences. In both cases, the aim is to discover, not just the truths, but the nature of the truthmakers for these truths.

This is the soul of realism. You can have good reason to accept a theory, yet that theory could turn out to be false. You can have good reason to accept a *true* theory, yet have no very clear understanding of what it is about the world that makes the theory true, the nature of the truthmakers. The universe disclosed by fundamental physics could turn out to be utterly surprising. The universe might be not at all as we currently take it to be. Whatever conception we develop, however, will be one encompassing substances and properties. The substances might be corpuscles, or electrons, or superstrings, or fields, or space–time itself, or something stranger still. Whatever the substance is or the substances are, it or they must *be* various ways: substances must have properties.

These assertions are meant to reflect not merely a limit on our thoughts about the universe, but a limit on the *universe*, a limit on how the universe, *any* universe, *could be*, hence a limit on scientific theorizing about the universe.

Think of fundamental physics as being in the business of telling us how the universe *is*. Fundamental ontology is in the business of telling us what the universe must be like if *any* theory is true. In this way ontology constrains science. Ontology does so, not by laying down immutable principles a priori, but by working in concert with science to discern the texture of being. One perhaps surprising consequence of this endeavour is the discovery of how often our predecessors were on target, even when they were wrong about the details.

1.2 The Province of Metaphysics

For millennia, philosophers engaged in metaphysical reflection without paying much heed to the nature of the enterprise. That changed with the advent of empiricism in the eighteenth century. Hume

argued that metaphysical theorizing divorced from empirical observation was empty, a projection of our own psychological makeup with no objective standing. Kant, taking Hume's point, made a virtue of necessity. To the extent that metaphysics affords an accounting of the structure of our thoughts about the universe, it provides an indispensable service. The mistake, Kant thought, was to imagine that metaphysical theses might reveal the nature of an objective, mind-independent reality.

Hume and Kant were right to be sceptical of the idea that metaphysics could provide a direct, unfiltered pipeline to reality. What knowledge anyone has of the universe is grounded in experience and observation tempered by scientific enquiry. When our interest is in the nature of things we turn to the sciences. When our interest extends to the deep story about those things we turn to fundamental physics. Fundamental physics provides an account of the truthmakers for scientific claims generally. Although some readers will hear this assertion as an expression of a kind of extreme reductionism, I argue that it is not. Confusion on this point has muddied the water, inhibiting progress in metaphysics, pure and applied. One symptom of this is the unfortunate tendency to conflate epistemological issues bearing on explanation with issues in metaphysics. Another is the lingering reluctance to distinguish metaphysics from philosophy of language. Talk about talk about the universe is not talk about the universe.

But this is to get ahead of the story. Back to Kant. Kant was right in supposing that the aim of metaphysics is, or ought to be, to disclose our fundamental categories. These categories are not merely artifacts of ways we have of thinking about the universe, however, parochial cognitive spectacles, ripe for debunking by social psychologists and experimental philosophers. Categories are required for any intelligible thought about the universe. The mistake would be to imagine that this turns the universe, or the universe as revealed by science, into a construct.

1.3 Substance and Property

A central theme in what lies ahead concerns the nature of substances and properties. Substance and property are basic, indeed *the* basic,

ontological categories. Think of these categories as complementary: substances are property bearers; properties are ways substances are. Every substance is some particular way or other, every substance is propertied; every property is a particular way some substance is.

Spelling out this thesis requires starting with everyday common-sensical examples: billiard balls, trees, tables. Such familiar cases provide a sense of the categories, a grasp of what it is for something to be a substance or a property. With this understanding in hand, you can proceed to reconsider the everyday cases, the billiard balls, trees, and tables. When you do so, you discover that substances, as property bearers, must be simple. Substances of necessity lack constituents that are themselves substances. For their part, properties, ways substances are, must be properties of simples. From this, together with what we think we know about such things, it follows that billiard balls, trees, and tables are not after all substances, their apparent properties are not in fact properties. The substances are propertied simples, the fundamental things. This conception of substances and properties flows directly from the idea of substances as property bearers and properties as particular ways substances are.

You might worry that this narrow conception of substance and property yields a kind of scepticism or anti-realism concerning every-day, medium-sized objects and their properties. But to say that a billiard ball is not a genuine substance or that the billiard ball's sphericality or redness are not genuine properties is not to say that there are no billiard balls or that it is false that this billiard ball is red and spherical. Truthmakers—*truth*makers—for judgements about billiard balls and their colours or shapes are particular arrangements of the fundamental substances. (At least this is how it is if the universe comprises arrangements of fundamental particles. If the universe is continuous, then the truthmakers might be particular ways the universe-substance is. I shall leave aside this qualification for the moment, returning to it in chapter 2.)

So, a second theme advanced here is that we philosophers are not in a position to ascertain truthmakers for everyday judgements about the universe a priori. You can know the application conditions for talk of billiard balls, trees, and tables, you can know how to use the terms, and you can know that claims about billiard balls, trees, and tables are often true, without having any clear idea as to the nature of

truthmakers for judgements about such things. In particular, there is no hope of recovering or 'reading off' the character of the truth-makers from an analysis of the application conditions for terms deployed in science and in everyday life to state truths about the universe. Only in the case of fundamental physics do you begin to get at the deep story, the nature of the truthmakers.

In advancing this view, I do not take myself to be embracing a radical thesis, but merely to be echoing Descartes, Locke, Spinoza, and countless other philosophers from Plato and Aristotle through the medievals to the late eighteenth century. Locke, for instance, thinks that the real substances are the corpuscles. Billiard balls, trees, and planets are fleeting arrangements of these, substances by courtesy. Descartes holds that there is but one extended material substance, space itself. Billiard balls, trees, and planets are local 'thickenings' of space, ways space is. (I owe this way of putting it to Jonathan Bennett.) For both Descartes and Locke, ordinary objects are not substances; ordinary objects are, to a first approximation, *modes*, particular ways substances are arranged or space is configured. Nei-ther Locke nor Descartes regarded this as casting doubt on billiard balls, tables, or trees. Judgements about such things can be, and often are, perfectly true. Their truthmakers, however, are not substances corresponding to 'billiard ball', 'table', 'tree'. Truths about the uni-verse do not require a substance corresponding to every singular term, a property corresponding to every predicate.

These topics are addressed in more detail in the chapters ahead. Here, the aim is only to fend off worries that I might be relying on premises concerning everyday objects and their characteristics to establish conclusions that, if true, would call those very premises into question. In getting clear about the universe, science and phil-osophy alike must begin, as Aristotle notes, with the familiar, the 'manifest image', and move to the unfamiliar, to the deep story. The deep story includes an account of the manifest image and why it is as it is. The manifest image is not, or need not be, false or illusory. Judgements about billiard balls, their shapes and colours can all be perfectly true. Such judgements can provide descriptions of the universe that we find indispensable.

So long as the goal is to produce true descriptions and explanations of phenomena of interest in everyday life and in the sciences

(save fundamental physics), taxonomies we deploy in such endeavours are, for the most part, entirely adequate. If you want to learn about trees, you consult the biologist, not the physicist. Yet there is an important respect in which the deep story about trees is what you would find were you to scrutinize trees through the lens of fundamental physics.

Again, this is not a reductive claim. This is not the claim that you could translate or analyse talk of trees into talk of electrons, quarks, fields. It is not the claim that you could replace biological taxonomies, concepts, or terms with taxonomies, concepts, or terms at home in fundamental physics. And it is very definitely not the claim that you could provide application conditions for biological or psychological predicates in fundamental physical terms. More generally, the suggestion is not that 'higher-level' judgements (or explanations framed in terms of these judgements) of the sort you find in the various sciences and everyday life could be translated into, or replaced by, judgements (or explanations) couched in the vocabulary of fundamental physics.

Thus, although it might be straightforwardly true that this billiard ball is red and spherical, the truthmaker for the claim is not the possession by a substance, a billiard ball, of a pair of properties, redness and sphericality. A billiard ball is not a substance, and its redness and sphericality are not properties.

Someone might complain that in treating substances as simples and properties as belonging only to these simples, I am making a hash of the ordinary notion of property: I am using 'property' in a technical, stipulated sense. Were that so, my conclusions would be uninteresting: there are *properties* (the ball's redness and sphericality, for instance) and Properties-with-a-capital-*P*. Why should anyone care about *those*?

This misses the point. In treating a particular billiard ball's redness and sphericality as properties of the ball, you are treating the billiard ball as a substance, as a *one*. *Treating* the ball as a substance and the ball's redness and sphericality as properties of the ball is one thing, however, the ball's *being* a substance, the ball's redness and sphericality's *being* properties, is another matter altogether. Suppose that the ball is, in fact, a particular dynamic arrangement of interacting substances. Then the ball is not a substance, but a *mode*, a particular way particular substances are arranged here and now: a particular arrangement of a

particular sort. The ball's redness and sphericality are not properties but consequences of this arrangement: what you get when you arrange these substances in this way.

Take three matches and arrange them so as to form a triangle. The triangle—the truthmaker for 'this is a triangle'—*is* these matches in this arrangement. You do not have the matches, with their properties, so arranged, *plus* a triangle and *its* properties. What goes for the triangle goes for the individual matches as well. You do not have *these* particles, with their properties, duly arranged, *plus* the matches and *their* properties.

We begin, as we must, with a common-sense conception of the universe that treats billiard balls as substances, property bearers, and redness and sphericality as properties of billiard balls. In pursuing the idea of substances as bearers of properties, however, we come to recognize that the common-sense conception contains the seeds of its own revision, revision in light of empirical discovery, revision in light of what we determine to be the deep story about billiard balls. We discover that billiard balls are not in fact substances, and properties we ascribe to billiard balls are not in fact properties. This is not to replace the notions of substance (as property bearer) and property with proprietary notions, but to recognize that plausible conceptions of substance and property mandate distinguishing between what we ordinarily *treat* as substances and properties and what the genuine substances and properties really are.

Talk of properties leads to thoughts of *universals*. Properties construed as universals have *instances*. A dozen billiard balls provide a dozen instances of a single universal sphericality. Although there is a place for talk of universals, truthmakers for such talk are fully particular. Properties are *modes*, particular ways particular substances are. One billiard ball's sphericality is distinct from another's. I prefer the traditional *mode* to *trope*, a term first deployed by D. C. Williams in a defence of a one-category ontology. 'Trope theorists', by and large, accept Williams's conception of substances as 'bundles' of tropes. I side with Locke and with Descartes, who noted that properties standing aloof from substances would themselves be propertied substances.

We have general truths—about sphericality and redness, for instance—but truthmakers for these truths are particular ways the

universe is, particular ways particular substances are. Generality is an important and irreplaceable feature of representations. Theories as standardly formulated consist exclusively of general assertions. But truthmakers for general truths—general truths that have truth-makers—are particular ways the universe is. The universe includes no general or universal *entities*.

The current philosophical climate is not friendly to a conception of properties as modes. Many contemporary philosophers regard 'property' and 'universal' as synonymous. Of the few who accept properties as particulars, even fewer embrace the kind of substance–mode ontology advanced here. It was not always so. Many, indeed *most*, of our most venerated philosophers regarded properties as particular ways particular substances are.

1.4 Relations

The billiard ball is red and spherical, but the billiard ball is not a substance and, consequently, its redness and sphericality are not properties. If you take the billiard ball's constituents—pretend for the moment that these are particles—and organize them just so, you have something of which it is true that it is a billiard ball, that it is red, that it is spherical. The organization is important. Differently organized, or widely dispersed, the billiard ball's constituents do not amount to a billiard ball. This is easier to see in the case of a complex object, a watch, for instance. You need more than the parts to have a watch, you need the parts assembled in precisely the right way.

All this makes it appear obvious that fundamental ontology requires, in addition to substances and properties, *relations*. I admit that this is how it *appears*. I admit, as well, that we need relational predicates to say all we have to say about the universe. Relations are, in this regard, ineliminable. What I am not prepared to admit, however, is that relational truths require relational truthmakers. At any rate, this is not something I am prepared to admit without a fight. Truthmakers for relational truths *could* turn out to be nonrelational features of the universe. Were that so, a two-category—substance–property—ontology would provide all anyone could want by way of truthmakers for everyday and scientific truths.

But why disdain relations? I follow a long line of philosophers, beginning with Aristotle, and running though the medieval period and the Enlightenment, who regarded relations as, at best, ontological anomalies. It is a measure of how far we have fallen that so few philosophers nowadays see any difficulty at all in the idea that relations have full ontological standing. 'What's the problem? There are relational truths; we *quantify over* relations; so we are *ontologically committed* to relations. No big deal.'

One of the aims of this volume is to encourage a visceral feeling for metaphysical hypotheses. Recent psycho-neurological studies suggest that metaphysics, ontologically serious metaphysics, is done in the gut, not in regions of the brain responsible for delicate, formal cognition (Below 1987). Metaphysical enquiry suffers when it is reduced to the kind of bloodless abstraction that results when metaphysics is replaced by conceptual analysis. Paraphrasing C. B. Martin (himself paraphrasing Locke), if you can't *live* it, you can live without it.

The discussion in chapter 7 of the ontology of relations incorporates an attempt to make salient reasons philosophers prior to the mid-twentieth century were so often uncomfortable with relations. This does not amount to a disproof of relations, whatever that might mean, but it does serve to motivate the attempt to identify nonrelational truthmakers for relational truths. Relational truths most resistant to this treatment are truths concerning causal, temporal, and spatial relations. My hope is that by the time the topic of relations is addressed head-on, the road to a successful treatment of relations will have been paved by discussion of various other topics, including causation.

1.5 Truthmaking

Truthmaking plays an indispensable role in the evaluation of ontological theses. You could think of truthmaking as an internal relation between a truth bearer—a judgement, a representation capable of truth or falsity—and some way the universe is, the truthmaker. To say that the truthmaking relation is *internal* is to say that, if you have the relata, if you have the judgement that snow is white and you have

snow's being white, you thereby have the relation. An internal relation is (as D. M. Armstrong puts it) 'no addition of being'. (It goes without saying that what it is for snow to be white could be, and assuredly is, a complicated matter.)

I do not regard truthmaking as a technical concept. A grasp of the notion of truthmaking goes hand in hand with a grasp of the notion of truth: each requires the other. Nor do I subscribe to the thesis that every truth requires a truthmaker. Mathematical truths and truths of logic are compatible with any way the universe could be. Mathematics and logic are informative, not in telling us how the universe is, but in enabling us to put two and two together.

Despite—or perhaps because of—its apparent obviousness, the notion of truthmaking has only recently come to be deployed self-consciously in discussions of metaphysical topics. In looking back over the history of various disputes in philosophy, however—disputes over the standing of relations, for instance—it is difficult to avoid the impression that, surprisingly often, philosophers, apparently at odds, were in fact struggling to say something like the same thing, struggling to identify truthmakers for truths of a particular kind.

Today confusion abounds in the philosophy of mind owing to the conviction that conceptually distinct truths require distinct truth-makers. The irreducibility of sociological, psychological, biological truths to truths of physics, however, provides no reason whatever to think that truthmakers for these truths could not be expressed in terms of fundamental physics.

This is a key point in all that follows. If you deny it, if you suppose that the impossibility of translating talk of flora and fauna, or talk of actions and emotions, into talk of electrons, quarks, and fields, shows either that there are no flora, fauna, actions, and emotions, or that these things must exist alongside, 'above', or in addition to, the quarks, electrons, and fields, you will find what I have to say here entirely unconvincing. My hope is that I can persuade impartial readers to share my conviction that metaphysics has been too long in the thrall of the linguisticizers, those who believe that deep truths about the universe are to be had by analysing ways we talk about the universe. Language takes us *to* the universe, to ways the universe is, but leaves open the nature of those ways. To discover that, we need science.

1.6   The Big Picture

I have emphasized that the metaphysical picture defended here is of a piece with conceptions implicit in Locke and Spinoza—and, in whole or in part, in the work of many other historically influential philosophers. This might lead naturally to the thought that I am merely repackaging mistakes of the past. Metaphysics developed alongside science. But science has come a long way since Aristotle. The science of Locke and Newton is not our science. Why should anyone imagine that categories that worked well enough for the Milesians, the medievals, and philosophers in the Enlightenment would work in an era encompassing relativity and quantum mechanics?

Metaphysics, and in particular ontology, sets the limits of scientific theorizing, at least in this sense: in plotting basic categories of being, ontology constrains science. These constraints are not externally imposed rules. The constraints are the expression of principles we are all—including the scientists—bound to accept. Such principles are not based on armchair reflection, but on our ongoing give and take with being.

What of the Big Picture? The strategy of divide and conquer, often useful in solving complex problems, can be self-defeating in the case of fundamental ontology. The Big Picture is not a patchwork of a lot of little pictures. The little pictures (concerning the nature of properties and property bearers, causation, consciousness, and the like) are aspects of a Big Picture. In my own case, I discovered years ago in the course of wrestling with the problem of mental causation, that, like it or not, in ontology one thing leads to another. In practice this means that, in working out a fundamental ontology, you must be prepared to adjust and readjust accounts of particular topics in light of their implications for all the rest. This is the thought animating the chapters that follow.

CHAPTER 2

Substance

Substance, in the truest and primary and most definite sense of the word, is that which is neither predicable of a subject nor present in a subject; for instance, the individual man or horse.

(Aristotle, *Categories*, v, 11–13)

Substance without knowing what it is, is that which supports *Accidents*. So that of Substance, we have no *Idea* of what it is, but only a confused obscure one of what it does.

(Locke 1690, II, xiii, 19)

Substances are unities if anything is. A single substance is a genuine, literal, case of *one thing*. Everyday objects are substances by a kind of courtesy; it suits us, for practical purposes to *treat* as single things knives and forks, whether they are of uniform stainless steel or have bone handles and blades of silver plate. Similarly for all the multitude of manifestly complex objects and artifacts which make up the familiar world.

(Campbell 1990, 10)

2.1 Substances as Property Bearers

Substances are property bearers; properties are ways substances are.[1] If there are substances, there are properties; if there are properties, there are substances. Every substance is some way or other, every property is a way some substance is. Substance and property are complementary categories of being. The idea is expressed by Locke's contention that substance and property are 'correlative'; they 'stand or fall

[1] On properties as ways, see Levinson 1978, 1980; Seargent 1985; Armstrong 1997, 30–1; Heil 2003a.

together' (see Martin 1980, 9). Might *properties* sometimes play the substance role, might properties be property bearers, might there be properties of properties, second-order properties? If a property is a way something is, a property of a property, *P*, would be a way *P* is. But this is just *P* itself. (See chapter 4 for further discussion.)

David Armstrong (1997) describes substances as 'thin particulars', constituents of states of affairs. States of affairs, on this view, are substances' instantiating properties at times. Socrates's being warm on Tuesday is a state of affairs. Socrates's being snub-nosed on Tuesday is a distinct state of affairs as is Socrates's being warm on Wednesday. Because the same substance can instantiate many properties at once, what we regard as ordinary objects turn out to be complex states of affairs: 'thick particulars'.

Armstrong's universe is a universe of states of affairs: states of affairs are the fundamental existents. Armstrong's substances and properties, however, do not make up states of affairs in the way stones make up a wall. Substances and properties alike are *abstractions*. Once you have the states of affairs, you can recognize that distinct states of affairs share constituents. Socrates is common to a host of distinct states of affairs that otherwise differ: Socrates's being snub-nosed, and Socrates's being warm, for instance. Similarly, being snub-nosed and being warm are common to distinct states of affairs that differ otherwise: Simmias's being snub-nosed, Glaucon's being warm.

Likewise, on the view I am recommending, substances are not ingredients, components, or parts of wholes that include properties. No substance could fail to be some way or other. A substance might cease to be some particular way, but it could cease to be that way (while remaining in existence) only by being some other way. A spherical lump of clay can cease to be spherical, but only by being some other way—disk-shaped, perhaps—or by ceasing altogether to exist, as when it encounters a solvent.

Armstrong's conception of substances is motivated in part by his commitment to the idea that properties are universals. If this is how you think about objects' properties, you will need something that is not a universal to provide for particularity. If states of affairs or objects were collections of universals, states of affairs or objects would be complex universals, not particulars; states of affairs or objects would themselves be repeatable, 'multiply locatable', entities. Armstrong's

universals have being only in their instances, however, only when paired with some particular in a state of affairs. This is the 'victory of particularity' (Armstrong 1997, 126–7).

I shall have more to say about universals in chapter 5. For the moment I merely register the thought that properties could themselves be particulars: *abstract* particulars. Socrates's whiteness might be a particular way a particular substance, Socrates, is. Philosophers nowadays embrace D. C. Williams's nomenclature in calling properties considered as abstract particulars *tropes*.[2] I much prefer the traditional label, *mode*, but I shall use the neutral term, *property*, in neutral contexts.

If you thought properties were universals, you would have a pressing need for substances as particularizers. If, in contrast, you are attracted to the idea that properties are themselves particular, particularity is guaranteed without the need for an additional something to play the particularizing role. Particular objects *could* be collections or bundles of properties without compromising their particularity. This line of thought is partly responsible for the prominence of bundle theories among philosophers today who embrace tropes. Such philosophers regard the substance category as dispensable: you can make do with a more parsimonious one-category ontology.

Another consideration at work here is traceable to Hume, who was impressed by the fact that, when we scrutinize an object, we seem only to encounter that object's properties, never a property bearer.

We have therefore no idea of substance, distinct from that of a collection of particular qualities, nor have we any other meaning when we either talk or reason concerning it. (Hume 1739 I, 1, vi)

Hume is echoing Locke's apparent qualms about substrata, property bearers that seem to lack any natures of their own. You have a substance's properties and you have the substance itself, a mysterious, inchoate 'something-we-know-not-what' (Locke, 1690, I, xxiii, 2). So, on the one hand, if you are happy with the idea that properties are

2 See, for instance, Campbell (1976, 1990); Martin (1993, 2008); Simons (1994); Bacon (1995, 2008); Robb (1997); Schaffer (2001); Maurin (2002); Levinson (2006); all of whom follow Williams's (1953) terminological lead.

particulars, you have no need to posit substances to introduce particularity into the universe. On the other hand, in dispensing with substances, you avoid an apparent ontological embarrassment and the prospect of denatured 'bare particulars' cloaked in properties. Indeed it might be thought that one important advantage a conception of properties as particulars enjoys over one countenancing universals is that it enables us to dispense with substances.

2.2 Substance Situated

Such considerations derive their plausibility in part from the fact that we began with universals and a need for particularity in working out the substance role. Substances are depicted as constituents of concrete objects or states of affairs, receptacles for properties. Properties condense around substances. States of affairs, or objects, comprise substances plus properties. Once the need for substances as particularizers vanishes—as it does when you regard properties as particulars—so does the point of positing substances. You are left with somethings-we-know-not-what that you could have no reason to think exist.

This is the wrong way to look at substance and substances, however. Substances are not bare, featureless entities to which properties attach themselves as limpets attach themselves to rocks at the seashore. Every substance is *itself* some way or other, indeed many ways. These ways are its properties. For a substance to possess a property is for it, the substance, to be a particular way. Properties—ways—do not make up a substance, they are not parts of substances. The charge, spin, and mass of an electron are not parts or constituents of the electron. As far as we know, electrons have no parts. Electrons might have spatial or temporal parts, but that is another matter, one I shall take up in due course. An electron's charge, spin, and mass are ways the electron is.

Considerations of this sort are of the first importance when you reflect on roles substances might plausibly play. A substance, an electron, for instance, is neither a compound including a 'bare particular' larded over with properties, nor a bundle or aggregation of properties. You can consider an electron as a substance, as a bearer of properties, as something that has a particular nature, something that is

various ways. You can consider ways an electron is. In so doing, you are abstracting, engaging in what Locke called 'partial consideration.' Imagine a ripe tomato illuminated by bright sunlight. You can consider the tomato's shape, its colour, its heft, or you can consider the tomato as something that *has* a particular shape, colour, heft. When you do so, you are not considering parts of the tomato, ingredients that, taken together, add up to a tomato. You are considering the tomato and ways it is. And just as you can consider the tomato's shape without considering its colour, even though these are, at any given time, inseparable, so you can consider the tomato without considering various ways it is.

The categories of substance and property are fundamental and complementary. To think of a substance is to think of something that is various ways; to think of a property is to think of a way a substance is or could be. A substance cannot be no way at all, and a property cannot fail to be a property of a substance, a way a substance is or might be. Philosophers who have tried to dispense with substances—bundle theorists—and those who have tried to dispense with properties—extreme nominalists—begin and end with the wrong picture.[3] If properties are ways, they must be ways something is. And if there is something, it must be some way or other. The idea that you might, in the interests of parsimony, get by with just a single category might be likened to the idea that you could have the smile without the cat—or the cat with no expression at all.

2.3 Substance and Property

I have contended that properties go hand in hand with substances: each requires the other. The contention begins with properties and ends with substances. Thus, although my position is that properties and substances are, so to speak, made for each other, you might worry that I have stacked the deck by helping myself to properties at the outset. Substances aside, do we really need *properties*? Perhaps all there is are objects, particulars that single-handedly combine the roles of substance and property.

---

3 Nominalism is usually identified with the denial of universals. By 'extreme nominalism' I mean a view that forgoes properties of any sort, whether universal or particular.

One way to interpret such a position is to understand it as a slightly misleading reformulation of the view I have been advancing. Even if this is the right thing to say, however, we can do better. Return to the tomato. The tomato looks red because of a particular way it is—its colour, how it is colourwise—and spherical because of another way it is—its shape, how it is shapewise. As C. B. Martin never tired of noting, it is not the tomato *holus bolus*, but the tomato's colour and shape, respectively, that are responsible for the truth of 'the tomato is red' and 'the tomato is spherical' (Martin 1980). Similarly, when the tomato depresses a scale, it does so in virtue of its mass—how it is masswise—and *not* in virtue of its colour or shape. Were we barred from saying such things, we would be unable to formulate truths about the fundamental things. An electron repels a fellow electron owing to its charge and not its mass; it accelerates as it does in a gravitational field because of its mass and not its charge.

Might talk of properties be replaced with talk of sets of objects? Might properties be eliminated in favour of sets? Might the tomato's being red be solely a matter of its being a member of the set of red things? Might the tomato's being spherical be solely a matter of its belonging to the set of spherical things? Might talk of properties just be an oblique way of talking of sets? Because sets could be regarded as a kind of ontological 'free lunch'—if you have the objects, you have the sets—the upshot would be an impressively parsimonious ontology.

A natural reaction to such a proposal is to pose a *Euthyphro* question: is something red because it belongs to the set of red things, or does it belong to the set of red things because it is red? Objects earn their way into set membership. You are in the set of mammals because you qualify: you are a mammal. Red objects qualify for membership in the set of red things because they are red. Perhaps it is possible to replace *talk* of properties with talk of sets, but this is ontologically unilluminating. (I shall have more to say about these issues in chapter 8.)

Although technical difficulties crop up in attempts to replace talk of properties with talk of sets (Goodman 1966, chap. 5, § 3), the mistake in such attempts is ontological, not technical. We perceive objects' properties—we see the tomato's redness and see and feel its sphericality. Given that perception includes a causal component,

perceptual awareness of these properties would seem to require their playing a causal role. Objects behave or would behave in various ways *because* they have the properties they have. The tomato would roll because it is spherical, not because it is red, and certainly not because it belongs to the set of red—or, for that matter, spherical—things. Properties are central to the universe's causal order in a way sets could not possibly be.

The idea that we should dispense with properties in the interest of parsimony reflects a confusion. Parsimony is, at best a tie-breaker, a consideration that has a role only in the endgame. To wield parsimony as a prior constraint on theorizing is to impose a crippling handicap on all but the most ontologically timid.

2.4 Simple Substances

I have suggested that the universe, whatever its ultimate nature as revealed by science, includes substances and properties. Properties are ways substances are. Substances are not congeries of properties; properties are not parts of, do not add up to, substances; substances are not made up of their properties. Nor are substances 'bare particulars', propertyless substrata that take on properties to yield objects or states of affairs. Substances and properties alike are abstract entities, abstract in the traditional sense. Abstraction, Locke's partial consideration, enables us to consider the tomato or its properties. You can have distinct—'separate'—thoughts about such things even though properties could not be prised apart from their bearers, and substances could not survive without bearing properties, without being propertied. Substances can gain or lose properties: the tomato can cease to be green, but only by coming to be some other colour (or by being obliterated).

I like to think that all this would strike most readers as plausible, even obvious. What I have to say next, however, is much more likely to meet with resistance. Substances must be simple. A substance is simple if it lacks substantial parts, parts that are themselves substances. If something could be said to be a part of itself, then the idea is that simple substances lack *proper* substantial parts. (This qualification is henceforth assumed.)

Substantial parts are to be distinguished from properties, on the one hand, and, on the other hand, from spatial or temporal parts. A simple substance need not be an extensionless point. A simple substance could have a right and left half, a top and a bottom. A simple substance, a Democratean atom, for instance, might be shaped like the letter 'Q'. The universe as a whole might be, as Spinoza (*Ethics*, I, 12–13) argued it must be, a single, simple, spatially extended substance, the *only* substance. If you think objects have temporal parts, that time is space-like, that objects extend through time as well as space, a simple substance could be temporally, as well as spatially, extended. An ordinary object, a tomato, say, is apparently made up of many substantial parts. But the tomato is not made up of its spatial or temporal parts. Substances are *mereologically* simple.

If substances are simple, then tomatoes are not substances. Indeed, most of the objects we talk about, manipulate, and investigate scientifically are not substances. In saying that tomatoes are not substances, I am not suggesting that tomatoes are wispy constructs, mind-dependent entities, nor am I suggesting that talk of tomatoes is false or misleading. I am not a 'nihilist', not an 'eliminativist', about tomatoes. To get a feel for what I am saying, think of Locke. For Locke the only genuine material substances are the corpuscles, Newtonian atoms. Tomatoes are particular dynamic, interrelated arrangements of corpuscles. It is true to say that there are tomatoes, that this tomato is red and spherical, but the truthmaker for this claim is not a substance, it is a fleeting, dynamic arrangement of substances, a particular way the substances—the corpuscles—are interactively arranged at a particular time.

Descartes provides an instructive contrast. For Descartes, there is but a single extended substance: space itself. On such a view, talk of tomatoes is made true by local, textured tomatoish thickenings of space. The tomato is not a substance, nor is it made up of substances. The tomato is a mode, a way the one extended substance is.

I mention Locke and Descartes here merely to dispel one thought you might have had when I announced that substances must be simple. You might have taken me to be calling for the abolition of tomatoes, planets, people, trees, and molecules. That is not my aim. The point, rather, is that configured simples or perhaps a configuration

of the big simple, provide truthmakers enough for claims about such things.

But then why insist that substances are simple? What is the point?

If you begin, as I have begun, with the idea that substances are, fundamentally, property bearers, and properties are ways substances are, it is none too easy to see how a substance *could* be complex. Consider the tomato's sphericality. This appears to be a property of the tomato, a way the tomato is. Indeed, in explicating the complementary notions of substance and property I appealed to this very example. But I am now suggesting that this is simply a loose, provisional way of talking, a way of thinking about the universe that is perfectly acceptable until you start doing serious ontology. It is true that the tomato is spherical, but what makes this true is not the possession of a property, sphericality, by a substance, the tomato. What makes it true is a particular dynamic configuration of substances, or a particular thickening of space, or . . . well, you get the idea.

Suppose you wanted to insist that sphericality was a genuine property possessed by a complex entity, the tomato. How might this work? What exactly possesses sphericality? When you attend closely to the tomato at any moment, you see that it appears to consist of endless tiny parts in a jiggly, dynamic arrangement. If you thought that sphericality belongs to this arrangement in the way that the mass of an electron, or its charge, belongs to the electron, the belonging in question would be different in kind.[4]

Let me try to give you a feel for what I see as the problem. Consider a simple substance, s_1, that possesses a property, F. You could represent this by means of a diagram (figure 2.1). Do not be misled by the diagram. F is meant to represent a way s_1 is, not a detachable element only externally related to s_1.

Figure 2.1

Now, imagine a complex object consisting of a particular arrangement of three simple substances, s_1, s_2, and s_3 (figure 2.2). Here s_1, s_2, and s_3 bear properties F, G, and H, respectively, and stand in particular relations to one another. This complex is not itself a substance, not itself a property bearer.
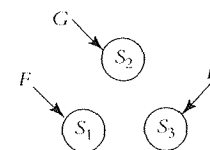
Figure 2.2

Pretend that the depicted complex is a tomato, and that the tomato is red. Is the tomato's being red a matter of the complex's possessing a property? Call the property in question, R. R is meant to be a property, not of s_1, s_2, or s_3, but of the complex comprising s_1, s_2, and s_3 (figure 2.3). It is hard to see this complex as belonging to the right ontological category. Ascribing a property to *it*, to the complex, as opposed to the substances making up the complex, has the aura of a category mistake.
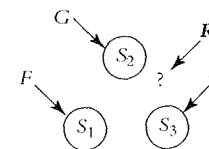
Figure 2.3

The thought here is that the sense in which a simple substance bears a property—the substance is some way or other—does not extend univocally to complexes made up of substances standing in particular relations. You could say that a complex is a particular way, the complex making up a tomato is the red way, but in that case the 'way' is the way these constituents are organized. The tomato's 'being the red way' just *is* a matter of the tomato's constituents being as they are, organized and interrelated as they are.

The same line of reasoning extends to Armstrong's states of affairs and to events, considered as the possession of a property by a substance at a time (see, e.g., Kim 1976). Suppose this is what events are, and suppose you have an event, s's possessing F at t—a particular

tomato's being red on Tuesday, for instance. You could represent this by means of another diagram (figure 2.4).

Figure 2.4

Now, it is common for philosophers to speak of events' having properties (see, for instance, Macdonald and Macdonald 2006). In fact, Donald Davidson is often credited with the view that every mental event is a physical event, where a mental event is said to be an event possessing a mental property and a physical event is one possessing a physical property (see Davidson 1970). I shall have more to say about this topic in chapter 9. For the moment, I mean only to call attention to an apparent difficulty for the idea that an *event*, regarded as a substance's possessing a property at a time—this tomato's being red on Tuesday— could possess a property. If a substance is a constituent of an event, it is easy to see how that *substance* could possess a property. But then, this is not the *event's* possessing the property. Consider figure 2.5.
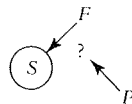
Figure 2.5

Here, an event, *s*'s possessing *F* at *t*, is itself taken to possess a property, *P*. But to what does *P* belong? Not to *s*; that would merely yield another event, *s*'s possessing *P* at *t*. No, *s*'s-possessing-*F*-at-*t* must possess *P*. But what could this *mean*? It, too, has the smell of a category mistake. Events, thus characterized, appear not to be apt property bearers.

These comments assume a corpuscular universe in which the substances are particles. Suppose, however, the universe were Cartesian: what we regard as objects are in fact thickenings in space. The tomato, then, would be a dynamic thickening of space, a mode, a property of space, a way space itself is, a wavelike entity. The tomato's characteristics—its shape and colour, for instance—would not be modes or properties of this thickening, but just aspects of the thickening itself. A red, spherical thickening differs from a green

spherical thickening, not by possessing different properties, but by simply being different in pertinent ways, by being a different kind of thickening. There are properties here, properties of the one extended substance, space itself, ways that substance is. Talk of objects' proper- ties—where objects are thickenings—is just to talk of the thickenings themselves. Objects, thickenings, are ways, and ways *ways* are, are just the ways themselves.

On a Cartesian conception, objects *in* space are not substances, they are modes, ways the one extended substance, space itself, is. On a corpuscularian conception, complex objects are modes as well, ways the corpuscles are arranged, apparently *relational* modes. Modes them- selves cannot have modes. A way a mode is would be just the mode itself.

I shall return to the topic of relations in chapter 7. For the moment, let me call attention to an ontological peculiarity of relations. A 'monadic', 'nonrelational' property, is a way some substance is. A relation, even a reflexive relation, a relation something bears to itself, is not a way a substance is. It would seem then that 'way', used to designate properties such as sphericality or charge, means some- thing different than it means when it designates arrangements.

I believe the point I am pressing here—that substances lack sub- stantial parts—was widely appreciated by medieval philosophers (see, e.g., Stump 2005, 39–44) following Aristotle (*Metaphysics Z* 16, 1040b5–16) and carried on into the early modern period and beyond. One reason, although certainly not the only reason, these philosophers might have held that living organisms possess no substantial parts is that they wanted living organisms to possess distinctive properties and this required apt property bearers. If something could not bear properties, it could not be a substance.

2.5 Properties and Properties by Courtesy

I do not expect this historical association to sway many readers, but before you flatly dismiss the idea that substances must be simple if they are to be property bearers, it is important to see what I am *not* arguing. I am not arguing that tomatoes are not red or spherical. I am not denying that there are indefinitely many *truths* pertaining to complex entities, including states of affairs and events. What I *am*

suggesting is that truthmakers for these claims are not properties, or at any rate not properties of the complex entities. If you arrange s_1, s_2, and s_3 (propertied as they are) in the right way, the result is something that answers to 'R'. Call the complex s. Then it is *true* that s is R. R-hood survives, but not as a property. R-hood is nothing in addition to this kind of arrangement of simples of these kinds; R-hood is, as Armstrong would say, 'no addition of being'.

The position defended here is sometimes described as one in which properties are 'sparse'. Note, however, it would be misleading to put this by saying that there are two kinds of property, 'sparse' and 'abundant'. 'Abundant' properties are not properties, and certainly not *kinds* of property.

Have I hijacked the term 'property', saddling it with a technical, proprietary sense it lacks in everyday use? There is, I admit, a relaxed use of 'property' according to which *any* truth about an object of the form 'a is ϕ' ascribes a property, ϕ, to a substance, a. But such everyday practices carry with them the seeds of an ontologically serious notion of substance and property. You acquire a feel for these categories in the course of interacting with ordinary objects: tables, trees, planets, tomatoes. When you engage in ontological reflection, you refine your understanding of things to which these categories apply. In taking a tomato to be a possessor of properties and taking the tomato's redness and sphericality to be properties, you are regarding the tomato as a unified substance, a one.

Once you move away from this view of tomatoes, you are in a position to recognize that tomatoes, really, are substances only 'by courtesy', the tomato's redness and sphericality are not genuine properties at all. In Keith Campbell's useful terminology, tomatoes are '*quasi*-substances', and tomatoes' colours and shapes, '*quasi*-properties' (Campbell 1990). The tomato's redness and sphericality are what you get when you arrange the tomato's constituents in the right way. By 'right way', I mean something more than a tomato-shaped arrangement. A wax tomato or a cloud resembling a tomato, would be a tomato-shaped arrangement of particles, but they are not tomatoes. To get a tomato, you need the right kinds of particle interacting in the right ways with one another and with the environment.

There are, then, no complex substances, no complex properties. What you might informally regard as properties of complexes are in

fact results of arrangements of propertied simples. Is this 'reductionism' run amok? I am not confident I understand what reduction amounts to, but whatever it is, reduction is not a relation among entities, not an ontological relation. What would it be to reduce one property, or state, or object to another? (The image here is of a philosopher in a white lab coat wedging entities into a vice, and turning the handle so as to squeeze them together.) Reduction evidently involves theories, categories, taxonomies. You do not reduce one thing, water, to another thing, H_2O. You discover that 'water' and 'H_2O' have a common denotation: what we call water *is* H_2O. Truthmakers for assertions about water turn out to be portions of H_2O; the deep story about water is that water is H_2O.

Regarded as a conceptual or analytical endeavour, the prospects of reduction are not promising. There is scant hope of replacing talk of complex objects with talk of arrangements of interrelated simples. This could be so even if truthmakers for judgements about complexes inevitably turned out to be arrangements of interrelated simples. You can, then, accept the thesis that substances must be simple and its corollary—the only properties are properties of simples—without calling into question scientific accounts of the universe, including accounts emanating from the 'special sciences', or, for that matter, ordinary assessments of ordinary things.

The task of ontology is to spell out the fundamental categories of being. The suggestion thus far is that these must include substance and property. Fundamental physics is in the business of telling us, among other things, what the substances and properties *are*. The universe as revealed by fundamental physics could turn out to be surprising indeed. The universe might be continuous and non-particulate, for instance, the universe might resemble Spinoza's universe. But if this is how it is, then you have, not the elimination of substance (or 'things'; see Ladyman et al., 2007) but a single substance: space–time, or the quantum field, or the cosmos, the One. What you might naturally think of as objects would be modes of this substance, ways it is. It could turn out that the universe is an arrangement of interpenetrating fields. In that case, fields would play the substance role and objects—electrons, tables, planets—would turn out to be fluctuations or disturbances in fields, particular ways the fields are.

The idea that the fundamental substances might be fields calls to mind the Cartesian conception of substance. Descartes holds that every substance is characterized by a distinctive, determinable attribute. Space, for instance, the one material substance, has the attribute of extension. What we regard as distinct bodies 'in' space are, in reality, determinate modes of extension, ways of being extended, ways space is. You could think of fields as having distinctive characteristic attributes, and their modes as being determinate ways of being thus attributed.

My conception of substance and property brings with it a host of arresting implications. To be sure, you might regard these implications as reasons enough to reject the conception. In that case, I invite you to offer a replacement. If you thought, for instance, that complex objects are themselves substances, bearers of properties, you would need to provide an account of the 'bearing' relation that encompasses simple and complex substances alike. And this, as many philosophers have recognized, is easy to say, hard to do.

2.6 Emergent Substances

One implication of the conception of substance and property defended here is that it opens the way to a plausible, if surprising, account of property emergence. The idea that conscious states of mind are emergent features of the universe is currently enjoying renewed popularity. Philosophers have been hard pressed to show how consciousness could be spawned by wholly non-conscious physical processes. You can see how the sphericity, or even the redness, of a tomato might result from particular arrangements of particles. But to many philosophers and neuroscientists it has seemed completely mysterious, indeed magical, that an arrangement of neurons could result in a conscious state. The apparent mystery has led philosophers and some neuroscientists to suggest that consciousness, unlike redness, or sphericity, or mass, or charge, is emergent.[5]

5 For a progression of views on emergence, see O'Connor 1994; O'Connor and Jacobs 2003; O'Connor and Wong 2005. See also Kim 1999; and, for an account of the history of the concept of emergence, McLaughlin 1992. See Macdonald and Macdonald (2010) for recent discussions of emergence.

Talk of emergence is notoriously slippery. When something—a property, for instance—is said to emerge, what exactly is being claimed? Emergence is often spelled out in terms of explanation: when the explanation for some phenomenon, Π, is not deducible from, or predictable on the basis of, various 'lower-level' phenomena, Π is said to be emergent relative to those phenomena. Philosophical proponents of emergence, however, take themselves to be defending an ontological thesis, not, or not merely, a thesis about prediction or explanation.[6] It is not merely that emergent phenomena are, given our current state of knowledge, novel or surprising, but that they are genuine 'additions of being'. The tomato is spherical, although none of its parts are. Yet the tomato's sphericity is not really emergent, not an addition of being in the way your conscious experiences might appear to be. If you arrange the tomato's constituents in the right way, you *thereby* have a red, spherical tomato. This is straightforward. Appropriate arrangements of neurons yield consciousness, but *how*? The connection between populations of neurons and conscious experience appears completely arbitrary. Colin McGinn, for instance, despairs of our ever coming to understand 'how Technicolor phenomenology could arise from grey soggy matter' (McGinn 1989, 349).

O'Connor and Wong (2005) make this point salient, borrowing, from Armstrong, the notion of a 'structural property'.

(S) A property, S, is structural if and only if proper parts of particulars having S have properties not identical with S and jointly stand in relation R, and this state of affairs *is* the particular's having S. (663)

I have denied that complex entities could bear properties, so I would deny that S is a property: structural properties are quasi-properties, properties only 'by courtesy.' S *is* a truthmaker for '*x* is S', however, and this is all you need to appreciate O'Connor and Wong's point. That point, as I should put it, is that genuinely emergent traits *are* properties, properties that are 'wholly nonstructural' (663). All physical properties, indeed all properties other than those characteristic of

6 The lines between explanatory concerns and ontology are easily blurred. Beware arguments for emergence that move from premises concerning the irreducibility of mental concepts to physical concepts to the conclusion that there is no basis in the physical for the mental. See Chalmers 1996.

conscious experiences, are, they contend, either fundamental physical properties or structural properties-by-courtesy resulting from combinations of fundamental things. Only in the case of consciousness do O'Connor and Wong think we need non-structural, emergent properties. These properties are fundamental. They stand alongside whatever other properties might be discovered in fundamental physics.

O'Connor and Wong offer an interesting account of the metaphysics of emergence: emergent states are *caused by* configurations of more fundamental items. Structural properties-by-courtesy, in contrast, are not *caused by* configurations, they *are* the configurations. I shall ignore the details of this view, although I shall return to these issues in chapter 11. My aim at present is simply to note that any conception of emergence is incomplete without an account of the *bearers* of emergent properties. Suppose I am right in thinking that complex entities are unfit to bear properties. Fundamental properties require, as bearers, fundamental substances. But it is hard to see complex biological systems as fundamental substances. In fact, their being complex, their being made up of substantial parts, disqualifies them as *substances*.

This point is partly addressed in an earlier paper, 'Emergent Individuals' (O'Connor and Jacobs 2003). There O'Connor and Jacobs argue that robust ontological emergence requires the emergence of fundamental *substances* as bearers of emergent properties. They are thinking of the emergent substances as persons or agents, causally dependent on, but ontologically distinct from, complex physical systems.

If my account of substances and properties is on the right track, this is the kind of move anyone who posits genuinely emergent properties is going to have to make. The move is, arguably, prefigured in Aristotle and Aquinas, who held that human beings—and organisms, generally—were causally dependent on physical states and processes, but in no sense identifiable with those states and processes. Aristotle and Aquinas regarded emergent substances as simple, lacking in substantial parts. More recently, this idea has been developed by E. J. Lowe (1996, 2008). One motive for such a view, a motive that has played a significant role in this chapter, is that it looks as though only a simple substance *could* be a property bearer. O'Connor and Jacobs, in contrast, allow their emergent individuals to be complex.

By my lights, O'Connor and Jacobs have made the hard move, but fallen short of the goal. They do not question the possibility that complexes might bear properties. Thus, their structural properties are *properties*. They recognize, however, that there is a problem about what exactly might bear an emergent property, something that appears not to have occurred to other proponents of emergence. If you are going to have emergent fundamental properties, you are going to need emergent fundamental substances as bearers of those properties. Fundamental substances, however, genuine substances, are simple, lacking in substantial parts.

Might all this be too much to hope for? Might the very idea of emergence be incoherent? Galen Strawson argues that, if emergence were a 'brute' phenomenon, if there really were *no* accounting for what it is about 'grey soggy matter' that facilitates consciousness, the notion of emergence would lack credibility.

If it really is true that Y is emergent from X then it must be the case that Y is in some sense wholly dependent on X and X alone, so that all features of Y trace intelligibly back to X (where 'intelligible' is a metaphysical rather than an epistemic notion). Emergence can't be brute. It is built into the heart of the notion of emergence that emergence cannot be brute in the sense of there being absolutely no reason in the nature of things why the emerging thing is as it is (so that it is unintelligible even to God). For any feature Y of anything that is correctly considered to be emergent from X, there must be something about X and X alone in virtue of which Y emerges, and which is sufficient for Y. (Strawson 2006, 18)

Strawson concedes that features of the fundamental things could be 'brute', but, given those features, the fact that fundamental things of one kind, the X's, could, under the right circumstances, give rise to further fundamental things, the Y's, must be written into the nature of the X's. If contemporary conceptions of emergence deny this, these conceptions are without merit.

Strawson could well be right about this. I am confident he *is* right. Nevertheless, I believe that there are cases of perfectly natural phenomena that satisfy what emergentists say they want in a concept of emergence. The cases I have in mind involve the production of some 'new' fundamental thing that could be accounted a genuine addition of being.

2.7 Real Emergence

My interest in the previous section was to clarify what the kind of emergence envisioned by philosophers who write on the topic would require. Emergent properties must be fundamental and so must their bearers. But now the way is open to see emergence as a straightforward, uncontroversial *natural* phenomenon. (I discuss emergence in quantum physics in § 3.7.)

Focus on the fundamental things, or what physicists currently regard as the fundamental things. Do these ever emerge? Consider an imaginary case in which a new kind of particle is produced in a collider.[7] When an α-particle encounters a β-particle, the upshot is the annihilation of the α- and β-particles, and the creation of a new kind of particle, a χ-particle, possessing properties emergent with respect to α- and β-particles. This is genuine, for-real, honest emergence! The set-up required to produce the particle is complex; it includes the collider and a host of supporting mechanisms. But the χ-particle does not mysteriously inform this complex, its properties are not properties of the complex. The particle, propertied as it is, emerges from a collision between α- and β-particles facilitated by the complex.

It would seem that χ-particles satisfy the hankering for entities 'over and above' those that go into their production. A χ-particle is not an arrangement of α- and β-particles, a χ-particle is a genuine addition of being. True, it is of the nature of α- and β-particles that, when they encounter one another at high velocities, they annihilate and produce χ-particles. But, as Strawson notes (and O'Connor, Jacobs, and Wong would agree), on any account of emergence, it had better be of the nature of the fundamental things that, under the right circumstances, they would yield a new kind of fundamental thing.

Now we are in a position to see clearly what the emergence of consciousness or conscious properties would require. It would need to be the case that, in the right circumstances, new fundamental kinds

7 Emergence of the sort described here occurred on a massive scale during the Big Bang and is a familiar feature of nuclear reactions.

of substance are produced via interactions of other fundamental kinds of substance. If emergentists about consciousness are right, then the brains of human beings (and undoubtedly the brains of many other species) might provide the right circumstances.

If you have emergentist sympathies, a view of this kind has at least three advantages.

First, it affords an account of emergence continuous with what is already present in fundamental physics. Emergence is not something that occurs exclusively in the nervous systems of sentient creatures. Appeals to emergence need not be regarded as dodgy, ad hoc manoeuvers employed solely to save the appearances.

Second, it offers an account of conscious properties that makes such properties fundamental in a way that does not imply panpsychism. You need not imagine that electrons and protons are dimly conscious, any more than you need to suppose that properties of emergent χ-particles are already present in α- and β-particles, or that α- and β-particles are 'proto-χ'-particles.

Third, the view fits nicely with the account of substance and property defended here, an account I hope you will eventually come to regard as inevitable. Emergent properties are properties of emergent, simple substances.

A caveat. Although it is clear that there *are* emergent substances and properties, it is much less clear that conscious qualities are emergent. Proponents of emergence are sure they must be; I am not. My aim at present is not to defend the thesis that consciousness is an emergent phenomenon, however, but merely to make clear what this might *mean*, what would have to be the case *were* consciousness genuinely emergent. I shall consider the status of consciousness in more detail in chapter 11.

2.8 Back to Basics

One lesson to be taken from this discussion is that it would be a mistake to imagine that you can get clear on a phenomenon such as emergence without first getting clear on the fundamental ontology of property and substance. If you begin by assuming that properties and property bearers can be complex—a substantive ontological

thesis—you risk losing your way forever in the metaphysical maze. Well, if not losing your way, you risk at least embracing a host of unexamined and arguably pernicious ontological theses. Their perniciousness stems from their effects on subsequent theorizing, effects themselves stemming from an ontological picture that threatens to mislead in a systematic way.

Of course, what I have labelled 'pernicious ontological theses' could turn out to be true, the ontology that I propose here could be utterly misguided. You will never know until you see how the theses stack up in a more encompassing ontological scheme. Thus far I have barely touched on a number of points that will receive much more attention in subsequent chapters. I have left numerous loose threads that will need to be knitted together. My hope is that the overall tapestry will bring illumination to some of the darker corners of metaphysics.

CHAPTER 3
Substance Stressed

The infinite turns out to be the contrary of what it is said to be. It is not what has nothing outside it that is infinite, but what always has something outside it. . . . A quantity is infinite if it is such that we can always take a part outside what has already been taken.

(Aristotle, *Physics* III, 206b 33–207a 11)

We have . . . been compelled to dismiss the idea that . . . a particle is an individual entity which retains its 'sameness' forever. Quite the contrary, we are now obliged to assert that the ultimate constituents of matter have no 'sameness' at all.

(Schrödinger 1951, 17)

I think I can safely say that nobody understands quantum mechanics.

(Feynman 1965, 129)

3.1 Ontological Dependence

Substances are property bearers. Substance and property are complementary categories. Properties are ways substances are, and every substance is some way or other. This is one way to introduce substances. Another tradition begins with a notion of independence: a substance is a non-dependent entity, an entity the existence of which does not require the existence of any other distinct entity (see, e.g., Campbell 1976, 1990; Robb 2009). Substances are basic entities, basic particulars; properties, in contrast, are dependent. Properties depend on the substances that bear them. If, like me, you regard properties as modes, every property would depend on its bearer: Socrates's whiteness depends on Socrates. This is what E. J. Lowe calls 'rigid existential dependence'. If you think of

properties as universals and you hold, with Armstrong and Lowe, that universals must have instances, then a property depends, not on a particular substance, but on some substance or other, what Lowe calls 'non-rigid existential dependence' (Lowe 2006 § 3.1).

One result of characterizing substances in this way is that complex objects, objects with substantial parts, will not count as substances. At least they will not count as substances if wholes depend ontologically on their parts. Complex wholes are *made up of* their parts. A whole— or some wholes—might be thought to survive gradual *replacement* of its parts, perhaps, but not their elimination. Your automobile could survive your gradually replacing each of its parts, but removal of parts without replacement results in the automobile's demise.

Speaking strictly, a whole is a definite collection of parts. Adding or subtracting a part results in a new whole. We are rarely interested in such wholes, however. Rather we deploy terms—'sortals'—that include much more lenient identity conditions. Part of learning what it is for something to count as an automobile, for instance, includes learning informal limits on the addition, subtraction, and rearrangement of parts. Your automobile today differs from your automobile last week: you have painted the roof, replaced the muffler, rotated the tires, removed the cigarette lighter, and added a set of roo bars. Nevertheless the wholes count as the same automobile.

This is not deep ontology, not something I intend to dwell on further. For the present I merely note that the idea that substances are non-dependent entities, in concert with the idea that wholes depend on their parts, provides independent support for the thesis, advanced in chapter 2, that substances must be simple. Whether you begin with a characterization of substances as property bearers or as non-dependent entities, it appears that substances could not be made up of other substances, substances could not be 'mereologically complex'.

In this context it is vital to bear in mind a distinction mentioned earlier between substantial parts, on the one hand, and, on the other hand, spatial or temporal parts. The idea was that a substantial part of a complex object is a part that is itself a substance. Although a spatial part could *coincide* with a substantial part, a simple substance could be spatially extended, a simple could have endless spatial parts. In fact, anything that has spatial parts must, it seems, have an infinity of spatial parts. (Unless space is 'granular': some spatial parts lack spatial parts.

Although this could be a genuine possibility, I shall ignore it here in the interests of simplifying the discussion.) In asking whether an object is divisible, then, it is important to make it clear whether the divisibility in question is substantial or spatial.

Spatial and substantial parts exhibit an important asymmetry. Although it seems right to think of a complex object as depending ontologically on its substantial parts, the opposite is so for an object's spatial parts. An object is not *made up* of its spatial parts: spatial parts are regions *of some object*. If objects persisting through time have temporal parts, these resemble spatial parts, not substantial parts.

3.2 Spatial and Temporal Parts

Complex objects, wholes, are made up of, and so depend on, their substantial parts. In the case of spatial parts, the order is reversed. Spatial parts depend on the objects of which they are parts; objects are not made up of their spatial parts. Complex objects of interest to us typically depend, not merely on the parts that make them up, but also on those parts being arranged and causally interacting in the right ways. Your body is not simply a collection of particles, but a dynamic, evolving collection of particles of the right kinds interacting in the right ways with one another and with particles nearby.

So an object's substantial parts can causally interact with one another. What of its spatial parts? Could spatial parts of a body causally interact? To think that they might is to confuse occupants of spatial regions for the regions themselves. The particles making up the right half of a tomato causally interact with particles making up its left half. But you do not have causal relations among the particles *together with* causal relations among the regions. One impediment to appreciating this point stems from the fact that it is easy to refer to the occupants of a spatial region by mentioning the region: southeastern Missouri is being buffeted by rain and high winds; the table's right half is on fire. But in so doing you are not referring to spatial parts, spatial regions, but to those regions' substantial occupants.

What if space were substantival? If, as Descartes thought, space itself is an extended substance and objects are thickenings of this substance, modes, ways space is, what we currently regard as causal

relations among objects would in fact be the propagating of modes in space, a kind of wave motion. But spatial modes, thickenings of space, are neither *occupants* of space nor substantial *parts* of space, parts that make up the whole. Whatever the causal story here, causal relations would be relations involving ways spatial regions are, not relations among the regions themselves.

What of *temporal* parts? Temporal parts enter the picture when you think of time as 'space-like'. A physicist might consider an object's persisting over time as a temporally extended thing, a space–time worm. If time were space-like, then objects would have, in addition to spatial parts, temporal parts. Were an object complex, it would have substantial parts as well.

The thesis that objects have temporal parts has come to be associated with a metaphysical doctrine discussed by David Lewis, according to which objects 'perdure' through time by having temporal parts (see Lewis 1986b, 202–5). The you of two years ago is not strictly identical with you today. The you of two years ago had properties you now lack and bore relations to endless other things you no longer bear. Strict identity requires sameness of properties and relations: A and B are one and the same, A and B are strictly identical, only if A has a property or stands in a relation if and only if B does.

Were you today and you two years ago distinct temporal parts of a single, temporally extended entity, it would be possible for the selfsame you to be happy and sad. Yesterday you were happy, today you are sad. Your being in these incompatible states would be analogous to a table's being both red and green by virtue of being red on one side and green on the other, or the Sydney Harbour bridge's being in both Kirribilli and Millers Point. A single object could possess apparently incompatible characteristics and stand in incompatible relations, provided the object did so by virtue of its distinct spatial or temporal parts (more accurately, the propertied substances coinciding with those parts) possessing the characteristics or standing in the relations. Lewis distinguishes this doctrine from conceptions of persistence according to which objects 'endure': the you of two years ago and the you of today *are* somehow strictly identical despite differing in any number of ways.

I am not concerned to defend the notion of temporal parts, although I suspect that, rightly understood, it is compatible with

assorted conceptions of time (see McCall and Lowe 2006). I intend only to consider the question whether it would make sense to imagine that causal relations could hold among temporal parts. The question arises because some philosophers have thought that it might be possible to account for objects' persistence through time by supposing that earlier temporal stages causally support later stages (see, e.g., Armstrong 1997, 73–4). This would not be a matter of the obtaining of causal relations among an object's substantial parts over time, but of causal relations between the object's successive temporal parts or stages themselves.

If temporal parts are analogous to spatial parts, then the idea that they could stand in causal relations to one another is fundamentally misguided. Temporal parts of an object extended in time are not like pearls in a necklace, each pearl encompassing a temporal slice of the temporally extended object. Temporal parts are not like dominoes standing in a row so that the toppling of one brings about the toppling of the adjacent domino. Temporal parts are akin to the spatial parts of a tomato or an electron. It would be hard to see how a tomato could survive unless its substantial parts causally interacted in the right ways (and interacted as well with external, supporting factors). But you do not have the interaction of these *plus* the interaction of the tomato's spatial parts. (See § 6.2 below.)

Similarly, if the tomato had temporal parts, if the tomato were temporally extended, it would hold together over time, at least in part, owing to causal relations among its substantial parts (*and* various relations those parts bear to states of the environment). The state of the tomato at one time might be causally responsible, in whole or in part, for its state at a later time, but this is not a matter of a temporal part of the tomato causally influencing a subsequent temporal part. The causal relations in question are just those that would be studied in molecular chemistry.

3.3 Complexity 'All the Way Down'

Bearing in mind a distinction between spatial (or temporal) and substantial parts, consider now the possibility that *every* apparent substance is infinitely complex. Every candidate substance has parts,

and each of its parts has parts. These parts are substantial parts—but not really. A substantial part is a part that is itself a substance. But if a part *itself* has substantial parts, it is not a substance. Were *everything* divisible this way, were everything *endlessly* divisible, there would, it seems, be no substances, no simple, non-dependent entities. As D. C. Williams puts it,

The universe might have been the same size that it is, and included exactly the same number of individuals, variously discrete, overlapping, and included, in the same way, and yet have been perfectly homogeneous throughout, one great blob of blanc-mange, say (i.e., of stuff which really is the way blanc-mange seems to be). (Williams 1959, 3–4)

In David Lewis's more pedestrian phrase, the universe might have consisted of 'atomless gunk' (Lewis 1991, 20–1).

The universe's being infinitely divided would seem to be a live possibility. But a commitment to substances—in the sense of ontologically non-dependent entities—is apparently at odds with this possibility. If wholes depend upon their parts, *everything* could not be a complex whole made up of parts. Ontological dependence of wholes upon parts apparently requires non-dependent entities, in this case substances, simples.[1]

Although the idea that wholes depend upon their parts, coupled with the idea that dependence of this kind must 'bottom out', implies the existence of simple substances, you might regard this as objectionable if you thought that an infinitely divisible universe were a genuine possibility. Jonathan Schaffer (2010) argues that the possibility that the universe is 'gunky' supports 'priority monism', the view that the universe is a single unified whole on which its parts depend, thereby reversing the order of dependency. If substances are characterized as non-dependent entities, then there is just one substance, the universe as a whole, the One. In that case, 'parts' of this substance would not be what they could not be: substantial parts.

This might or might not be what Schaffer has in mind, but it *is* close to Spinoza's monism. What are ordinarily considered to be self-standing substantial parts of the universe (electrons, pebbles, trees,

galaxies) would be, not substances, but modes, ways the universe is. This is how Descartes thinks of extended bodies. For Descartes, there is but one extended substance, space itself. What are commonly regarded as material substances are, in reality, modes, determinate ways space itself is. You might think of 'gunky', blancmange universes in this way. Such universes are not infinitely divided into substantial or *quasi*-substantial parts themselves infinitely divisible. Blancmange universes are effectively seamless. Blancmange universes are themselves simple substances. Their complexity is 'non-mereological', their 'parts' are not genuinely parts, not *substantial* parts, they are no more substantial parts than the left and right halves of a tomato are substantial parts of the tomato. A tomato's left and right half might *coincide* with substantial parts of the tomato, but that is another matter.

If you begin with the idea that substances are non-dependent entities, you seem to be faced with a choice between (a) a denial of the possibility that the universe could be blancmange-like; and (b) regarding blancmange universes as simples. What if you begin with a characterization of substances as simple property bearers? A blancmange universe might seem to be a universe with no substances, hence no properties! Now the problem would be to find truthmakers for truths holding in such a universe. If there are simple substances, then we can have truthmakers for 'The tomato is red and spherical', a particular dynamic arrangement of simple substances, for instance, or, if there is but one substance, a particular way that substance is.

What is emerging here is the idea that it is a mistake to think of a blancmange universe as a universe with literally endless parts. The parts of such universes would not be substantial parts. They would resemble—no, they would *be*—spatial parts. A blancmange universe would be an extended simple substance, a One.

You might be suspicious of this line of reasoning on the grounds that it apparently assumes without warrant that the universe is finite. We started with the universe and imagined it divided. But suppose you began, so to speak, at the other end with simple substances, the fundamental particles, say, and imagined these multiplied infinitely. Although such a universe would be a universe with an infinite number of parts, these would 'bottom out' in simple substances, substances themselves lacking substantial parts.

---

[1] The dependence problem does not stem from infinite divisibility alone, but from the lack of a 'ground floor' of non-dependent simples. Cameron (2008) disagrees, arguing that, although an infinitely divided universe is possible, there are good reasons to think that the actual universe includes simples.

I take up this possibility in § 3.5. Meanwhile, it is important to see that blancmange, 'gunky' universes are not just universes with an endless number of parts. Rather, a blancmange universe is a universe with *no* simple parts, no parts themselves lacking parts.

3.4 Infinite Divisibility

The possibility, then, that the universe is 'infinitely divided', where the divisibility in question is substantial divisibility, not spatial divisibility, is unpromising. You might be tempted to think of such a universe as one in which every object has parts that are themselves objects with objects as parts. Such a universe would contain no simples, hence, by my lights, no substances and, given that properties are ways substances are, no properties. Would such a universe be genuinely conceivable? Parts of an object make up the object. But it is hard to see how anything, including the universe as a whole, could be made up of parts that are themselves made up of parts, that are themselves made up of parts that are. . . . The difficulty here is not simply a failure of imagination, however.

Talk of 'infinite divisibility' is cheap. A line of whatever length is infinitely divisible in the sense that any of its segments, however small, could be regarded as having further spatial divisions. These segments do not *make up* the line in the way grains of sand make up a sand dune, however. The sand dune depends on individual grains of sand that make it up. In an important sense, however, a line's segments depend on the line rather than it on them.

The notion of infinity deserves a closer look.[2] An infinite number is not a very, very large number. An infinite series is not a very, very, *very* long series, but a series that has *no* end, *no* last member. Infinite numbers differ qualitatively, not merely quantitatively, from finite numbers, however large. For this reason, it is at least misleading to think of an 'infinitely divided' universe as a universe made up of an inordinately large number of parts. 'Infinite divisibility' is not just divisibility into tiny parts, it is divisibility *without end*.

2 A. W. Moore (2001) offers a trenchant account of the notion of infinity and its surprisingly troubled history.

A universe, 'infinitely divisible' or otherwise, might be made up of a finite number of simple substances or it might be a continuous blancmange universe lacking parts—substantial parts—altogether. The universe might contain a very large number of simples (the fundamental particles, the electrons and quarks), for instance, or only a few (the fundamental fields), or the universe might turn out to be a single unity (space–time itself, or the One). The divisions in question would not be substantial divisions, however.

How the universe is substance-wise is an empirical matter concerning which the ontological picture advanced here is meant to be neutral. Whatever conception of the universe we find in fundamental physics, that conception will include simple, propertied substances. We look to science to fill in the blanks, to tell us what the fundamental substances and properties are.

These issues are frustratingly elusive. Philosophers have grown accustomed to talk of infinite numbers and sets. Infinity is mathematically indispensable. Infinity has earned its keep. The difficulty is to understand possible applications of the notion of infinity to the universe, to understand the *ontology* of infinity. With this in mind, it might be useful to come at topics addressed in §§ 3.1–3.4 from a slightly different direction.

3.5 An Infinity of Substances

Could the universe contain an infinite number of substances? Could there be an infinite number of electrons, for instance? You might regard this as an obvious possibility: why on earth *not*?

Consider, first, how this question is related to the possibility discussed earlier, the possibility that everything might be infinitely complex: everything might have an infinite number of proper substantial parts. Start with any object, this tomato, say. The tomato has parts—stem, seeds, pulp, skin—and each of these parts has parts—living cells. Eventually you reach the atomic level. But atoms, too, have parts: electrons, protons, neutrons. Similarly, protons and neutrons have nucleons and quarks as parts. Suppose there were no end to this kind of division, suppose every object could be decomposed into parts, and these parts into parts, and so on to infinity.

I have insisted that a sensible discussion of this topic requires distinguishing spatial (or temporal) parts from substantial parts. Geometry tells us that any finite line segment comprises an infinite number of line segments. But these segments do not make up the line, they are not elements you assemble to create a line. Any actual line is made up of physical constituents, ink droplets, for instance, which are not to be confused with purely spatial magnitudes.

The Greeks distinguished two kinds of infinity: what was infinite by addition and what was infinite by division. A Euclidean line is infinitely divisible, but it does not follow that any actual line could be made up of the addition of an infinity of elements. How, exactly, would that work? A line with infinite divisions is one in which every division could be further divided. The divisions are, quite literally, endless. Although you might be sanguine about divisibility of this kind, there is something deeply puzzling about the prospect of a finite object's being made up of an infinite number of substantial parts. If you thought that a given line were infinitely divisible, you could think of it as encompassing all the infinite segments into which it is divided. But this is not to regard the line as being made up of an infinite number of parts.

One difficulty about such cases stems from apparent constraints on composition. Wholes appear to 'depend metaphysically' on parts that make them up. Complex wholes (tomatoes, for instance) are made up of other complex wholes (seeds, pulp, skin) that are made up in turn of other complex wholes (cells, molecules), and so on. I have re-commended an ontology of substance and property, arguing that substances must be simple, substances lack parts that are themselves substances. Substances are the fundamental building blocks. Charac-teristics of complex objects result from arrangements of substances possessing various properties. If you take away the substances, if you take away the properties and the arrangements, you take away the universe.

It goes without saying that there is no bullet-proof argument for any of this. If you reject it, however, you are under obligation to provide an account of the resulting ontology: the ball is in your court. And in these cases, appeals to the fact that infinite numbers are mathematically well established is not to the point.

Once again, nothing I have said here implies that the universe is corpuscular, that the universe *is* made up of a very large number of proper parts that are themselves simple substances. This is important, because I believe that the idea that the universe might be infinitely divisible makes perfect sense provided the divisions in question are spatial or temporal divisions. In that case, the divisions do not mark off substances, the universe is not made up of, does not depend on, proper substantial parts: the parts depend on the whole (see Schaffer 2010).

An appeal to proprietary notions of substance and property—albeit ones that I believe stem naturally from reflection on the nature of things—might arouse suspicion about my conclusions. Consider, however, an independent line of reasoning that supports the idea that there really is something deeply puzzling about wholes made up of an infinity of substantial parts.

First, return to a question broached in § 3.1. Might there be an infinite number of electrons? Suppose there were. Suppose, further, that the electrons were arranged so as to form a cosmic sphere. We know that electrons have finite masses. We know, as well, that it is of the nature of an infinite set that it has a proper subset, members of which can be put into one–one correspondence with members of the whole set. In the case of the integers, for instance, the even numbers can be paired one–one with all the integers: the odds *plus* the evens. This feature of infinite sets makes it clear why it is a mistake to think of an infinite number as a very large number. Infinite numbers differ *qualitatively* from finite numbers.

The mass of all the electrons is what you get by combining the masses of individual electrons. Now consider all the even-numbered electrons. These might be the electrons making up the left hemi-sphere of the electron sphere. These electrons, a proper subset of all the electrons, stand in a one–one correspondence to all the electrons. But then it would seem that the mass of one of the hemispheres would have to equal the mass of the sphere itself, the whole that includes each hemisphere as a proper part. But wait! How could the mass of a homogeneous whole be the same as the mass of one of its proper parts?

Imagine God extracting the even-numbered electrons and arran-ging them so as to form a second sphere. This sphere would have the

same mass as the original sphere. But now God, merely by rearranging the electrons, would have doubled their mass![3]

Does it help that both masses are infinite? I do not see how. You still have a hemisphere with the same mass as the whole sphere that includes that hemisphere as a proper part. You still have a second sphere constructed from a proper part of the original sphere that has the same mass as the original. Imagine God annihilating all the odd-numbered electrons. This action would in no way diminish the number of electrons in existence, in no way diminish their mass. How strange!

Have I muddied the water in supposing that a one–one correspondence yields this result? The question might be turned around: why *shouldn't* this result follow from a definitive characteristic of infinite numbers?

3.6 The Weirdness of Actual Infinities

Perhaps, given the nature of infinity, these results are neither puzzling nor in any way objectionable. The results are not puzzling in one respect: they are just what you would expect under the circumstances. What they call into question is whether you can muster a coherent conception of composition that covers cases in which wholes are taken to be *made up of* an infinite number of proper substantial parts. Ordinary tenets of composition suggest that objects that lose proper parts suffer diminution, and that collections of an object's proper parts cannot be put into correspondence with all of the object's parts. In attempting to imagine objects with an infinity of proper substantial parts, we have, I think, lost a grip on the thought that complex objects have parts, that complex objects are made up of their substantial parts.

Some readers will dismiss such concerns. Maybe the normal rules of composition apply only to finite collections. Maybe infinite collections differ compositionally from finite collections in a way analogous to the way infinite numbers differ from finite numbers.

---

3 This possibility was suggested to me by Roy Sorensen.

Such a response betrays a lack of sensitivity to the manifest weirdness of actual infinities. One way to desensitize yourself is to start with a notion of the infinite-by-division. It is relatively easy to regard a line or a surface as being infinitely divisible into spatial segments or regions. This puts infinity on the table. That would seem to be the hard part. Now imagine an infinity by addition. Start with an electron and imagine adding electrons until you cross the threshold and have an infinite number of electrons. But, as noted earlier, an infinite number of electrons is not just a very large number of electrons, it is a literally *endless* number of electrons. Just as there is no greatest integer, if the electrons are infinite in number, there would be no definite number of them. But how could that be? In a universe containing an infinite number of electrons would mass-energy be conserved? In a universe comprising an infinite number of electrons, electrons could come and go without affecting the total mass-energy. Under the circumstances, it is hard to know what to say.

My recommendation is that we take all this to heart and reject the possibility that the universe could be made up of an infinite number of substances. You can allow an infinite number of spatial or temporal divisions, perhaps, but it would be a mistake to move from these possibilities to the possibility that the universe could be infinitely complex, the possibility that the universe could contain an infinity of distinct substances.

If this is a constraint on scientific theorizing, it is not much of a constraint. It allows for the possibility that there could be many, many particles. These particles could have many, many parts, and the parts could have many, many parts. But these possibilities do not take us down a road to an infinity of substances or to objects with an infinite number of proper parts. That would require a leap for which it is hard to imagine empirical warrant.

3.7 From the Weirdness of Infinity to Quantum Weirdness

In a universe made up of multitudinous distinct substances, truths about the universe are made true by the substances, their properties, and their interactive arrangements. Complex objects—tables, trees,

planets, atoms—would not be substances, but complexes made up of substances. What we regard as properties of complexes are just what you get when you arrange substances of these kinds in these ways. Just as complex objects are no addition of being but simply the substances duly propertied and duly arranged, so characteristics—'properties'—of complex objects are no addition of being. Truthmakers for all the truths concerning complex objects are particular dynamic arrangements of interacting simples.

This conception of substances and properties is compatible with, but does not imply, corpuscularism, the doctrine that the post-Big Bang universe is made up of arrangements of particles. The conception is compatible with the possibility that the fundamental substances, the fundamental property bearers, are fields, and with the possibility that there is just one substance, space or space–time, the One, the universe as a whole. These are matters to be settled by fundamental physics.

Let me pause here to note that I distinguish *atomism* from *corpuscularism*, a species of atomism. Atomism is the view that the universe is made up of simple substances. This is compatible with the possibility that the universe is made up of a very large number of simple substances, the corpuscles, but also with the possibility that the universe is a single simple substance, and with the possibility that there is some small number of simple substances, the fields.

What of quantum physics? What happens to particles in a quantum universe? Hard cases make bad law. Quantum physics makes bad metaphysics—bad metaphysics, not bad science. Competing 'interpretations' of quantum physics make it exceedingly unclear how the universe must be if the quantum theory is true. A number of philosophers have argued that quantum theory as it now stands is sharply at odds with corpuscularism. The phenomena of 'entanglement' and 'nonseparability' force us to accept the idea that quantum systems are 'holistic' (see Teller 1986, 1995; Healey 1991). Such systems appear to possess properties not founded on properties of their constituents. Were it to turn out that the universe is best seen as a single entangled system, then you would have something approaching the universe of Spinoza: the One.

Before discussing why physicists or philosophers might be attracted to such a position, let me say something about how it bears on the conception of substance and property advanced here.

Suppose quantum systems, systems of 'entangled' particles, are genuinely 'holistic', suppose their characteristics really do outstrip characteristics of their ingredients. In that case, it would be at least misleading to describe such systems as being *made up of* their parts—electrons, photons, quarks. The 'parts' of such systems would have the status of modes: the wholes of which they are parts would not depend on the parts, the parts would depend on the wholes.[4]

On a 'classical' model, complex systems are invariably made up of, hence dependent on, the elementary parts that make them up: electrons, protons, neutrons, and the like. These parts are simple substances. Truths about such wholes are made true by particular arrangements of, and interactions among, these substantial parts.

Quantum physics affords a different picture. Although it is convenient to speak of electrons, for instance, as particles or elementary substances, when electrons enter into relations, they can 'lose their individuality' (see the Schrödinger quotation at the beginning of this chapter). An electron becomes, on this conception, a kind of 'abstract particular', a way a given system is, a mode. The system must be various ways, but the ways do not *make up* the system.

The situation is sometimes described, as Schrödinger describes it, as one in which an electron, in entering into a relation with another electron, 'loses its identity'. This suggests that, left to themselves, electrons are well-behaved simple substances possessing definite identities. When electrons enter into various relations with other particles, however, their individuality is compromised. Here is how E. J. Lowe puts it:

Suppose . . . that in an ionization chamber a free electron *a* is captured by a certain atom to form a negative ion which, a short time later, reverts to a neutral state by releasing an electron *b* . . . [A]ccording to currently accepted quantum-mechanical principles there may simply be no objective fact of the matter as to whether or not *a* is identical with *b* . . . [W]hat is being proposed here is *not* merely that we may well have no way of telling whether or not *a* and *b* are identical, which would imply only an epistemic indeterminacy.

---

4 This point is defended forcefully by Schaffer (2010) in a defence of 'priority monism'.

It is well known that the sort of indeterminacy presupposed by orthodox interpretations of quantum theory is more than merely epistemic—it is *ontic*. The key feature of the example is that in such an interaction electron *a* and other electrons in the outer shell of the relevant atom enter an 'entangled' or 'superposed' state in which the *number* of electrons present is determinate but the identity of any one of them with *a* is not, thus rendering likewise indeterminate the identity of *a* with the released electron *b*. (Lowe 1994, 110)

The idea that an electron could 'lose its identity'—and perhaps later reacquire some definite identity or other—could be avoided if you took electrons—and the particles, generally—to exist *only* in entangled states. In this regard, the particles would more naturally be regarded as waves, the position towards which Schrödinger himself was inclined.[5] Waves do not exhibit identity crises of the kind associated with entangled particles. On such a conception, fundamental 'particles' are not like water droplets that cease to exist as definite individuals when they combine to make up a puddle. Rather, they are abstractions, the nature and identity of which is dependent on systems to which they belong.

I have followed other authors who have written on this topic and spoken of 'holism' and holistic systems. But what is holism? Explicating holism is subject to the same sorts of difficulty associated with conventional explications of 'emergence'. In fact the ideas are related. A holistic system is one in which a property emerges, a property of the *system,* a property not founded on properties of the system's components. This, by my lights, requires the emergence of a substance, the bearer of the emergent property. Emergent substances, like any substances, are simple; emergent substances lack substantial parts. What might initially be thought of as the system's parts are abstractions, ways the system is. Particles would be abstractions in the sense that a billiard ball's redness would be an abstraction, a way the billiard ball is, or a Cartesian billiard ball would be an abstraction, a way space is configured.

With all this in mind, I propose to consider very briefly one aspect of quantum physics that has led its interpreters to abandon the classical picture and embrace some form of holism. My discussion ignores a

5 See French and Krause 2006, § 3.6. This might also be close to what Schaffer has in mind.

host of technical distinctions, and disagreements in interpretation that characterize the literature. My aim, however, is merely to depict the kinds of consideration that might be thought to undermine corpuscularism, the idea that the universe and objects in it are made up of myriad fundamental simple substances.

3.8   Quantum Statistics

Consider two coins, *a* and *b*, and two 'states' of the coins, *H* and *T*, heads and tails. Let *Ha* indicate that *a* is in the *H* state and *Ta* represent *a*'s being in the *T* state. When you toss both coins, you anticipate four possible outcomes:

(1) *Ha* and *Hb*
(2) *Ha* and *Tb*
(3) *Ta* and *Hb*
(4) *Ta* and *Tb*

Assume that the coins are 'fair': the probability of *Ha* is ½ and the probability of *Ta* is ½; similarly, the probability of *Hb* is ½ and the probability of *Tb* is ½. Each of these paired outcomes, (1)–(4) is equally probable, each has a probability of ¼. This means that the probability that the coins are in state (2) *or* state (3), a state in which one coin is *H* and the other *T*, is ½. Sure enough, when you observe many sequences of tosses, this is what you find.

$$\text{Prob}(Ha) = \tfrac{1}{2} \qquad\qquad \text{Prob}(Ta) = \tfrac{1}{2}$$
$$\text{Prob}(Hb) = \tfrac{1}{2} \qquad\qquad \text{Prob}(Tb) = \tfrac{1}{2}$$
$$\text{Prob}(Ha + Hb) = \text{Prob}(Ha) \times \text{Prob}(Hb/Ha) = \tfrac{1}{4}$$
$$\text{Prob}(Ta + Tb) = \text{Prob}(Ta) \times \text{Prob}(Tb/Ta) = \tfrac{1}{4}$$
$$\text{Prob}((Ta + Hb) \text{ or } (Ha + Tb)) = \tfrac{1}{2}$$

Classical Statistics

When you look at paired quantum particles, matters are different. Pretend that *a* and *b* are, not coins, but pairs of electrons and that *H* and *T* are states of these electrons. Stunningly, the theory predicts, and observation confirms, that the probability that both electrons will be in the *H*-state is ⅓, the probability that both will be in the *T*-state is ⅓, and the probability that one electron will be in the *H*-state, the

other in the T-state is $\frac{1}{3}$. This encourages the thought that (2) and (3) above must be treated as a *single* state.

$$\text{Prob}(Ha) = \frac{1}{2} \qquad \text{Prob}(Ta) = \frac{1}{2}$$
$$\text{Prob}(Hb) = \frac{1}{2} \qquad \text{Prob}(Tb) = \frac{1}{2}$$
$$\text{Prob}(Ha + Hb) = \text{Prob}(Ha) \times \text{Prob}(Hb/Ha) = \frac{1}{3}$$
$$\text{Prob}(Ta + Tb) = \text{Prob}(Ta) \times \text{Prob}(Tb/Ta) = \frac{1}{3}$$
$$\text{Prob}((Ta + Hb) \text{ or } (Ha + Tb)) = \frac{1}{3}$$
Nonclassical Statistics

This is what leads to talk of a 'loss of identity'. Considered together, the particles 'shed their identity'. There continue to be two of them, two particles, but there is no 'fact of the matter' as to which is which.

In quantum physics a permutation of indistinguishable objects between states is typically not regarded as leading to a new arrangement. Hence, it was concluded, such objects should not be regarded as individuals, or at least in the sense that they are not taken to be in [classical] statistics. (French and Krause 2006, 84)

True, there are other ways of construing such cases. Return to the original example of tossed coins. You might expect to obtain the $\frac{1}{3}$, $\frac{1}{3}$, $\frac{1}{3}$ probability distribution if the tosses were not in fact independent, if outcomes causally influenced other outcomes. The difficulty of supposing that this is what is going on in the quantum case is that entangled particles can be widely dispersed. Particles would need to influence one another instantaneously at a distance. As Reichenbach notes, 'these causal relationships would represent action at a distance, since the particles could be far apart; that is the dependence relations would constitute causal anomalies' (Reichenbach 1956, 234).

Another possibility, one suggested by Einstein and mentioned already, is that the particles in question are not particles after all, but *waves* (French and Krause 2006, 88–9). In this case, what were initially called particles, would be modes, not substances, not substantial constituents of systems. *Systems* would take on the role of substances. Such systems would not be made up of particles. Particles would be abstractions, dependent on systems to which they belonged. Were quantum 'entanglement' universal, the universe would amount to a single dynamic substance. What might naturally be regarded as objects—common-sense substances, including the particles—would be

modifications of the one substance, modes. Paul Teller puts this in terms of 'inherent relations':

The universe is a more deeply intermeshed web than we had thought. Indeed, according to quantum mechanics, the extent of entanglement through inherent relations is all pervasive. (Teller 1986, 80–1)

I prefer to leave open the ontological lessons to be extracted from quantum physics. Quantum physics supplies us with equations that physicists regard as indispensable. The difficulty is to understand what these equations say about the universe. Attempts to bridge the gap between theory and ontology abound in the form of multiple 'interpretations' of quantum theory. However, there is scant agreement among physicists as to the correct, or best interpretation. My contention is that, whatever ontology you draw out of fundamental physics, it will be an ontology of substance and property. The scientific question is, what are the substances? Are they the particles? The fields? The universe as a whole? And, what are the properties? Mass, charge, spin? The electrons and quarks, themselves?

3.9 Taking Stock

The universe—*any* universe—must include substances and properties. Substance and property are reciprocal, complementary categories. Substances are the bearers of properties, properties are ways substances are. Substances and properties alike are abstractions: you can consider a substance as a property bearer or you can consider its properties, but substances and properties could not exist apart. Substances are not compound entities made up of 'thin' or bare particulars clothed in properties.

Substances must be simple, lacking substantial proper parts, although a simple substance can bear many properties and have endless spatial (and perhaps temporal) parts or divisions. Substances are *mereologically* simple. Complex entities, if there are any, could not be bearers of properties. If simple substances are genuine atoms, you might put this by saying that atomism is true of necessity. The question is, what are the atoms? And this is an empirical question. There might be many atoms, as Locke and Leibniz thought, or one

big atom (Descartes, Spinoza), or there might be a handful of atoms (the fundamental fields).

A complex object consisting of propertied substances standing in particular relations can serve as a truthmaker for a range of truths. This billiard ball is apparently a complex entity, a quasi-substance, not a substance. You can say truly that the billiard ball is red, that it is spherical, that it has a particular mass. But the truthmakers for such claims are the constituent substances, with their properties, arranged and interacting as they are. Only in the case of simples are judgements of the form 'a is ϕ', made true by a substance, a's, bearing a property, ϕ.

What the substances are, what the properties are is an empirical question, one you must consult fundamental physics to answer. Thus, although the categories of substance and property prove indispensable, this does not imply that corpuscularism is true. The substances could be fields, space or space–time itself, or Spinoza's *Deus sive Natura*, the universe as a whole. Were that so, what you might ordinarily regard as paradigmatic substances would turn out to be properties: modes, ways the substance or substances are. Trees, electrons, planets, then, could turn out to be particular fluctuations in the quantum field or thickenings in space–time. The guiding idea here is that, whatever story you tell, it will be one featuring substances and properties.

The focus in this and the previous chapter has been on substance. What of properties? This is the topic to which I shall now turn.

CHAPTER 4
Properties

Whatsoever the Mind perceives in it self, or is the immediate object of Perception, Thought, or Understanding, that I call *Idea*; and the power to produce any Idea in our mind, I call *Quality* of the Subject wherein that power is. Thus a Snow-ball having the power to produce in us the *Ideas* of *White*, *Cold*, and *Round*, the Powers to produce those *Ideas* in us, as they are in the Snow-ball, I call *Qualities*.

(Locke 1690, II, vii, 8)

The possibility of action . . . is the criterion of existence, and the test of substantiality.

(Santayana 1930, 107)

All things are full of gods.

(Thales of Miletus)

4.1 Substances, Properties, Truthmakers

In previous chapters, I have promoted the idea that substance and property are complementary categories. Locke describes substance and property as 'correlatives'. Substance and property alike are abstractions. They are abstractions, not in the sense that they are 'abstract entities' residing 'outside' space–time, whatever that could mean. They are abstractions in the traditional sense, what we pick out by engaging in 'partial consideration'.

The *Oxford English Dictionary* tells us that to abstract is 'to separate in mental conception', and that an abstract entity is 'the result of abstracting; the idea of something which has no independent existence'. Thus you can consider an object's *properties*, and you can consider the

object as a property *bearer*, a *substance*. You can consider the redness and sphericality of the tomato you hold in your hand, for instance, and you can consider the tomato itself as something red and spherical.

In explicating abstraction as Locke does, in terms of 'partial consideration', I do not mean to suggest that substance and property are in any sense mind-dependent. Substances are 'out there' mind-independently with all their myriad properties. At any given time, a substance and its properties are separable, however, only in conception or thought, not in reality. In abstracting we take note of perfectly real facets of the universe. Any account of the universe we could hope to have from the sciences and, in particular, from fundamental physics, is going to be framed in terms of substances and properties. What the substances and properties *are* is a decidedly empirical question. The substances could turn out to be particles, or fields, or space–time itself, or the One, the cosmos as a whole. Whatever they, or it, turn out to be, they or it will be some way or other, they or it will be propertied.

I suspect that the most controversial aspect of my conception of property and substance is the thought that substances must be simple, lacking in substantial parts, and consequently that the only genuine properties are properties of simple substances. This means that tables, tomatoes, and planets, because they are not simples, are not substances; characteristics of such things are not properties. This will seem shocking only if you are concerned that it requires giving up on tables, tomatoes, and planets. But this is not what I think, nor is it implied by the position I am advancing here. Judgements about tables, trees, and planets and their various characteristics can be perfectly true. They are made true, however, not by substances identifiable as tables, tomatoes, and planets, and properties of these substances.

The central idea could be put theologically. Assume, for the sake of illustration, that the universe comprises corpuscles interacting with one another and occupying regions of space–time. When God creates the universe and arranges the corpuscles in space–time, God has *thereby* made it the case that many of our judgements concerning tables, tomatoes, and planets are true. The universe comprises atoms arranged as they are in space–time, not these *plus* tables, tomatoes, and planets.

I have put this in terms of substances, but the very same point holds for properties. In arranging the corpuscles in space–time, God is

arranging substances with various properties. These arrangements serve as truthmakers for judgements concerning all the characteristics of tables, tomatoes, and planets. When God arranges the corpuscles in the right way, the result is a configuration that answers equally to 'this is a tomato' and 'this is red'. Tomatoes are not substances, red is not a property. But this in no way undermines judgements about tomatoes and their colours.

I shall return to these issues in chapter 8 when I take up the truthmaking relation. I mention them here only to head off the thought that I am advocating a kind of nihilism or anti-realism about tables, tomatoes, and planets. On the contrary; my belief is that fundamental physics discloses the deep story about such things, fundamental physics enables us to see what must be the case if judgements about tables, tomatoes, and planets are, as they often enough are, *true*.

4.2   Properties as Qualities

In chapter 2, I suggested that there are good reasons to regard properties as particulars, *modes*, particular ways particular substances are, what philosophers nowadays are likely to call *tropes*. The arguments of that chapter did not depend on this supposition, however. My aim there was to move away from the idea that the most important role for substance was the particularizing role: substances are required for particularity if properties are universals. Even if properties are particulars, however, even if we do not need substances as particularizers, if you have properties, you need substances, and, correlatively, substances require properties.

In this chapter, I am going to continue to suppose that properties are particulars, but, again, little of what I have to say will depend on this supposition. I shall reserve discussion of whether properties are best regarded as universals or as particulars for chapter 5.

In chapter 2 I alluded to the fact that 'property' can be used both in a relaxed sense—to mean whatever answers to a true assertion about an object—and in an ontologically serious sense. The discussion was meant to nudge you in the direction of a conception of properties according to which genuine properties are ways substances are.

Substances are simple, so the only genuine properties are properties of simple substances. Our conception of substances and properties is born in the course of interactions with ordinary objects, however. You can treat something as a property whenever you can treat it as a way some substance is, which in turn requires your treating something that might or might not be a substance as a substance, as a single, unified thing. In this way you can speak in a relaxed way about a tomato's redness or sphericality as properties of the tomato. I shall continue to dwell on simple examples of this kind to illustrate general truths about properties, then extend the points to properties considered with due ontological seriousness.

When you consider a tomato's redness and sphericality, it is natural to regard its redness and sphericality as *qualities* of the tomato. A long-standing philosophical tradition regards sphericality as a *primary quality*, redness as a *secondary quality*. Early modern philosophers and scientists—Galileo and Descartes, for instance—held that only primary qualities are genuine; secondary qualities belong only to the appearances.

A better way of thinking about the primary–secondary division aligns with the aforementioned distinction between genuine properties and properties in the relaxed sense, 'properties by courtesy', *quasi-properties*, Armstrong's 'structural properties', what you get when you organize substances in the right way. It is important to see that the traditional view was not that the tomato's shape, for instance, but not its colour, was genuine, something the tomato really possessed. The idea, rather, was that the fundamental substances had shapes, but not colours. Colours result only from *arrangements* of the fundamental things. Were you a corpuscularist, you would think that a corpuscle must have some shape or other, some size, some mass, but no colour, no odour, no taste.

Berkeley lampooned this view on the grounds that it was impossible to conceive of an object's having primary qualities, but no secondary qualities, having a shape, for instance, without having a colour. However, this argument appears to conflate features of *representations* of corpuscles, including perceptual representations, with features of corpuscles. If you create a visual representation of a corpuscle, if you paint a portrait of a corpuscle, you will need to give it a colour to distinguish it from the background. If you observe a

corpuscle in a scanning electron microscope, the image produced will have some colour. But it by no means follows that corpuscles *themselves* must be coloured. When you look at features of tomatoes responsible for their red appearance, you find complex chemical and structural relationships among the molecules making up the tomato's surface and insides. If these are responsible for objects' colours, it makes no sense to suppose that corpuscles could be coloured: they lack the requisite structural complexity.

So the fundamental things—the corpuscles, perhaps—possess primary qualities. Secondary qualities are no addition of being, hence of no concern to fundamental physics. The view can be updated by supposing that the primary qualities are those possessed by the fundamental things, whatever they might turn out to be. An electron's mass, its charge, its spin might be primary qualities.

All this was a lead up to a single point, namely that properties are *qualities* of their possessors. Different kinds of property differ *qualitatively*. This is patent in the case of the tomato's redness and sphericality. What of the properties of an electron? Is it clear that an electron's mass or charge are qualities of the electron?

Given that properties are properties of the fundamental things, any argument to the conclusion that properties are qualities confronts the fact that our access to the fundamental things is inevitably indirect. We have no choice but to begin with observable complexes such as tomatoes. The tomato's shape and colour do appear to be paradigmatic qualities. You might wonder where such qualities come from if the fundamental things are qualitatively barren. It will not do to imagine that qualitativity emerges, at least not if the account of emergence in chapter 2 is correct: emergent properties must be properties of fundamental things, simple substances. One option is to suppose that qualities are not properties: truthmakers for judgements concerning qualities are non-qualitative ways the universe is. I shall offer considerations against this possibility in due course.

4.3 'Categorical Properties'

Before moving on, let me tie up a terminological loose end. I have been speaking of qualities, but many philosophers prefer to talk of

'categorical' properties. The notion of a categorical property has an ontologically unsavory ancestry to which its latter day proponents seem oblivious. Categorical properties were introduced originally by philosophers bent on distinguishing properties possessed 'categorically', that is, *actually*, by objects from mere if–then, *conditional* properties, mere potentialities. Here is Ryle in a representative passage.

When we describe the glass as brittle, or sugar as soluble, we are using dispositional concepts, the logical force of which is this. The brittleness of the glass does not consist in the fact that it is at a given moment actually being shivered. To say that it is brittle is to say that if it ever is, or ever had been struck or strained, it would fly, or have flown, into fragments. To say that sugar is soluble is to say that it would dissolve, or would have dissolved, if immersed in water. (Ryle 1949, 43)

Much has been written about the prospects of 'analysing' talk of powers or dispositions conditionally (see Martin 1994). The analytic project encourages the idea that powers or dispositions are non-categorical, that is, non-actual. On such a view dispositions are *mere* potentialities to be contrasted with here and now, real, categorical properties of objects. Ryle puts it this way:

To possess a dispositional property is not to be in a particular state, or to undergo a particular change; it is to be bound or liable to be in a particular state, or to undergo a particular change, when a particular condition is realized. (Ryle 1949, 43)

Objects are as they are categorically, and objects would do this or that, presumably, owing to their here and now categorical natures. But this is not to *identify* dispositions with categorical features of objects. Doing so would involve a 'category mistake'. Actuality is one thing, potentiality something else altogether.

Not all philosophers have been as sanguine as Ryle about detaching truths about dispositions or powers from truths about the categorical makeup of objects. The idea is that there must be *something* categorical about the glass, something about the glass *here and now*, responsible for its being true that the glass is brittle, that it would break were it struck. It is no wonder, then, that so many philosophers have followed Prior, Pargetter, and Jackson (1982), who defend the idea that dispositions are 'grounded' in, or 'realized by', objects'

categorical properties. Until very recently, this conception of the relation of dispositional and categorical properties was so widely accepted as to constitute the default view.

I shall have more to say about this approach—and its sisters, cousins, and aunts—in the next section (and in subsequent chapters). For the moment, my aim is simply to note that talk of categorical properties might usefully be understood as an oblique, arguably misleading, way of talking about *qualities*. At any rate, qualities are categorical; qualities are here and now, actual, not merely potential, features of the objects of which they are qualities. I prefer to speak of qualities or the qualitative because 'categorical' has taken on the sense of 'non-dispositional', and I mean to leave open the possibility that qualities, or some qualities, are *themselves* dispositional. Were that so, categorical properties could be dispositional, potentialities could be actual, a thesis that appears, unfairly, to condemn itself.

'Qualitative' is less tendentious than 'categorical', not least because it leaves open the possibility that powers or dispositions are actual, 'categorical' features of objects. If you start with a categorization of properties as dispositional *or* categorical, you will be embracing a terminology that forecloses conceptions of properties that were once commonplace. This is a nice example of how a focus on conceptual analysis as practised by Ryle and armies of other linguistically inclined philosophers can cast a pall over serious ontology. Who could doubt that categorical and non-categorical, conditional locutions are importantly distinct? The mistake is to move unselfconsciously from this linguistic platitude to the idea that such locutions must designate utterly distinct kinds of property.

4.4 Properties as Powers

Suppose properties *are* qualities. You might wonder whether this could be the whole story. One important role played by properties is explanatory: we appeal to objects' properties to explain what those objects *do* or *would do*. This suggests a connection between properties and *powers* or *dispositions*, terms I have been using interchangeably. Consider the tomato's sphericality. In virtue of being spherical, the tomato rolls or would roll. In virtue of being red, the tomato would

look red. Similarly, properties of an electron—its mass and charge, for instance—appear to equip the electron with powers. In virtue of having a particular negative charge an electron would repel other electrons; in virtue of having a particular mass an electron would accelerate in a particular way in a gravitational field of a particular sort.

Considerations of this kind have suggested to many philosophers that properties, or some properties, are *powers*. Philosophers who embrace this view include those who hold that every property (or every property of objects residing in space–time) is a power (Bird 2007), and those who hold a mixed conception: some properties are qualities, some are powers (Ellis 2001; Molnar 2003; Unger 2006). The latter contingent includes philosophers who see powers as 'grounded in' or 'realized by' qualities (Prior, Pargetter, and Jackson 1982; Mumford 1998). What all these views have in common is the idea that powers and qualities are mutually exclusive: if a property is a power, that property is not a quality; if a property is a quality, it itself is powerless, inert.

This last assertion calls for qualification. Philosophers who 'ground' powers in objects' 'categorical' properties—Prior, Pargetter, and Jackson, for instance—argue that because powers or dispositions are grounded in or realized by non-dispositional categorical properties, dispositions themselves must be causally inert. Any causal contribution a disposition might make to the behaviour of its possessor is pre-empted by its ground or realizer.

A striking feature of this conception is that it implies that dispositions or powers *themselves*, as distinct from their grounds, are powerless. Thus, Frank Jackson (1998, 92), an architect of the idea that powers are grounded in or realized by categorical properties, regards it as obvious that powers are powerless, describing the alternative as requiring 'a curious, ontologically extravagant kind of overdetermination'. The argument will be familiar to anyone who has followed recent debates over the causal efficacy of 'multiply realized' mental properties. I shall return to this issue in chapter 9.

One unwelcome consequence of regarding powers and qualities as mutually exclusive is that qualities would be thereby rendered unobservable—at least they would be were observation to include a causal component. If your seeing the tomato's shape requires that the

tomato's shape be part of what is causally responsible for your perceptual experience, then it appears to follow that either shapes are powers or shapes are unperceivable. The same point would hold for colours or any other apparent quality. This way lies Berkeleyan idealism.

Such considerations might lead you to conclude that shapes, colours, and the like are not qualities after all, they are powers. Once you start down this road, it is easy to conclude that, really, properties, or at least properties capable of making a causal difference, must be powers, *not* qualities. Qualities would be causally inert, they could make no difference to what anything does, or would do, they would be unperceivable, unknowable *quiddities*.

We are, it would seem, left with three possibilities:

(1) Properties are (powerless) qualities.
(2) Properties are purely powers.
(3) Some properties are qualities, some are powers.

My suggestion is that the most promising option is one that doesn't occur, one not on this list:

(4) Properties are *powerful qualities*.

The tomato's sphericality is a quality, but it is in virtue of possessing this quality that the tomato would roll. The tomato's redness is equally a quality, and in virtue of possessing this quality, the tomato would look red in bright sunlight, black in blue light. A property's dispositionality and its qualitativity are not aspects or properties of the property. Rather they inseparably constitute its nature.

One conception of the property–power relation I have not mentioned is a conception defended by Sydney Shoemaker (1980, 1998, 2007). According to Shoemaker, properties are not to be identified with powers; properties *confer* or *bestow* powers on their possessors. Thus, although it is true that the tomato would roll because it is spherical, the tomato's sphericality is not itself a power. The power, rather, is bestowed on the tomato by virtue of the tomato's being spherical.

I do not pretend to have a firm grip on Shoemaker's ontology, but here is one way you might think about it. Suppose properties were universals, instantiated by individual objects. A property's being instantiated is a matter of the object's possessing certain powers,

those powers, namely, associated with the property. Properties are individuated by the powers they bestow, powers possessed by their instances. Powers are the here-and-now representatives of properties, property stand-ins in the universe of physical objects. Properties remain aloof from the causal mix: powers do all the work, properties get all the credit.

If this is how Shoemaker is looking at it, his conception of properties is sharply at odds with the conception advanced here: properties are ways substances are, and, or so it seems, substances do or would do what they do owing to the ways they are, the ways themselves, not their representatives. As I hope to show, the closer you look at properties, the less attractive a Shoemaker-style view lookt.[1]

4.5 'Scientiphicalism'

In working up to what I believe is a promising view of properties I will use Peter Unger's recent discussion of powers and qualities as a stalking horse (Unger 2006). My sense is that, although Unger's heart is in the right place, he succumbs eventually to kinds of error—traceable, ultimately, to Ryle's linguisticism—that continue to plague contemporary discussions of properties and powers.

Unger's primary target is something he calls 'scientiphicalism'. Scientiphicalism is less a clearly articulated thesis than a *Weltanschauung*, a collection of implicit doctrines that inform the way we think about the universe and our place in it. Central to scientiphicalism is the idea that the physical realm is bereft of qualities. The concern of science is to provide an accounting of objects' capacities for interaction. Qualities play no role in this accounting. To the extent that our conscious experiences are imbued with qualitative character, experiences stand 'outside'—or 'at the limits of'—the physical.

Electrons have a definite mass and charge. We characterize these (as John Foster has put it) 'topic neutrally' by spelling out their contributions to the capacities of electrons to act and be acted

---

[1] Shoemaker is a longstanding exponent of the doctrine that many properties are 'multiply realized.' His most recent account of the 'realizing' relation appeals to powers as realizers. See Shoemaker 2007; for discussion see Heil 2003b, 2011.

upon. When you reflect on electrons' 'intrinsic nature', it seems obvious that once you have set out their propensities to do this or that in concert with other elementary things, you have said all there is to say about them. So it is with the other particles, the forces, and the fields. We are left with a picture of the physical universe as a shadowy domain in which objects, acting in ways that reflect their propensities, combine to form more complex objects the nature of which is determined wholly by propensities of *their* constituents.

This dreary picture leaves no room for Technicolour vistas, no room for experienced boomings and buzzings. To the extent that we tolerate qualities in our thinking about the universe, such things appear to be, at best, epiphenomenal add-ons, byproducts of complex physical processes, hazy characteristics of the conscious states of sentient creatures. To imagine that qualities themselves could affect or be affected by physical goings-on is to make a mistake of a fundamental sort, the mistake of imputing efficacy to the inherently inefficacious.

The point might be put in terms of causal completeness or 'closure', the idea that the physical universe is 'causally autonomous'. If the physical realm excludes qualities, the idea that qualities might 'make a causal difference' would require abandonment of the idea that the physical universe is causally self-contained. To the extent that you accept scientiphicalism, however, you are more likely to regard the thought that mental qualities, or indeed qualities of any kind, could affect anything physical as something akin to a category mistake: no need to appeal to a contentious closure thesis.

All this is at least superficially reminiscent of a conception of the universe you find in Galileo. There are the primary qualities—mass, size, shape, a capacity for motion—and the secondary qualities—colours, sounds, tastes, odours. Science is concerned only with objects' primary qualities. It is these that determine how objects behave or would behave. The rest, the secondary qualities, are subjective mental ephemera falling outside the purview of science, hence outside any respectable ontology of the physical universe.

A conception of this kind dramatically bifurcates the mental and the physical, the qualitative and the powerful. This bifurcation is codified in Descartes' characterization of the physical universe as extended and unthinking, and minds as thinking and unextended.

It is reflected today in the widespread presumption that mental and physical properties differ absolutely.

I need not remind you of the many philosophical difficulties that attend this kind of division of reality. Attempts to surmount these difficulties so as to produce a unified picture of mind and universe have, for the most part, embraced reductive strategies. Physicalists hope to 'reduce' the mental to the physical. Idealists take the opposite tack, 'constructing' the physical from the mental. The most prominent recent twist, 'nonreductive physicalism', endeavours to find a middle ground by conceiving of the mental as wholly dependent on and determined by, but nevertheless distinct from, the physical. The most compelling reason for embracing nonreductive physicalism is negative: the alternatives look so much worse.

I might mention two philosophers who seem not to fit the scientiphicalist mould: Spinoza and Davidson.[2] Both embrace a thoroughgoing monism that turns mental–physical dualities into creatures of reason. This is not idealism, but neither is it physicalism, reductive or otherwise. There is no interesting sense in which the physical is privileged. How could it be, if the mental–physical distinction amounts to a distinction in mode of conception, not a distinction in reality?

I mention Spinoza and Davidson in order to set them to one side. The focus at the moment is on Unger's scientiphicalism, a conception of reality that accepts a sharp distinction between powers and qualities, and between the physical and the mental. The suspicion is that these are just two ways of making the same distinction. In banishing qualities, scientiphicalism excludes the mental, or whatever portions of the mental resist 'reduction', from the physical domain. Features of the universe as we experience it that fail to turn up in our best physical theories are relegated to the mind. The task of making sense of mental phenomena is left to philosophers and to superannuated physicists and neuroscientists with time on their hands. If mental phenomena make no physical difference, they are excluded from the purview of serious science. Hence a conception of physical reality stripped of qualities. There are the objects—the

particles, the fields, space–time—and propensities of these to affect and be affected by one another in various ways. Qualities are, at best, mental projections, unwelcome embarrassments that, so far as the physical realm is concerned, we could live without.[3]

4.6 Idealism and the 'Mystery of the Physical'

Unger argues that scientiphicalism, in excising qualities from the physical universe, renders the physical universe 'humanly unintelligible'. His argument (which I shall discuss in more detail below) belongs to a family of arguments traceable at least to Berkeley. Berkeley argues, in effect, that the elimination of qualities from the physical realm leads directly to idealism. He asks us to consider what follows from the supposition that qualities, or the secondary qualities, are exclusively mental. When you attempt to conceive of a material body, you inevitably conceive of it as having various qualities: it has a colour, a shape, it makes a sound, it feels warm. Subtract these, and you are left with an empty conception.

Those who assert that figure, motion, and the rest of the primary or original qualities do exist without the mind, in unthinking substances, do at the same time acknowledge that colours, sounds, heat, cold, and such like secondary qualities, do not, which they tell us are sensations existing in the mind alone, that depend on and are occasioned by the different size, texture, and motion of the minute particles of matter. This they take for an undoubted truth, which they can demonstrate beyond all exception. Now if it be certain that those original qualities are inseparably united with the other sensible qualities, and not, even in thought, capable of being abstracted from them, it plainly follows that they exist only in the mind. But I desire any one to reflect and try, whether he can by any abstraction of thought conceive the extension and motion of a body, without all other sensible qualities. For my own part, I see evidently that it is not in my power to frame an idea of a body extended and moved, but I must withal give it some colour or other sensible quality which is acknowledged to exist only in the mind. In short,

---

2 On Davidson, see Heil 2008b and chapter 9 below. Although he describes himself as a physicalist, Galen Strawson could be added to the list; see Strawson 2006.

3 David Chalmers, an influential exponent of scientiphicalism, regards the reconciliation of mental qualities with the scientiphicalist world view as the last remaining 'hard problem' (see his 1996). The alternative, he suggests, is either 'something like epiphenomenalism' or panpsychism.

extension, figure, and motion, abstracted from all other qualities, are inconceivable. (1710, § 10)

Berkeley's aim is to parlay the apparent inseparability of primary and secondary qualities into a defence of idealism. This theme is echoed by John Foster (1982) who argues that (as Unger might put it) scientiphicalist attempts to specify the intrinsic nature of material bodies yield at best 'topic neutral specifications', specifications that could be satisfied by immaterial states of affairs. We depict material bodies as spatially located, for instance, and extended three dimensionally. Space is presumed to be material because it serves as a medium for material bodies; bodies are taken to be material because they occupy regions of space. But

the two specifications cancel out, leaving us with a combined specification which is … topic neutral—a specification which characterizes matter-in-physical-space as a 3-dimensionally extended substance (of whatever intrinsic nature) in a 3-dimensionally extended medium (of whatever intrinsic nature). (1982, 57)

When you turn to science for help in ascertaining the intrinsic qualitative nature of occupants of the physical realm you find little help.

Scientific analysis uncovers spatiotemporal arrangement and nomological organization, but does not reveal the intrinsic nature of the fundamental space-occupying substance or substances which are thus arranged and organized. It specifies the intrinsic nature of those substances only opaquely, in terms of their causal powers and sensitivities—the powers and sensitivities which, in the framework of natural law, their intrinsic properties sustain. (1982, 65)

The intrinsic qualitative nature of matter is, Foster holds, 'inscrutable'.[4] Mental qualities, which constitute the intrinsic nature of states of mind are, in contrast, immediately evident. If you think of powers as being grounded in the intrinsic natures of objects, you are left with an empty, purely formal conception of physical—that is, non-mental—reality. This is Unger's 'Mystery of the Physical'.

---

4 Ladyman et al. 2007 advance a similar argument, albeit in the course of defending a version of 'structural realism.'

4.7 Powers without Qualities

Berkeley and Foster, along with a bevy of philosophers unsympathetic to idealism, object to attempts to detach powers from qualities on the grounds that descriptions of a universe wholly in terms of powers manifested by objects populating the universe omit the intrinsic qualitative nature of the bearers of powers.[5] As Foster puts it, scientific descriptions are topic neutral, leaving the intrinsic character of the physical 'inscrutable'. The idea is that there is no reason to suppose that whatever items play the 'physical role' are non-mental. If you couple this with the thought that the only intrinsic qualitative natures we have any conception of are those we encounter in conscious experience, you are on your way to the Berkeleyan thesis that objects lacking mental qualities are flatly inconceivable.

One kind of scientiphicalist response to this line of reasoning begins by conceding its premises. We have no conception of what an intrinsic physical quality could be. Perhaps this poses no threat to the scientiphicalist project, however. Perhaps it is of the nature of physical objects altogether to lack intrinsic qualities. A physical object is wholly constituted by powers to affect and be affected by other physical objects. In telling us about objects' powers, science is telling us all there is to know about the objects. As Foster puts it, on such a conception 'each particle is, in itself, no more than a mobile cluster of causal powers, there being no "substantial" space-occupant which possesses the powers and on whose categorical nature the powers are grounded'.[6]

Something like this seems to have been what Priestley, for instance, had in mind in defending the thesis that the physical realm is made up of

certain *centres of various attractions and repulsions*, extending indefinitely in all directions, the whole effect of them to be upon each other; … a *compages* of

---

5 Versions of the argument *not* aimed at a defence of idealism can be found in Armstrong (1961, chap. 15; 1999); Smart (1963, 73–5); Campbell (1976, 93–4); Blackburn (1990); Martin (1997, 213–17; 2008, 61–9); and Heil (2003a, 97–110).
6 Foster 1982, 67–8. Foster cites Leibniz, Boscovich, Kant, Priestley, and Faraday as among those who have defended this kind of view. See Harré and Madden (1975), chap. 9.

these centres, placed within the sphere of each other's attraction, will constitute a body that we term *compact*.[7]

Following Berkeley, Foster contends that such a conception of the physical universe is incoherent.

The main problem is that if all the fundamental particles are construed in this way, there seem to be no physical items in terms of whose behavior the content of the powers could be specified, and consequently, it seems that, in the last analysis, there is nothing which the powers are powers to do. (Foster 1982, 68)

The difficulty, he thinks, manifests itself in a regress. Pretend the fundamental physical things are Newtonian atoms regarded as bundles of powers, 'mobile spheres of impenetrability'.

The problem arises when we ask: 'To what is a sphere of impenetrability impenetrable?' The answer is 'To other atoms, i.e. to other spheres of impenetrability.' But this means that the specification of the content of the atom-constituting power is viciously regressive: each atom is a sphere of impenetrability to any other sphere of impenetrability to any other sphere of impenetrability . . . and so on ad infinitum. From which it follows that the notion of such a power is incoherent, since there is nothing which the power is a power to do. To conceive of a sphere of impenetrability, we have to postulate some other type of space-occupant whose passage it is empowered to obstruct. (68)

Keith Campbell, no friend of idealism, had, six years earlier, advanced a similar argument in a discussion of Roger Boscovich's conception of the universe as an arrangement of material points the intrinsic nature of which consists solely of a power to accelerate other points. Campbell wonders

What is at a material point? What distinguishes a location in space where there is a point from one where there is no such thing? All we can say is: at a material point there is something which accelerates other somethings which in turn accelerate somethings (including the first) which in turn . . . But what an odd object this is. Its *only* feature is to have an effect on things which have an effect on things which have an effect on things which . . . We

7 Priestley (1782, 239), a version of which is cited in Harré and Madden (1975, 172).

seem to be caught in a regress or circle, forever unable to say what these things *are* which have an effect on each other. (1976, 93)

If all there is to a material point is its power to accelerate other points, *what* is accelerated?

When one point moves another, all that has been shifted is a power to shift powers to shift. . . . But powers to shift *what*? To be coherent, I consider that Boscovich's points must be *somethings* which have the power to shift one another. They must have some intrinsic features which make them things in their own right, and they must in addition have the power to shift one another. Then, and only then, will there be something to move about. There must be some answer to the question What is at a point? independent of accelerative capacity. (1976, 93)

Note that it is not simply that Boscovich's material points or atoms must have 'intrinsic natures'. Boscovich's powers *are* intrinsic, right there, on the scene at the spatial points at which they are located. The idea rather is that bearers of powers must have some intrinsic *qualitative* nature.

4.8 Purely Relational Universes

The situation resembles what we encounter in attempts to think about purely relational universes (see e.g. Dipert 1997; Ladyman et al. 2007; cf. Heil 2003a, 102–5). Consider a universe in which objects are constituted by relations into which they enter. Richard Holton (1999) describes a universe comprising four objects, *A*, *B*, *C*, and *D*, each wholly constituted by relations in which it stands to other objects:

A is directly to the left of B and directly above C; B is directly to the right of A, and directly above D; C is directly to the left of D, and directly below A; D is directly to the right of C, and directly below B.

Holton's universe might be represented this way:

A•   *B*•
C•   *D*•

The labeled points here are meant only as visual aids, however. 'There really is nothing more to *A*, *B*, *C*, and *D* than that given by the

descriptions' (10). You get from this representation to the universe itself by erasing the points (and their labels) while leaving the relations intact. The result is a universe evidently bereft of qualities and qualitied individuals.

This is a thrilling prospect, especially when you couple it with the perennially seductive thought that everything is what it is owing to ways it is related to everything else. The difficulty is to get a grip on the ontology. This could be due in part to its being natural to think of relations as dependent on relata in the sense that, without the relata, there is nothing to relate. Imagine the Myth Busters setting out to subtract the cat while keeping the smile.

One possibility is that the attraction of a purely relational universe (or a universe in which relata are 'constructed' from relations) rests on an illicit move from the idea that a universe could be given a purely relational *description*—via graph theory, for instance—to the much stronger thesis that this might be *all there is* to the universe. Such a move might be especially tempting to philosophers who start with scientific formalisms and proceed to extract an ontology directly from these.

Think of functionalism in the philosophy of mind according to which all that matters are relations among 'nodes' in a system. The intrinsic qualitative nature of the nodes themselves is presumed irrelevant. Transistors could be swapped for neurons, provided only that the transistors preserve connections to other elements of the system implemented by neurons they replace. It is one thing, however, to allow that the qualitative nature of what occupies a system's nodes could vary across systems, another matter altogether to imagine that nodes are qualitatively empty, wholly constituted by relations 'they' bear to one another.

My tentative suggestion is that the giddy feeling accompanying thoughts of purely relational universes and universes in which objects are wholly constituted by powers to affect other objects, springs from a common source. In the case of purely relational universes, it is hard to see what could distinguish one relation from another. What distinguishes Holton's directly-to-the-right of relation from his directly-above relation? How do you count instances of each relation? Purely relational universes arguably lack sufficient individuative resources. The same problem bedevils scientiphicalist pure powers

universes. The real difficulty lies not in the threat of a regress, but in the fact that qualities play a central role in the identity and individuation of powers. Strip away the qualities, and it is no longer clear what, if anything, you are talking about.

4.9 Scientific Abstraction

Before turning to considerations favouring this diagnosis, let me note that proponents of qualitatively empty universes appear sometimes to reason from the fact, or alleged fact, that science is silent about the F's to the conclusion that science tells us there are no F's. This illicit line of thought calls to mind David Armstrong's 'headless woman': you move from a lack of awareness of a woman's head (the woman is a magician's assistant with a black felt bag over her head) to the thought that you are aware of the woman's lacking a head (Armstrong 1968). The question to be faced is whether we can make sense of the idea that the physical universe is bereft of qualities, a realm of pure powers. Suppose that science *were* interested exclusively in objects' powers, suppose scientific descriptions of objects concerned *only* what those objects do or would do. The issue is not whether we should accept what science tells us, but whether we should regard this as the end of the story.

Science is characterized by abstraction, what Locke calls 'partial consideration'. You might think of Newton's laws of motion, for instance, as concerning the behaviour of objects *qua* massy. The laws are meant to provide a precise account of the contribution an object's mass makes to the object's overall complement of dispositionalities. Abstraction of this kind lies at the heart of scientific theorizing and indeed of much ordinary thought about the universe. The mistake is to imagine that such abstraction gives us a complete picture.

One way to begin unraveling this topic is by considering an example deployed by Unger in his assault on scientiphicalism and 'The Mystery of the Physical'. Unger, unselfconsciously invoking Berkeley, Armstrong, Campbell, Foster, C. B. Martin, and others, contends that we could form no 'adequate' notion of a universe of material bodies lacking qualities. Any 'humanly graspable' conception

of material bodies endows those bodies with 'extensible colors'.[8] You could conceive of translucent 'extensibly red' spheres, for instance, moving about in space and interacting in various ways. Subtract the qualities—sphericality and 'extensible redness'—and nothing remains of your conception.

Unger drives the point home by inviting the reader to envision universes populated by mobile, extensibly coloured spherical particles juxtaposed with universes made up of bubble-filled extensibly coloured 'plena' in which motions of bubbles mimic particle motions. Because particle universes differ from bubble universes *only* qualitatively, scientiphicalism is unable to distinguish the one from the other. Scientiphicalism is blind to a momentous distinction. The alternative recommended by Unger is to reject the 'Denial of Qualities':

Against those who've assumed the Denial to hold, I've argued that, without conceiving a concrete reality that's Qualitatively endowed, we humans, at least, can't conceive at all adequately, any physical reality at all. . . . Once we reject the Denial, we may resolve our Mystery of the Physical . . . [thereby removing] an obstacle to our having a tolerably clear conception of a World that's the way Scientiphicalism claims our actual world to be. (137)

Only by endowing scientifically specified objects with spatially extensible qualities can we frame a 'Humanly Realistic metaphysic'.[9]

4.10 Reciprocity

Unger's 'Humanly Realistic Metaphysic' includes objects 'variously qualitied and propensitied'. But how are propensities—powers—and qualities related? Philosophers have been conditioned to think it obvious that powers and qualities must be distinct species of property: powers endow their possessors with capacities to inflict or suffer various changes; qualities are impotent. A conception of this kind would be natural for anyone accustomed to distinguishing dispositional and categorical properties.

---

8 'Spatially extensible color' (158) is a central feature of Unger's 'humanly realistic' conception of material reality. 'Spatially extensible colors' need not be actual *colours*, however, but only objects' qualitative clothing, Foster's intrinsic 'filling'.

9 Unger assumes what Foster questions: that spatiality is a paradigmatic physical characteristic.

The idea that qualities and powers are distinct kinds of property fits nicely with the view that objects' powers and qualities covary contingently. In our universe spherical objects roll, but in other universes they need not. The ease with which we concoct such imaginary universes leads Unger to describe in loving detail dozens of universes populated by dynamic, 'extensibly colored' objects (typically translucent spheres) endowed with powers of attraction and repulsion. One such universe might comprise red, blue, and yellow spheres moving about in a three-dimensional space. Blue spheres might attract yellow spheres, but not red spheres. Here, the power to attract yellow spheres covaries with extensible blueness. Blue spheres could, in other universes lack this power, possessing, instead, other powers or no powers at all.[10]

Two points are worth emphasizing here. First, the attractive and repulsive powers in Unger's imaginary universe are individuated *qualitatively*. A power to attract a yellow sphere is a power to attract a *yellow* sphere, not a *red* sphere. Second, connections between qualities and powers are contingent. Although (in Unger's imagined universe) blue spheres attract yellow spheres, blue spheres might have attracted red spheres. (As Unger would put it, there are universes in which blue spheres *do* attract red, but not yellow, spheres.) I shall argue that these points are in tension.

Notice first that, if blue spheres are equipped with a power to attract yellow spheres, yellow spheres must, of necessity, possess a reciprocal power, the power to be attracted by blue spheres. So in any universe in which blue spheres *contingently* possess the power to attract yellow spheres, yellow spheres in that universe *must* harbour the reciprocal attractive power. The result, unremarked by Unger, presents us with an odd kind of necessary correlation. Although it is contingent that blue things have the power to attract yellow things and contingent that yellow things have the power to be attracted by blue things, in any universe in which blue things have the power to attract yellow things, yellow things could not fail to have the reciprocal power to be attracted by blue things (and vice versa).

---

10 If blue spheres could have different powers in different universes, might there be 'mixed universes' in which some otherwise identical blue spheres have one power, others another? If not, *why* not? (See Blackburn 1984, 186.)

You can see the prima facie oddness of such a view by thinking of it this way. Suppose God creates a universe containing blue spheres and elects to endow these blue spheres with a power to attract yellow things. If God subsequently creates yellow things, God could not fail to give these yellow things the power to be attracted by blue things, a power they could have lacked otherwise.

You might be sceptical that powers have the reciprocal character I have attributed to them. Maybe all it takes is that blue spheres have an attractive power; yellow spheres passively cooperate. What, then, distinguishes the power possessed by blue spheres from other powers? It is a power to attract yellow things, a power to attract certain spheres, not in virtue of their size or their sphericality (blue spheres do not attract spheres of the same size that are not yellow), but in virtue of their *yellowness*. Blue spheres are 'yellow attractors'. We have a qualitative mode of individuation for a power. This, I think, is how Unger sees it. A qualitatively individuated power is contingently possessed by our imagined blue spheres, but it is a power blue spheres could have lacked, a power that could have been possessed, instead, by red spheres.

A magnetic bar has the power to attract iron filings. One way to think about such cases is to ascribe an *active* power to the bar, and *passive* powers to the filings, a distinction familiar to readers of Locke.[11] Let me suggest that this asymmetrical picture is inappropriate, a byproduct of our explanatory practices and contingencies of experimental manipulation (I shall return to this topic in chapter 6). A magnetic bar attracts iron filings in virtue of some feature of those iron filings. The attractive event is a manifestation of *reciprocal* powers possessed by the bar and by the filings. Similarly, in Unger's imagined universe, blue spheres attract yellow spheres in virtue of reciprocal powers possessed by blue spheres and yellow spheres.

One source of difficulty facing anyone trying to move from consideration of pretend 'extensibly colored' spheres to serious ontology arises from the fact that we are focusing on toy universes in which objects' powers are whimsically stipulated. In such universes anything

goes. Consider an actual power, however, the power possessed by a billiard ball to roll (in a particular way) down an inclined plane. In this case it seems obvious that the ball's manifesting this power is a mutual manifestation of reciprocal powers possessed by the ball, the inclined plane, the gravitational field in which they are located, and assorted other factors. You would want to think about powers in this way if you thought, as I do think, that one and the same power could manifest itself differently with different kinds of reciprocal disposition partner.[12] This is one important consequence of identifying powers with properties. Suppose the billiard ball possesses the power to roll in virtue of being spherical. In virtue of being spherical, the ball would make a concave circular imprint in soft clay, would reflect light so as to look spherical, would cast a particular sort of shadow. A ball's powers are intrinsic to it, but how these powers manifest themselves depends on their reciprocal manifestation partners.

When you think of powers this way, you will be thinking of them as reciprocal. In virtue of being as it is, each power would manifest itself in a particular way with particular kinds of reciprocal partner.[13] Suppose we accept reciprocity and see where it leads.

4.11 Contingency

Unger individuates powers qualitatively, thereby rejecting scientiphicalism. He insists, however, that the relation between qualities and powers is contingent. Consider again an imagined Unger-style universe featuring red, blue, and yellow spheres, and assume, as before, that blue spheres have the power to attract yellow things—blue spheres are yellow attractors—and yellow spheres possess a reciprocal power to be attracted by blue things. The contingency thesis requires that these powers be only contingently related to qualities possessed by their bearers. On Unger's view there is nothing about blueness *itself* that connects blueness to the power to attract yellow things. Blue spheres could have failed to enjoy this power. Similarly for

---

11 Shoemaker (1980, 1998) follows suit in speaking of 'forward-looking' and 'backward-looking' powers. (For a contrary view, see Davidson 1973, 64.) In fact, the distinction between a yellow sphere's having the power to be attracted by blue spheres and a yellow sphere's having the power to attract a blue sphere is elusive.

12 This is not to deny that some powers might manifest themselves spontaneously. Manifestations of such powers would appear to be uncaused. Many powers, however, evidently manifest themselves as they do only in concert with the right kinds of reciprocal partner; see chapter 6.

13 This way of thinking about powers has long been advocated by C. B. Martin; see his 2008.

yellowness: there is nothing about being yellow that ensures that its possessors would be attracted by blue things.

It is time to pay the piper, time to ask how all this is supposed to fit together. In the imagined universe, blue spheres have the power to attract yellow things and yellow things a reciprocal power to be attracted by blue things, but these powers are only contingently related to blueness and yellowness, respectively. This means that the power possessed by blue spheres could not be a power to attract yellow spheres *as such*, a power to attract yellow spheres in virtue of their *yellowness*. To think otherwise would be to turn yellowness itself into a power: in virtue of *being yellow* a sphere would be attracted to blue spheres. It must be that yellow things possess some reciprocal power that contingently covaries with yellowness. This power, and not the spheres' yellowness, is, so to speak, the target of the attractive power possessed by blue spheres. The power is one that, in the imagined universe, happens to be possessed by yellow spheres. Parallel reasoning establishes that the power possessed by yellow spheres to be attracted by blue things is a power to be attracted by something with the right power, a power that contingently covaries with blueness.[14]

Does this manoeuver help make sense of contingency? Consider the powers just described. What exactly are these powers supposed to be powers *for*? Each is a power to attract or be attracted by objects possessing the other. If contingency holds, the powers are swappable. The power possessed by blue things might be possessed in some universes by yellow things, and the power possessed by yellow things might be possessed elsewhere by blue things. Indeed the powers might belong to red and green things respectively: there are universes in which red and green things possessed these very powers and yellow and blue things possessed some *other* pair of attractive powers.

My sense is that such possibilities amount to distinctions without differences. Unger's initial thought was right: qualities figure ineliminably in the individuation of powers. A blue sphere that possesses a power to attract yellow things, what we began by calling a yellow

14 Unger distinguishes quality-directed powers ('Propensities [that] *concern only something as to the Quality of disposition partners*') from power-directed powers (212–13), but it is difficult to see how, given contingency, a power could be anything *but* power-directed.

attractor, possesses a power to attract objects in virtue of their *yellowness*. Yellow things possess, by virtue of being yellow, a reciprocal power to be attracted by blue spheres in virtue of their *blueness*. It would seem to follow that the connection between 'extensible yellowness', a quality, and a power, the power to be attracted by yellow attractors is not, after all, contingent. Indeed it appears that this power is inseparable from the nature of the quality. More generally, when you appeal to qualities in individuating powers, it appears that the pertinent reciprocal powers are powers possessed by objects in virtue of their possession of the pertinent quality.

If this sounds far fetched, it does so partly because we are dealing with fanciful, stipulated powers. We have no reason at all to think that blueness must be bound up with the attraction of yellow things, or that yellowness has a special affinity for blueness. The point I want to make could be framed conditionally, however. *Were* the universe as Unger describes it, *were* the universe such that blue spheres were yellow attractors and yellow spheres were subject to attraction by blue things, blueness would not be contingently associated with the power to attract yellow things in virtue of their yellowness, yellowness would not be contingently subject to blue attraction.

Unger regards scientiphicalism as hopeless because scientiphicalism excludes qualities, and qualities are required for the individuation of powers. If qualities and powers contingently covaried, however, qualities would cease to play this individuative role. When a blue sphere is described as possessing a power to attract yellow spheres, what the blue sphere really possesses is a power to attract something with the power to be attracted by something possessing a power contingently possessed by the blue sphere. It so happens that this power-to-be-attracted is possessed by yellow spheres, but, because the power could fail to be possessed by yellow spheres, its being a power possessed by yellow spheres in Unger's pretend universe is no part of its identity.

I fear the preceding discussion invokes just the kind of esoteric mumbo-jumbo I have inveighed against in other contexts. For that reason, it would be distressing to think that the position I am advancing in this chapter depended on arguments of this kind. I prefer to think that you could completely ignore my discussion of Unger and still find the conception of properties sketched here plausible, even compelling.

Why, then, address these matters at all? I see Unger's antipathy toward scientiphicalism as affording a way of grasping what really lies behind latent worries about the idea that all there is to the universe is a network of powers, an idea especially popular among philosophers who regard science as providing us with all the metaphysics we need. Arguments advanced by philosophers as different as Berkeley, Armstrong, Campbell, and Foster point to an important lesson, although it is not always easy to see what that lesson is. On my reconstruction, what the arguments all tell us is that the very identity of powers binds them to qualities. Far from being distinct, powers *are* qualities, qualities powers.

4.12 Qualities Unleashed

The widespread presumption that qualities are one thing and powers something else altogether, and with it, the idea that powers and qualities, at best, contingently covary, springs, I suspect, from a Humean undercurrent evident in so much contemporary philosophical thinking. Abetted by Ryle, it begins with a basic confusion of powers themselves with *manifestations* of powers. In general, a power or disposition requires for its manifestation, a suite of reciprocal disposition partners. *How* a disposition manifests itself depends both on its nature and on the nature of its reciprocal disposition partners. A billiard ball's sphericality is responsible for the ball's rolling, but only on a solid sloping surface situated in a gravitational field. This same sphericality is responsible for the ball's reflecting light so as to look sphericality and for the ball's making a concave circular impression of a distinctive sort in the carpet.

This feature of dispositions—one disposition, many kinds of manifestation with many kinds of reciprocal disposition partner—is easily missed if you imagine that subjunctive conditionals commonly used to pick out dispositions provide exhaustive characterizations or analyses of the dispositions they pick out. Distinct conditionals are presumed to pick out distinct dispositions. In identifying dispositions solely by reference to their manifestations, you might naturally be led to suppose that different dispositions must underlie different kinds of manifestation.

Suppose I am right about this. Now think about Humean considerations apparently favouring contingency. You can easily imagine a billiard ball failing to roll or failing to look spherical. This is because the ball's rolling or looking spherical requires a host of cooperating reciprocal partners. If you vary these, you will vary the way the ball's sphericality is manifested.

Another factor at work in these cases stems from the possibility of a ball closely resembling a billiard ball that does not roll down an inclined plane. One reason you can do this is that balls (and inclined planes) are complex objects. Two balls made of very different materials could nevertheless *look* and *feel* just alike, could be superficially indistinguishable. A magnetized, plastic-covered, steel billiard ball might fail to roll down an inclined plane made of iron. It is however, rather more of a challenge to imagine a ball intrinsically indiscernible from a regulation billiard ball that does not roll, but tumbles, as a cube might, down an inclined plane; reflects light so as to look cubical rather than spherical; feels cubical; makes a square-shaped impression in the carpet. If you grant that one and the same power could manifest itself differently with different kinds of disposition partner, it is much harder to hold the qualities constant while varying the powers.

Let me mention one more, related, difficulty for anyone who regards qualities and powers as distinguishable, contingently related features of objects. Suppose perception of something involves a causal connection between what is perceived and the perceiver. If you accept an ontology of powers, you will want to say that perceiving is a mutual manifestation of dispositions of us, dispositions of our perceptual systems, and dispositions of perceptual stimuli. In the case of visual perception, this might be structured light radiation, which is in turn a mutual manifestation of dispositions present in ambient light and those present in objects seen. If qualities and powers are distinct, how could anyone perceive qualities? Suppose that red spheres are equipped with assorted powers, one of which covaries with redness. It is not the sphere's *redness* that causes you to perceive something red, but a power *accompanying* redness. You do not perceive the redness of a sphere, then, but only—what?—some power contingently possessed by red things? This turns on its head the idea that qualities are observable, powers unobservable.

The moment you divorce qualities and powers, you have little use for qualities (at least until you start worrying about consciousness). I have suggested that there are good reasons to think that qualities are required for the individuation of powers by reference to what those powers are powers *for*. Given the reciprocity of manifestations this, perhaps surprisingly, encourages the view that powers and qualities are not merely contingently associated. But if powers and qualities are associated of necessity, what is the nature of this necessity, what is its basis? Is it simply a 'brute' necessary correlation? That would be hard to swallow. It would, in addition, keep qualities out of the causal picture. Qualities would be necessary, but epiphenomenal, accompaniments of powers. This, I have suggested, is a source of unwelcome difficulties and, indeed, implausibilities. You can accommodate the necessity and resolve the difficulties by *identifying* powers and qualities, turning properties into *powerful qualities.*

Philosophers are apt to regard such a view as contentious in the extreme, but I would wager that non-philosophers accept it as too obvious to bear mention. Things do what they do because they are as they are, and ways things *are* are qualities. When you cite qualities in causal explanations, when you say that the bull charged because the matador's cape was red, you are not citing features of objects you take to be correlated with their powers. The cape's redness, you think, sparked the bull's anger: in virtue of being red, the cape has the power to attract the attention of aggressive bulls.

This natural way of thinking of qualities as *themselves* powers has been compromised by a strain of Humeanism that runs through so much contemporary philosophy, and by accounts of powers that distinguish powers or dispositions from their 'categorical bases': dispositions are presumed 'higher-level', 'multiply realized' properties (again, Prior, Pargetter, and Jackson 1982). I have argued at length elsewhere against the idea that powers are 'higher-level' properties (see chapter 9 and Heil 2003a, 2005a; Heil and Robb 2003). 'Higher-level' properties collapse into their realizers. What of Humeanism?

Twentieth-century analytic philosophy inculcated a Humean picture of the universe as a default. The universe is characterized by what Santayana (1930, 110) calls 'radical contingency'. We take it as obvious that, in philosophical argument, contingencies can simply be postulated without cost. However intimately related the F's and G's

are in our experience, it seems *easy* to conceive of their failing to be so related. The burden of proof falls to those who would deny contingency. Unless you can establish that the F's and G's are related of necessity, it is reasonable to suppose that their relation is contingent. The situation resembles the situation surrounding the heavyweight boxing crown: to be dethroned, the champ must be decisively defeated.

It is time we accepted that, in serious ontology, there are no default views, no heavyweight champs. Substantive claims of contingency, no less than claims of necessity, must earn their keep. Spinoza might be right: apparent contingencies could turn out to be merely apparent. Telling a story according to which the F's and G's are related of necessity, then, need not oblige us first to provide a definitive refutation of Humean contingency. The story's plausibility depends on how well it does what we want such stories to do: provide an illuminating ontological picture that makes sense of the universe and our place in it. Against this background, I believe a conception of properties as powerful qualities has much to recommend it.

Of course the kind of necessity at play here is a kind even a lover of contingency could appreciate. Powers—dispositions—and qualities covary of necessity because powers—dispositions—*are* qualities. Really, what is at issue is not a supposedly dodgy necessary correlation among distinct existences, but the idea that powers or dispositions and qualities *could* be identical. I have suggested reasons to think, not merely that they *could* be identical, but that they *must* be identical. I admit that my reasons are not conclusive. The chief virtue of the position is not that it incorporates decisive refutations of its competitors, but that it allows us to make sense of so much.

4.13 Primary and Secondary Qualities

Berkeley attacked materialists who held that secondary qualities (tastes, sounds, colours, smells) were mind-dependent. When you try to imagine a material object altogether lacking in such qualities, you find yourself brought up short. If secondary qualities are mind-dependent, and if you cannot imagine objects lacking in such qualities, the objects must be mind-dependent.

Whatever you might think of Berkeley's argument, it is worth noting that the underlying assumption is that primary and secondary qualities alike are *qualities*. Berkeley, echoing Galileo, takes the secondary qualities to be mental. Let me propose another way to think about the distinction.

Assume that the primary qualities are qualities possessed by the fundamental things: the real properties. Assume, as well, that these qualities are powers—powerful qualities. The fundamental things might possess shape, mass, and size, for instance. And, in virtue of their possession of these qualities they do and would do various things in various circumstances. What of the secondaries? Think of secondary qualities as arrangements of the primaries. If you arrange the corpuscles in a particular way, the result is something red—something that looks red in virtue of reflecting light in a particular way. Once God has created the fundamental things and arranged them in a definite way, God has thereby created all the secondary qualities. The secondary qualities are secondary by virtue of being 'no addition of being'.

If something like this is right, you could agree with Berkeley's premises without accepting his conclusion. In depicting to ourselves any object we depict it as possessing various secondary qualities. This is because we depict medium-sized objects, and medium-sized objects are complex. When you imagine viewing a simple corpuscle, you imagine it as it appears on a photographic plate or a computer monitor, and images on photographic plates and computer monitors are complex entities exhibiting secondary qualities. More fundamentally, when you imagine what something—a corpuscle, for instance—might *look* like, you are conjuring a visual appearance, an experience of a complex entity. The corpuscle is not red, or any colour at all. But your image is as of something red, something billiard-ball like. None of this implies that secondary qualities are mind-dependent, none of it yields an argument for idealism.

4.14 Powerful Qualities

We are left with the pleasing idea that properties are powerful qualities. To be a powerful quality is to be a way some substance is. Properties are abstractions, ways particular substances are. Because

substances must be simple, lacking in substantial parts, the only properties are properties of fundamental, simple substances. What we naturally regard as properties of complex objects are not in fact properties. Or, if you like, they are properties in a relaxed sense, properties by courtesy, *quasi*-properties. Although it is true that the tomato is red and spherical, what makes this true is a particular dynamic arrangement of fundamental things. This is not to cast aspersions on tomatoes, or their shapes and colours. It is merely to locate these in an ontological picture of the universe, a picture the seeds of which can be found in everyday experience, scientific investigation, and in self-conscious philosophical reflection.

Although I have expressed a preference for the idea that properties are particulars, I have thus far left open the possibility that properties are universals. The time has come to address this vexed topic in more detail. As will become clear in the next chapter, the issues are at once simpler and more complicated than most discussions of properties admit.

CHAPTER 5
Universals

Number, when it is considered simply in the abstract or in general, and not in any created things, is merely a mode of thinking; and the same applies to all the other *universals*, as we call them. . . . These universals arise solely from the fact that we make use of one and the same idea for thinking of all individual items which resemble each other: we apply one and the same term to all the things which are represented by the idea in question, and this is the universal term. When we see two stones, for example, and direct our attention not to their nature but merely to the fact that there are two of them, we form the idea of the number which we call 'two'; and when we later see two birds or two trees, and consider not their nature but merely the fact that there are two of them, we go back to the same idea as before. This, then, is the universal idea; and we always designate the number in question by the same universal term 'two'. In the same way when we see a figure made up of three lines, we form an idea of it which we call the idea of a triangle; and we later make use of it as a universal idea, so as to represent to our mind all the other figures made up of three lines.

(Descartes, *Principles* I, §§58–9)

All things that exist are only particulars. (Locke, *Essay*, III, iii, 6)

5.1   *Comme il Faut* Philosophy

Chapter 4 provided an account of the nature of properties according to which properties are at once qualitative and dispositional: properties are powerful qualities. Questions remain. You might wonder about properties of 'abstract' entities: sets, numbers, propositions.

These seem to be neither qualitative nor dispositional. I shall address this topic in chapter 8. Before making the leap to *abstracta*, however, let me turn to another question that has been simmering in the background, the question whether properties are *universals*.

Locke's and Descartes's properties are *modes*, 'abstract particulars'. The redness of this apple and the redness of a separate but indistinguishably coloured apple are similar but numerically distinct ways each apple is. The alternative is to regard properties as universals, sharable, 'repeatable' entities. Easy to say; harder to understand.

Today, philosophers who look with favour on properties more often than not take it as given that properties are universals. The alternative, that properties are particulars, strikes these philosophers as next to unintelligible. Peter Van Inwagen, for instance, observes that

a universal is supposed to be a thing that has 'instances', and a property—redness, say—is a species of universal. One might suppose that 'instances' of the property redness would be things like red apples. According to those who believe in 'property instances', however, this red apple and that red apple are not two instances of the property redness; there are, nevertheless, two instances of redness in the vicinity (so to speak) of the two apples. Those who believe in these items call them 'the redness of this apple' and 'the redness of that apple'. These items, these 'particular rednesses', are supposed to be two distinct things even if the two apples are of exactly the same shade of red. (Van Inwagen 2007, 38)

Van Inwagen proceeds to heap scorn on the very idea of property instances regarded as particularized properties.

Belief in 'property instances' must represent some perennial tendency of the human mind, since it has arisen independently in several philosophical traditions. I can find no trace of it in my own mind, however. Consider two apples that are exactly alike in every respect. To my mind, saying that 'the redness of this apple' and 'the redness of that apple' are names of two distinct things makes about as much sense as saying that 'the diameter of this apple' and 'the diameter of that apple' are names of two distinct things. (Van Inwagen 2007, 38)

Jerrold Levinson echoes this sentiment in a discussion of 'tropes', the currently fashionable label for Van Inwagen's property instances: 'tropes cannot be particularized properties, since the notion of a particularized property, or condition, is simply an oxymoron.

Hence there are no tropes' (Levinson 2006, 564). And the Macdonalds, Cynthia and Graham, speak for many when they note that 'in standard philosophical usage, a property is construed as a universal and an instance of a property is not a trope of that universal but a thing that has (instantiates, exemplifies) that property' (MacDonald and MacDonald 2006, 547).

Such comments would give pause to anyone who has struggled to get first-year philosophy students to comprehend universals. Is it that students who go on in philosophy eventually catch on? Or is it that they simply learn to talk the talk and repress their initial misgivings? Becoming a licensed philosopher requires learning many things, including what questions *not* to ask, what issues *not* to press, what battles *not* to fight. Postgraduate training in philosophy requires adeptness at repression.

Anthropological considerations aside, what is most striking about these comments is the extent to which they distort philosophical tradition. The suggestion is that only philosophical cranks and outliers have taken properties to be particulars. Serious philosophers would not give such a conception the time of day. The idea that properties are universals has bubbled to the surface as the default view, the view to be accepted in the absence of compelling reasons against it.

Philosophy is in trouble when philosophers rely for their inspiration on what everyone knows, what is accepted by all right-minded—meaning like-minded—professionals, what is doctrinally *comme il faut*. In philosophy, everything is up for grabs; there are no default views, no heavyweight champs. The thesis that properties are universals is a substantive ontological proposal on all fours with the thesis that properties are particulars. In neither case is there any question of a decisive refutation of competitors. Fundamental ontology is a matter of give and take, a matter of weighing costs and benefits. On this score, I suspect that universals as philosophers today commonly understand them do not fare well. My aim here, however, is not to disparage universals but to explain why, despite various temptations to regard properties as universals, you might find the doctrine that the universe, indeed, any imaginable universe, is a universe of particulars so much more appealing. To set the stage, I shall begin with a brief historical survey.

5.2 Historical Reminder

To hear philosophers today tell it, the belief that properties are not universals but particulars has a spotty philosophical history, a thought embraced here and there by the odd philosopher, the philosopher outside the mainstream, the philosopher cut off from philosophy's deeper currents. *Real* philosophers, the great ones, would see through the idea quickly and move on.

To view the history of philosophy this way, however, is to view it in manifestly revisionist terms. Start with Plato, the original friend of universals. Plato left room for, indeed insisted on, the 'moving forms', instances of universals in the objects. There is Socrates, there is the universal whiteness, and there is Socrates' whiteness, whiteness *in* Socrates. Aristotle might or might not have embraced universals, but he very definitely believed in 'individual accidents': the sphericality of this ball, the horseness of this horse.[1]

By the medieval period, many prominent philosophers flatly rejected universals, regarding properties as particulars: modes or accidents. The list reads like a who's who of medieval philosophy: Boethius, Avicenna, Anselm, Abelard, Averroës, Aquinas, Scotus, Ockham, Buridan, Suárez. This attitude carries over into the EnlightEnlightenment. Descartes, Spinoza, Leibniz, Locke, Berkeley, Hume, and Kant accepted properties but rejected universals. Not until the nineteenth century did universals make a comeback of sorts, and even then Husserl, and the early Russell, resisted the trend. In the twentieth century, G. F. Stout, John Cook Wilson, D. C. Williams, P. F. Strawson, and Wilfrid Sellars embraced particularized properties. This is hardly a catalogue of fringe figures.[2]

I mention all this, not in hopes of currying favour for the view that properties are particulars, but merely to encourage doubts concerning

1 For discussion of Plato, see Demos (1948) and Morrison (1977); on Aristotle see Sellars (1957) and Albritton (1957); Mertz (1996, 83–117) addresses both. Although I shall not attempt to make it, a case could be made for Aristotle's embracing a conception of universals close to the one to be spelled out below.

2 See Mertz (1996, 83–162) and Bacon (2008). Williams (1953, 189–91; 1966, 106–8) cites—in addition to Plato, Aristotle, Locke, Hume, Descartes, Spinoza, and Leibniz—James, Santayana, the Cambridge realists, A. C. Benjamin, H. W. B. Joseph, Dickinson S. Miller ('R. E. Hobart'), William Savory, Wilfrid Sellars, and E. B. McGilvary as countenancing particularized properties.

the prevailing doctrine that conceptions of properties as particulars are found only in philosophical backwaters, promoted by poseurs and eccentrics. Many of the greatest, most venerated philosophers regarded universals with suspicion; most regarded universals as unintelligible. These philosophers found it altogether natural to think of properties as particulars, and took themselves to have sensible reasons for doing so.

5.3 Terminological Interlude

Perhaps I have said enough to convince you that a conception of properties as particulars deserves a hearing. Before diving in, let me pause long enough to comment on a potentially distracting terminological matter. One reason they are difficult to trace historically is that, unlike universals (or 'forms'), particularized properties have flown under a multitude of banners.[3] Plato's 'moving forms' became, for Aristotle, 'individual accidents', and for medievals, 'modes'. G. F. Stout (1921, 1936) and Keith Campbell (1981, 1990) call them 'abstract particulars'; Gustav Bergmann (1967), 'perfect particulars'; Nicholas Wolterstorff (1970), 'cases'; Guido Küng (1967), 'concrete properties'; Gareth Matthews and Marc Cohen (1968), 'unit-properties'. Van Inwagen, in the passages quoted earlier, follows a number of philosophers in speaking of 'property instances'.

Nowadays, the philosophical community favours the term 'trope', a label originally proposed by D. C. Williams.[4] Williams cites one of his predecessors at Harvard, George Santayana, as the inspiration for his use of the term:

Santayana . . . used 'trope' to stand for the *essence* of an *occurrence* . . . ; and I shall divert the word, which is almost useless in either his or its dictionary sense, to stand for the abstract particular which is, so to speak, the *occurrence* of an *essence* (1953, 7; 1966, 78).

---

3 The list here owes much to Armstrong (1989, 113). Mertz (1996) and Bacon (2008) catalogue additional labels.

4 Williams's best-known discussion of tropes occurs in 'On the Elements of Being', which originally appeared in two installments in the *Review of Metaphysics*, 1953, and was subsequently republished in a modestly revised form in Williams's *Principles of Empirical Realism*, 1966. Citations are provided for both.

Santayana? Santayana's *The Realm of Matter* (1930) includes a chapter, 'Tropes', in which he makes clear that he means the term to denote 'essences' or universals as opposed to their particular 'occurrences'. Williams's motive for inverting Santayana's usage is less than transparent. This is especially so in light of the fact that Williams's own preferred ontology reflects central features of Santayana's (see below and Williams 1954).

We are left with a pint-sized literary puzzle. Williams adopts Santayana's term 'trope', but deploys it self-consciously to mean the opposite of what his Harvard precursor meant. Why? The obvious possibilities—that Williams was unaware of traditional labels or that he was poking fun at Santayana, for instance—appear unlikely for several reasons.

(1) Williams was certainly familiar with (and indeed explicitly mentions) terms used by other philosophers—Aristotle's 'individual accident', the medieval 'mode', Stout's 'abstract particular'—that meant what he seemed to mean by 'trope'.
(2) Williams understood Santayana's use of the term.
(3) Williams's ontology closely resembles Santayana's.
(4) Williams is a perceptive and subtle wielder of language, not given to lexicographical carelessness or gratuitous flouting of tradition.

What then might account for Williams's choice of terminology?

In reflecting on his use of 'trope' to mean what others meant by 'abstract particular', Williams makes the following observation:

That the category of abstract particulars thus indicated conforms to the logic of whole and part, or the so-called calculus of individuals, that they have logical sums and products, and so forth, and that being by definition finer or lesser parts than the concreta in which they occur they are in an important sense the 'elements of being', I once argued in print in the *Review of Metaphysics* 2, where I called them 'tropes', which has a nice historical connection with the Latin 'modes'. (1959, 4)

This connects tropes with modes or 'ways', but what 'historical connection' does Williams have in mind?

The *Oxford English Dictionary* tells us that a trope is 'a figure of speech which consists in the use of a word or phrase in a sense other

than that which is proper to it; also, in casual use, a figure of speech; figurative language.' So, whatever else he is doing, Williams is using 'trope' as a trope!

The *OED* does not stop there, however. In addition to its familiar use, there is an 'obs. rare' use in logic to mean 'mood' or 'mode'—as the mood of a syllogism.[5] There is, then, a connection, albeit an 'obscure' one, between Williams's 'trope' and the traditional 'mode'. But why take the trouble to seek out a term with this tenuous relation to a more familiar (or at least more traditional) term? Why not, for instance, settle for 'mode', thereby refusing to contribute to a frustrating proliferation of terms?

I believe Williams was after a term that lacked potentially misleading connotations and historical associations. He notes explicitly that, by 'trope', he means what others—Stout, for instance—had meant by 'abstract particular' (1953, 174; 1966, 92). He sees a problem, however, with the term 'abstract' which had, he thought, come to be used in a way that diverges from its 'true meaning': '*partial, incomplete*, or *fragmentary*, the trait of what is less than its including whole' (1953, 15; 1966, 85; see also Williams 1959, 5).

The many meanings of 'abstract' which make it repulsive to the empirical temper of our age suggest that an abstractum is the product of some magical feat of mind, or the denizen of some remote immaterial eternity. . . . Logical philosophers proclaim their 'renunciation of abstract entities' without making clear either what makes an entity 'abstract' or how one goes about 'renouncing' an entity. (1953, 14; 1966, 84)

This billiard ball's sphericality is abstract, not in the sense of its falling short of being fully 'concrete', but in the sense that its 'separation' from the ball is a matter of abstraction, Locke's 'partial consideration', a matter of our capacity for considering the ball's shape as distinct from considering its colour or considering the ball itself. The idea is not that, in abstracting, we *manufacture* abstracta; abstraction (partial consideration) is what enables us to *apprehend* abstracta. Because philosophers have too often lost sight of this 'root meaning' and come to think of abstract entities as non-concrete, existing

incorruptibly apart from or 'outside' space and time, Williams thinks the label 'abstract particular' is apt to sow confusion.

Fair enough. But why should Williams shy away from 'mode', or 'individual accident', or even Santayana's 'occurrence'? Here is one possibility. With few exceptions, philosophers who have used such terms have invariably embraced two-category, substance–attribute ontologies. A mode, for Aquinas, or Descartes, or Locke, is a way, a particular way, some substance is. Williams wants a one-category ontology. The role of substances is to be filled by spatio-temporally 'concurrent' particularized properties. What is called for is a new term that lacks inconvenient, even embarrassing, historical ancestors. 'Trope' serves this purpose while at the same time connecting Williams's ontology with Santayana's. I shall have more to say about this connection later.

In speaking of particularized properties, I vastly prefer 'mode' to 'trope' (as does E. J. Lowe; see his 2006) for two reasons. First, thanks to Williams—and to Keith Campbell (1981, 1990), Peter Simons (1994), David Robb (1997, 2005), Anna-Sophia Maurin (2002), and others—most self-described 'trope theorists' today are also 'bundle theorists', conceiving of objects as bundles of 'compresent' tropes.[6] Williams himself defended an austere ontology of tropes and part–whole—mereological—relations. I side with Martin and with Lowe in thinking of particularized properties as particular ways—from the Latin *modus*, way—particular substances are.[7] Following Locke, and indeed all those historically prominent philosophers who discussed modes, I have argued that substance and property are correlative categories: each requires the other. This places me in the camp of traditional substance–attribute theorists.

Second, by thinking of properties as modes—ways—it is easier to see them as dependent items the identity of which is partly determined by substances of which they are modes. Socrates' whiteness is *Socrates'* whiteness, a particular way *Socrates* is. There can be no

---

6 Most, but not all. Armstrong (1989, 136) argues that 'we do better, with Locke and C. B. Martin, to hold the trope view in a substance–attribute form'. See Martin 1980, 2008; Heil 2003a.

7 The conception of properties as 'ways' made explicit by talk of modes is common to most, but not all, conceptions of particularized properties prior to the twentieth century. Hume is a notable exception. In recent years talk of ways was revived by Jerrold Levinson (1978, 1980) and Daniel Seargent (1985) and endorsed provisionally by D. M. Armstrong (1989, 96–8).

question of Socrates' whiteness migrating to Simmias or Xantippe. On such a conception, a conception defended in chapter 2, properties are not assembled to form objects; objects are not made up of properties. Objects—substances—are the basic particulars. Every substance is various ways. These ways are modes.

You can see why a philosopher attracted to the idea that properties are particulars might be attracted as well to a bundle theory. If you thought properties were universals, you would need some way of introducing particularity into the universe (see chapter 2). Bundles of universals would seem just to be complex universals. You can achieve particularity by introducing substances as 'non-repeatable' items that instantiate or exemplify universals. If properties themselves are particulars, however, the particularizing role of substance is no longer required. Why not economize and settle exclusively for properties? Why not opt for a Spartan one-category ontology?

G. F. Stout (1921, 1936) provides a nice example of this line of thinking. Stout regards the conception of objects as substances—substrata—equipped with properties, hopeless. This pushes us toward a bundle theory of objects. As the reflections above suggest, however, bundles of universals amount to composite universals, not particular objects. One way to think about particularity is by reference to the principle of the identity of indiscernibles: indiscernible particulars can be distinct. In contrast, indiscernible universals are, of necessity, identical, one and the same universal. If you are attracted to the idea that objects are bundles of properties, if you are suspicious of 'substrata', you will want properties themselves to be particulars.

One difficulty with this option lies in understanding properties as parts that *add up to* objects. Some philosophers, Descartes for instance, have regarded the very idea as self-evidently incoherent. If there is sphericality, there must be something to be spherical, some substance extended in the spherical way; if there are thoughts, there must be something that thinks, some thinking substance.[8] The conception is of properties as items the identity of which depends on the substance that bears them.

8 If you thought that sphericalities are *themselves* spherical, that the sphericalities are what is spherical, you would be turning sphericalities into substances, what Armstrong, following Ayer, calls 'junior substances' (Armstrong 1989, 115; see also Descartes 1641, 176; Garcia 2015).

Philosophers who debate the pros and cons of Williams's tropes too often appear oblivious to their revolutionary nature. Williams's view is thought to be unorthodox chiefly because he treats properties as particulars. Williams is in the majority party in this regard, however. What makes Williams a revolutionary is his embracing a one-category ontology, an ontology bereft of substances. *This* is the move that Williams's philosophical ancestors would have regarded as flatly unintelligible.

These observations are not meant to constitute an argument against the bundle theory. I provide them only to dispel the misguided impression that a commitment to abstract particulars, or modes, or tropes is inconsistent with the acceptance of a traditional substance–attribute ontology. Although I have argued at length that properties require substantial bearers, nothing in what follows requires taking a stand on the matter, taking a stand on whether properties are parts of objects or ways particular substances are. I shall omit this qualification and continue to speak of properties as 'ways', however, in the interests of expository simplicity.

5.4   Similarity and Identity

We have on the table a conception of properties as modes, particular ways particular objects are. Modes, like substances, are numerically non-repeatable particulars. Socrates' whiteness and Simmias' whiteness, although exactly similar perhaps, are distinct whitenesses. Socrates and Simmias could both have 'the same' mass, *m*; but when they stand together on a scale it registers 2*m*. (And two balls with 'the same' diameter, *d*, would not both fit into a square box the sides of which were *d*.) Where does this leave us with respect to the venerable problem of 'the one over many'? We want to say of Socrates and Simmias, distinct particulars, that they have something *in common*, that they *share* a colour, that they have *the same* colour. This suggests to impressionable minds that there is some one thing common to Socrates and Simmias (and anything else with *the same* colour).

Thoughts of this kind provide one source of inspiration for the idea that properties are universals. There is a single universal whiteness, and there are endless (actual and possible) *instances* of whiteness.

Socrates and Simmias are *the same* colour, they *share* a colour, they have a feature, colour, *in common* because they, Socrates and Simmias, are each instances of the selfsame universal, whiteness. Socrates and Simmias each *instantiate* or *exemplify* whiteness.

A philosopher who takes properties to be modes and rejects universals must say that Socrates' and Simmias' having 'the same' colour is akin to two debutants arriving at the ball wearing 'the same' dress, or a son's having his father's nose (Williams 1953, 5) where sameness is not a matter of self-sameness, identity, but of more or less exact similarity. Socrates and Simmias' 'sharing' a colour, on such a view, is not at all like their sharing an umbrella; and their having a colour in common is not analogous to their joint ownership of a vacation cottage.

Detractors of modes are apt to weigh in here and point out that, whereas a believer in universals can explain similarities as rooted in identities, a proponent of particularized properties must appeal to 'brute', 'ungrounded' similarities. Everyone needs identity, self-sameness, but a denier of universals needs, in addition, brute similarity.

This line of attack is, or certainly ought to be, unpersuasive. The suggestion is that identity, so to speak, comes for free, but 'ungrounded' similarity is an unseemly ontological whisker vulnerable to Ockham's Razor. Consider the similarity relation, however, in particular, consider precise or perfect similarity. Suppose that Socrates' whiteness and Simmias' are precisely similar. Must we imagine that this similarity is grounded in some deeper, more fundamental feature of Socrates, or Simmias, or their respective whitenesses? Or is it rather that, if you have Socrates' whiteness and Simmias' whiteness, you *thereby* have these whitenesses being perfectly similar? Similarity, no less than identity, is an *internal* relation, a relation founded on its relata: if you have the relata, you *thereby* have the relation. To regard similarity, but not identity, as ontologically weighty is to mistake an internal for an external relation. (Chapter 7 is devoted to a discussion of relations.)

This mistake is all too common, chiefly because it is easy to misidentify the relata. Think of an internal relation as one in which, given the relata, you have the relation (see, e.g., Moore 1919). You could have Socrates and Simmias without their being similar

colourwise. God could have created Socrates and Simmias without making it the case that Socrates and Simmias are similarly coloured. But, in this example, Socrates and Simmias are the relata only derivatively. If Socrates and Simmias are similarly coloured, this is because Socrates is coloured in a particular way and Simmias is coloured in a particular way, and these two ways of being coloured are similar. If God creates objects coloured in just these ways, God has thereby created similarly coloured objects.

The idea that the invocation of unreduced 'brute' similarities would in some fashion be ontologically profligate is a red herring. Even if this were not the case, however, it appears likely that a proponent of universals needs 'brute' similarities no less than someone sceptical of universals. This can be seen if you consider *im*perfect similarities.[9] We allow that objects could be imperfectly similar by virtue of possessing imperfectly similar properties. This is arguably different from imperfectly similar objects' possessing some perfectly similar properties and some dissimilar properties. Imagine three billiard balls, one red, one green, one orange. We recognize that the red and orange balls, while not perfectly similar, are *im*perfectly similar, more similar to one another than either is to the green ball. Red, we think, is more similar to orange than to green. On the face of it, this kind of attenuated similarity could not be grounded in identity. If that is so, a proponent of universals will need bare similarities: the universal orange, although not perfectly similar to the universal red, is more similar to the universal red than either is to the universal green.

One way round this difficulty—if you think it *is* a difficulty—is to insist that imperfect similarities are grounded in imperfect or partial *identities* (Armstrong 1989, 103–7). Suppose that red, green, and orange are, appearances notwithstanding, 'complex properties'. Red comprises properties [A,B,C,D,E], orange [C,D,E,F,G], and green [G,H,I,J,K]. You could then say that the imperfect similarity of red to orange stems from the identity of some, but not all, of the components of these properties.

Before taking up this line of thought, it is worth pausing briefly to consider the notion of a *complex property*, a property made up of other properties. I have argued that property bearers, substances, must be

---

9 The argument sketched here is developed in Heil (2003a, 151–68); see also Armstrong 1989, 136–8.

simple. Suppose that is so. Might a simple substance possess a complex property? If a property is a way a substance is, a complex property would seem to be a substance's being more than one way. But that is just to say that the substance has more than one property. In fact, this point would hold even if I were wrong about substances, even if property bearers could be complex: a complex property would seem to be a complex of properties. Thoughts of complex properties stem in part from the fact that objects can be described by means of complex predicates (or syntactically simple predicates 'analysable' into complex predicates). It would be a mistake, however, to imagine that a syntactic feature of predicates must mirror an ontological feature of objects to which the predicates apply.

So one source of the thought that some properties are complex might be an expression of the latter-day tendency to read ontology off language. Another source is more innocent. We have in play both relaxed and ontologically serious conceptions of properties. Ascribing properties to complex objects—to tomatoes, for instance—is unobjectionable, provided only that we recognize that truthmakers for such ascriptions are interrelated arrangements of the fundamental things. Truthmakers for such ascriptions *will*, typically anyway, be complex. Truthmakers for 'the tomato is red' or 'the tomato is spherical', are complex indeed.

Even allowing the possibility of genuinely complex properties, however, the idea that imperfect similarities can be explained by invoking perfect partial identities falls flat in the case of imperfect similarities among *simple* properties. If there are, or could be, imperfectly similar simple properties (and why not?), reduction of similarity to identity is not in the cards. You might doubt that there could be such properties. It is, however, no less doubtful that every case of imperfect similarity among properties is grounded in an imperfect identity. The example above of red, orange, and green, is at best a schematic suggestion of how it might be with colours. An ontologically candid proponent of the thesis that imperfect similarity is grounded in imperfect identity would admit as much (see Armstrong 1989, 103–7).

Whether proponents of universals require 'brute' similarities, similarities not grounded in identities, is not something that ought to incline you one way or another in choosing between universals

and modes. Properly understood, similarity relations are founded on internal relations among properties. And what goes for similarity goes, as well, for *dis*similarity. Dissimilarity must be 'brute', When God creates the substances and endows them with properties, God *thereby* creates all the similarities *and* dissimilarities. Similarity and dissimilarity, alike, are ontologically recessive.

5.5 Costs and Benefits

Ontological theses are assayed, not by measuring them directly against reality, but by considering their relative power. One thesis bests another when it proves more adept at making sense of our experiences of the universe in light of our most promising scientific theories. What of parsimony, Ockham's Razor? Aren't simpler, more parsimonious ontologies to be preferred? Some find it comforting to think of parsimony as a tie-breaker: other things equal, the more parsimonious theory is to be preferred. The trouble is, things are never equal. Parsimony figures in the endgame, not at the outset of theorizing. Parsimony wielded as a theoretical constraint is a strait-jacket. The question here is whether an ontology of particulars (tropes alone or substances and modes) could accomplish what an ontology of particulars plus universals could accomplish.

You might doubt that it could. The thought that the sciences are engaged in a project of uncovering universals appears promising. Universals would seem to provide a tidy explanation of natural law and of the regularities we find in nature.[10] Hume wondered how it could be reasonable to expect similar objects to behave similarly. A proponent of universals has a ready response: similar objects behave, or would behave, similarly because they encompass identical elements.

Suppose this ball rolls owing to its sphericality and suppose that *this* ball's sphericality is distinct from *that* ball's sphericality. Why imagine *that* ball's sphericality would dispose *it* to roll? If sphericality were a universal, this question could not arise: this ball's sphericality *is* that ball's

---

10 Recent proponents of such a conception include Armstrong (1983, 1997), Brian Ellis (2001), E. J. Lowe (2006), and Alexander Bird (2007).

sphericality. If this ball's sphericality is in any way responsible for its rolling, and if that ball possesses the *very same* sphericality, then it is no wonder that it, or *any*, ball rolls or would roll.

This elegant, one-step solution to Hume's problem is not available to a philosopher who regards properties as particulars. This ball's sphericality and that ball's sphericality are distinct ways distinct objects are. The same holds for the charge and the mass of individual electrons. How could such unrelenting particularity yield the kind of generality that characterizes scientific theorizing?

Reflect first on how a proponent of universals might be thinking about the relation between universals 'instantiated' by objects and laws of nature. You could follow Armstrong (1978, 1989, 1997) and embrace an *externalist* conception of laws. Laws, according to Armstrong, are 'second-order universals'. Suppose F and G are universals. There might be a law, $N(F,G)$, to the effect that the F's necessitate the G's. In fact, matters will need to be considerably more complicated. All sorts of additional factors might be required for the necessitating of a G by an F, and all sorts of factors might intervene to inhibit or block the bringing about of an F by a G. Let us pretend what is probably false, however, that such things complicate the picture without changing its fundamental character.

A conception of this kind leaves open the possibility that laws are deeply contingent. We discover that the F's necessitate the G's, but we recognize it could have been otherwise. In the idiom of 'possible worlds', there are worlds in which there are F's and G's, but the F's fail to necessitate the G's. In these worlds, there is no higher-order universal linking F and G. (Permit me, in passing, to note a contrast between (a) the ease with which such thoughts can be formulated using the jargon of possible worlds and (b) the difficulty in saying what truthmakers for these thoughts might be, what their being true would tell us about the universe.)

A rather different approach to laws builds them—or, more accurately, truthmakers for formulations of laws—into 'first-order universals': it is of the nature of the F's, *qua* F's, to necessitate the G's, and of the nature of the G's, *qua* G's, to be necessitated by the F's (see, for instance, Ellis 2001; Bird 2007; see also chapters 3 and 6 in this volume). On this conception, properties are *powers*. A property's identity is bound up with the contribution the property would make

to what its bearers do or would do. 'Nomological necessity' collapses into 'metaphysical necessity'. Contingency survives, if at all, in the possibility that different universes might include different kinds of property.

An externalist conception of laws makes laws out to be entities in the universe—second-order universals—in addition to propertied objects. In creating the universe, God creates the objects and first-order universals, then adds the laws. An *internalist* conception, a conception that regards properties as powers, encourages the thought that laws are more aptly regarded as linguistic items: equations, formulae, or generalizations that are meant in effect to codify the contribution made by particular properties to the dispositional makeup of their possessors.[11] Newton's law of universal gravitation, for instance, expresses the contribution mass makes to what objects do or would do—how objects would affect one another—*qua* 'massy'. To a first approximation, externalists think of laws as governing objects and holding under 'ideal' circumstances; internalists think of objects as self-governing and law statements as attempts to distill the contribution particular kinds of property make to objects' capacities.

What happens to laws if properties are particulars, modes? In that case, it is hard to see how an externalist conception of laws—laws as higher-order properties—could get off the ground. The notion of a higher-order mode or trope is difficult to credit. A property of a property, regarded as a particular way a particular property is, would seem just to be the property itself: a way a way is, would be the way itself. If properties are modes, it is much more natural to regard the properties *themselves* as powers.[12] And, again, where does this leave the laws?

If properties are modes, not universals, similarity stands in for identity. If properties, modes, are powers, similarly propertied objects will be similarly empowered. Here, identity fares no better than perfect similarity in grounding the kinds of generalization important

11 A conception of this kind is associated with Nancy Cartwright; see her 1989, 1999. See also Chakravartty (2007) and Bird (2007), although Bird holds that laws 'supervene' on propensities, which he regards as transcendent universals.

12 A reminder: to say that a property is a power is *not* to say that this exhausts its nature, not to say that properties are *purely* powers. Chapter 4 supplies reasons to regard properties as powerful *qualities*.

in scientific theorizing and everyday judgement. As I shall argue presently, this fits comfortably with an account of universals propounded by Williams, an account that finds modes (or tropes) pregnant with generality. In fact, it looks as though *any* advantages thought to be provided by universals will be matched in an ontology of modes. This, coupled with Williams' suggestion, a suggestion that echoes Aquinas, Descartes, Locke, and countless others, that *thoughts* of universals must be understood as being made true, not by 'general entities', transcendent or immanent universals, but by particulars, provides all the reason any neutral party could need to abandon an ontology of universals. Particularity rules: the deep story about the spatio-temporal universe is that 'all things that exist are only particulars'.

5.6 Williams on Universals

Williams, a principal source of current interest in particularized properties, follows Stout in distinguishing conceptions of properties as universals from conceptions of properties as modes or tropes by invoking the principle of the identity of indiscernibles.[13] Where universals are concerned, the identity of indiscernibles holds: if universals F and G are indiscernible, F *is* G. Indiscernibility—perfect similarity—in the case of tropes, in contrast, is compatible with distinctness.

Williams finds the idea that universals exist in a 'Platonic' realm, outside space–time, deeply unsatisfying—as, in truth, it would have been to Plato (Plato's heaven is not the logician's). If properties are universals, if properties make perceptible differences to concrete spatio-temporal objects, it is hard to see how they could be situated 'outside' the objects. Indeed, if you separate universals from objects that 'instantiate' them, you will need representatives of the universals in the objects to do their dirty work. In Plato's case, these are the 'moving forms'. In general, however, philosophers who regard

properties as universals have disdained particularized property instances.[14] This encourages the view that universals are 'immanent': universals are present 'in' objects that instantiate them. The whiteness of Socrates and the whiteness of Simmias are the selfsame universal whiteness. Whiteness is not a 'scattered' entity, however; the two whitenesses are not *parts* of the universal, they *are*, each of them, the universal itself in its entirety. Thus we have the view, most ably defended by Armstrong, that a universal is 'wholly present' in each of its instances: *universalia in rebus* (Armstrong 1975, 1978, 1989, 1997).

The difficulty most of us feel when confronted with this idea is in grasping what it could possibly *mean*. You can understand the *words*, and you can understand that they paint a picture of spatio-temporal entities that could be wholly present in distinct places at once. But what is this picture meant to be a picture *of*?

Williams suggests one possibility. Start with the tropes and consider classes or sets of exactly resembling tropes. These classes or sets will be precisely coextensive with instances of the corresponding universal—where 'instance' denotes, not the object 'instantiating' the universal, but the universal-in-the-object. Many things are spherical. Suppose the individual sphericalities are modes or tropes. Now consider the class or set of these individual sphericalities. This class or set would be a kind of functional equivalent of the universal sphericality. To say that *this* sphericality and *that* sphericality are one and the same sphericality might be, on such a view, to say that both sphericalities are members of a set of precisely resembling sphericalities, each of which is, to be sure, *wholly* spherical, no sphericality is missing from it. Such resemblance classes of tropes are what Armstrong (1989, 122) calls 'ersatz universals'. Speaking of the universal sphericality, then, might be understood as an oblique way of speaking of the class, or set, or community of particular sphericalities.

The trouble is that this appears not at all to be what proponents of universals take themselves to mean when they speak of the universal sphericality as wholly present in each of its instances. Believers in immanent universals evidently have something much more robust in mind. Williams offers a subtle diagnosis.

<hr/>

13 See Williams 1959. Although I much prefer 'mode' to 'trope', in discussing Williams, I shall defer to his terminology and speak of tropes.

14 Plato and Lowe (2006) being notable exceptions; perhaps Aristotle was as well. Williams (1959, 6) notes that 'the Nayaya philosophy of India' might have countenanced both universals and tropes.

Assume a trope ontology, and imagine that you are inspecting two balls, more particularly you are attending to the sphericality of each ball. In perceiving the first ball's sphericality, you might naturally describe what you perceive by saying 'this is sphericality'. Turning to the second ball, you might think, 'this is [sphericality] *too* ... the whole entity all over again' (Williams 1959, 8). Such thoughts call attention to an abstract entity, a *characteristic* of each ball, rather than the balls themselves. When we consider objects of a particular kind— the two balls, for instance—we accept a principle that '*a* is identical with *b* if and only if every part of *a* is a part of *b* and conversely' (1959, 8). Identity in this sense, strict identity, 'entails but is not entailed by exact resemblance' and applies to particular balls. When you 'abstract' a ball's shape, however, when you engage in Locke's 'partial consideration' and observe that distinct balls have *the same* shape, you employ a notion of 'identity which *is* just exact resemblance' (1959, 8). This is the sense of identity in play when you think of a father and son as having the same nose. The noses are *the same* in the sense of being exactly similar.

Williams's suggestion is that 'universals are not made nor discovered but are, as it were, acknowledged by a relaxation of the identity conditions of thought and language' (8). Williams draws a parallel with 'temporal parts' of concrete particulars:

Similar relaxations occur in our treatment of ordinary proper names of concrete particulars, especially in the common idiom which, innocent of the notion of temporal parts of a thing, finds the whole enduring object, a man or a stone, in each momentary stage of its history. For here and now, we say, *is* the person called 'John', not just part but all of him, and now again is the same 'John', all present at another instant, though in strict ontology the 'John' of today is a batch of being as discrete from the 'John' of yesterday as he is from the moon. The relaxation of conditions which acknowledges universals, however, ... is much more firmly seated in the facts of language and its object than any other I know. (1959, 8–9)

The thought is not that universals are linguistic contrivances, but that we employ a linguistic contrivance, 'generization', to mark off kinds: kinds of object, kinds of property. These kinds are not exotic 'general' entities; they are the individuals considered as ways things are or could be. This is evidently Locke's view (*Essay*, III, iii), and, I believe,

of many of Locke's scholastic predecessors. Reality is uniformly particular; general terms designate, not general entities, but particular entities falling under the term. Generality stems from our capacity to think indifferently about members of classes of similar particulars, where the particulars can be substances or, as is more common, properties.

That universals are determined by a 'weaker' identity condition than particulars does not even mean that they have an inferior or diluted reality. A tabulation of universals is just one way of counting, as it were, the same world which is counted, in a legitimately different and more discriminating way, in a tabulation of particulars. (Williams 1959, 9)

One way to put Williams's point might be to say that truthmakers for assertions concerning kinds or universals are *cases*, particulars through and through. What of classes and sets? Does the universe include, in addition to particulars, classes or sets? Again, truthmakers for claims about classes or sets are just the particulars (I shall have more to say on this below).

We are left with what Williams dubs 'the trope-kind theory'.

As there is nothing in anything which is not either a trope or resolvable into tropes, so every trope, of whatever level of complexity, manifests its universal or kind. Generization, moreover, does not even stop short of concreteness, and does not therefore in the least depend upon *de facto* similarity or the recurrence of kinds. That is, having a general readiness to contemplate, by the right quirk of attention or description, either the case or the kind of any given occasion, we can identify a universal once for all in a single instance, only conceiving *ipso facto* that it is capable of other instances. (1959, 10)

Williams, note, regards trope-kind theory as a flat out, albeit 'modest', realism about universals,

in as much as it holds that universals are real entities, and it is an immanent realism in as much as it holds them to exist *in rebus*—to be present in, and in fact components of, their instances. To make plain the sense in which it holds that an abstract universal is 'in' a concrete particular we need only make explicit the analysis of predication, characterization, or instantiation which has been barely implicit here all along. That Socrates is wise, i.e., that he is an instance of Wisdom, which is an 'instantiation' or 'characterization'

in the full sense, is sufficiently expanded in the formula that the concrete particular Socrates 'embraces' [an] abstract particular (trope) which 'manifests' Wisdom. (1959, 10)

5.7 'Painless Realism'

Could this be right? Could a 'trope-kind theory' give us a *realism* about universals? That will depend on what a 'realism about universals' amounts to. If realism requires that terms used to designate properties or kinds are made true by the presence in the universe of general or universal entities intermingled with the particulars, Williams is no realist. If Williams is right, if universals are particulars, modes or tropes, considered without concern for their particularity, if this is what universals *are*, then Williams ought to be accounted a realist.

Whether previous immanent realists would recognize their view in this opinion that universals are immanent because they are, to speak crudely, the similarity roles (or 'adjectival identities') of abstract occurrents, I have some doubt. I am sure, from experience with myself, that an immanent realist begins by thinking he means more, but can bring himself to see, or think he sees, that he *couldn't* mean more—that every attempt to state an alternative results in something verbally but not significantly different from just redefining 'identity' by resemblance. (Williams 1959, 10)

Statements concerning universals can be true, but their truthmakers are wholly particular. Universals, as Armstrong would say, involve 'no addition of being'. Given Socrates' whiteness, a case or particular, you have a way Socrates is. This is a way other things might be in virtue of being exactly similar colourwise. Every case is a kind; every way something *is* is a way something else *could* be—by virtue of being an exactly similar way. Campbell (1990, 44–5) dubs this 'painless realism', explaining it by reference to a 'rule for counting' that can treat modes or tropes sometimes as cases, particulars, sometimes as kinds.

As an earlier comment on Locke suggests, I suspect that many philosophers in bygone days who spoke of universals had something of this sort in mind. Consider the remarks of an historically sensitive

proponent of universals, E. J. Lowe. Lowe regards universals as 'abstract entities'. However,

abstract entities are not denizens of some 'Platonic' realm which is 'separated' from the world of things existing in space and time. According to this view, to say that abstract entities do not exist 'in' space and time is not to say that they somehow exist 'elsewhere', a notion which is doubtfully coherent in any case. It is merely to say that when we speak of abstract entities we must 'abstract away' from all spatio-temporal determinations and distinctions. (2002, 66)

Lowe appeals to sets by way of analogy:

Although the planets are concrete objects, each one occupying some particular spatial location at every time during its existence, the set whose members are the planets cannot be assigned a spatial location and cannot be said to persist through, or undergo change in, time. A set of objects exists, timelessly and without spatial location, in any possible world just in case those objects exist in that world. Time and place simply do not enter into the existence- and identity-conditions of sets and that is why they qualify as 'abstract' objects. (2002, 66)

The same will be true of

properties, conceived as universals, such as the property of being red or the property of being square. Even if the properties are, like these ones, properties exemplified by concrete things, such as flowers and books, the properties themselves are abstract entities because time and place do not enter into their existence- and identity-conditions. According to 'Aristotelian' realism concerning universals, it is a necessary condition of a property P's existing in a world w that some object should exemplify P in w—and if that object is a concrete one, P will be exemplified by it at some time and in some place. But this does not imply that P itself exists at any time or in any place. By implication, then, I am rejecting here the doctrine, strangely popular just now, that universals exemplified by concrete objects are 'wholly present' in the space–time locations of those objects—a view which I have elsewhere argued to be incoherent. (Lowe 2002, 66; see also Lowe 1998, 155–6)

In discussing substance and property, I suggested that both are abstractions. You can consider an object—this tomato, for instance—as a substance, as a bearer of properties, or you can consider ways the object is, its sphericality, its redness. You can consider ways this tomato

is, as ways *it* is—*its* redness, *its* shape—and you can consider those ways purely as ways, as ways other things might be. In so doing, you are considering them as universals. And they *are*, so considered, universals. The colour of this tomato is a particular way a particular object is; but this way is, as well, a way that another object—another tomato, say, or a convertible, or a jumper, or a sunset—could be. Another object could be this way—be *the same* way—by exactly resembling the tomato colourwise.

A quick reminder: talk of abstraction and partial consideration does not make what is abstracted mind-dependent. In abstracting, you attend to what is there to be considered. But does this really yield mind-independent universals? Think of tables and trees. Does realism about such things require that the universe include, in addition to the particles duly arranged (or space–time locally thickened, or the quantum field in flux), tables and trees? Or might it be enough that thoughts of tables and trees can be, and often are, literally true? Such thoughts could be true, yet their truthmakers be particular arrangements of particles or regional thickenings of space–time. In the same manner, reflections on universals need not require that the universe contain, in addition to all the particular things, general or universal things. To think that is to perpetuate a philosophical myth, a myth that, strangely enough, seems more influential today than at almost any time in the past.

5.8   Modes and Tropes

I have described properties as ways substances are and indicated a preference for the traditional 'mode' over Williams's 'trope'. Most philosophers of my acquaintance who are willing to discuss such topics regard 'mode' and 'trope' as equivalent, two terms with a common denotation. True, modes are best understood as at home in a two-category, substance–attribute ontology; modes are particular ways particular substances are. Tropes, in contrast, bundle together to make up objects: tropes are modes minus the substance. This easy distinction, I now believe, harbours a mistake of a fundamental sort, a mistake that threatens incoherence.

You cannot subtract the tomato from a way the tomato is. It is of the nature of a mode to be a mode, a modification, of some substance. This feature of modes underlies medieval debates over whether modes are *entities*. Ways the tomato is are entities in the sense that the tomato's being the ways it is makes a qualitative and dispositional difference to the tomato. Is your smile an entity? Your smile is a real feature of your face, a way your face is, but calling it an entity is apt to mislead. Certainly, modes are not entities in the sense in which the tomato's substantial parts—its stem, its skin, its seeds—are entities, items that make up the tomato, constituents of the tomato. You might construct a tomato by assembling the right kinds of particle in the right way. But a tomato is not made up of an assemblage of modes bonded to a substance.[15]

If tropes are kinds of entity that could, in combination, make up objects, then tropes differ fundamentally from modes. Properties, considered as tropes, are, as Williams says, *parts* of objects, *mereological* parts, in the sense that they collectively make up or compose objects: objects are nothing more than bundles of tropes.

Although modes have been around since Plato, tropes are a twentieth-century invention. 'Trope theory' emerged hand in hand with another dubious innovation, 'bundle theory'.[16] Indeed, it is precisely because properties were widely regarded as modes or accidents, that bundle theories seem never to have been seriously considered, much less defended. It would make no sense to regard objects as bundles of accidents or modes. My reading of history suggests that the closest earlier philosophers came to a bundle theory was in discussions of the Eucharist. God transforms the wine into the blood of Christ by miraculously subtracting the substance and leaving behind the properties. The properties left, however, are not modes. Not even God could eliminate a substance and leave ways the substance is, its modes. Properties that can survive the miraculous subtraction of the substance to which they belong are a species of *individual accident*, so-called *real accidents*.

---

15 Another reminder: on the view defended here the tomato is not *really* a substance and its redness and sphericality are not *really* properties. I am treating the tomato as a substance for illustrative purposes only.

16 It is, I suggest, anachronistic to call Hume a bundle theorist. Berkeley's material objects were bundles of ideas, but ideas are modes of immaterial substances.

In a discussion of real accidents, Descartes notes that such entities would be propertied substances. 'The human mind cannot think of the accidents of the bread as real and yet existing apart from its substance, without conceiving of them by employing the notion of a substance' (Descartes 1641, 176). The whiteness of the bread considered as a stand-alone entity is something white, something that is the white way.

This is not the place to take up the troubled history of the metaphysics of transubstantiation: metaphysics in the service of a specific theological doctrine. The important point here is that tropes resemble real accidents in being very special kinds of entity, utterly different from traditional conceptions of particularized properties as modes. Even real accidents begin their careers belonging to a substance. They persist 'unsupported' only by a miraculous intervention. For these reasons, I have come to regard it as a bad idea to lump modes and tropes together. Substance–attribute trope theories (Martin 1980, 2008; Armstrong 1989, 136) are more perspicuously thought of as taking properties to be modes.[17]

The sweeping, guileless, historical claims advanced in the preceding paragraphs are sure to raise eyebrows. The history of disputes over modes and accidents is nuanced and open to sharp interpretative disagreement. Matters are made even more complicated by the deployment of cumbersome circumlocutions aimed at diffusing potential ecclesiastical distractions arising over the metaphysics of transubstantiation and the Trinity.

At the risk of digging myself in even more deeply, I will add that it is difficult not to see particular metaphysical disputes among medieval and early modern philosophers as often as not as resting on terminological infelicities. This thought first occurred to me in the course of reading Jeffrey Brower's perceptive discussion of 'Medieval Theories of Relations' (Brower 2009). I shall take up relations explicitly in chapter 7, where I suggest that, in some cases, philosophers defending apparently inconsistent accounts of relations might better be understood as struggling to express a single, fundamental thesis: irreducible relational truths might turn out to have non-relational truthmakers.

17 The situation is complicated by the fact that self-styled trope theorists are not always as clear as they might be on the nature of tropes. See Garcia (2015).

Something like this is, I believe, going on in discussions of modes, individual accidents, and universals as well.

I could be utterly wrong in these historical reflections. Indeed, it is a safe bet that I *am* wrong. This would not affect the philosophical points at issue here. These stand or fall on their own. I will be happy if I have at least given pause to those philosophers who imagine that talk of particularized properties is historically anomalous.

5.9 Santayana to Lewis via Williams

Many friends of universals will remain unpersuaded that an ontology that admits universals, but only as 'no addition of being' in a universe of particulars, could possibly be taken seriously. Under the circumstances, it might be helpful to conclude this chapter with a brief look at some of the deeper historical currents circulating around Williams's position and at their influence on ways of thinking that many of those inclined to smile at Williams's brand of realism continue to find perfectly unexceptionable. You can, as it happens, discern a clear line of influence from Santayana to Williams, and from Williams to David Lewis. I offer these observations as evidence for the significance of a strain of thought that has, in our time, remained largely subterranean. The story begins with Santayana.

You have seen already that, in borrowing Santayana's term 'trope' to designate individual accidents, modes, or abstract particulars, Williams turns Santayana's terminology on its head. Santayana uses 'trope' to designate 'essences', Santayana's label for universals, which he distinguishes from the 'flux' of 'existents'. The flux includes events, non-repeatable particulars such as 'the birth of Christ or the battle of Waterloo' (1930, 102).

Each event is a particular and can occur only once. Only the *type* of such a sequence, composed of such moments, is the *form* of the event, and this form is a universal. It need never have occurred; if I had said the resurrection of Christ, instead of his birth, the reader might have his doubts about it. The fact that such an essence was exemplified somewhere in an event would be a historical truth; in order to substantiate it, the flux of matter must assume that form; without this material and incidental illustration that type of sequence would remain in the air, in the realm of fiction or of theory.

It is especially important at this point to dispel that confusion between essences and facts which makes a quicksand of all philosophy. I will therefore give a separate name to the essence of any event, as distinguished from that event itself, and call it a *trope*. (1930, 102–3)

Earlier, I speculated as to why Williams would borrow a term from Santayana then invert its meaning. The idea that Williams might have been confused about Santayana's conception of tropes, or that he did not take Santayana seriously, or that his usage was meant to be ironic, does not survive scrutiny. Williams (1954) provides a compelling reading of Santayana on 'essence and existence' in which it is clear that he understands what Santayana is up to. One striking feature of Williams's account is the extent to which the doctrine he finds in Santayana closely resembles Williams's own ontological scheme. Indeed, it is clear that Santayana's conception of tropes—i.e., universals—*is* Williams's. When this is coupled with the fact that the ontological picture we today associate with David Lewis can be found in fledgling form in Williams, we have a surprising progression from Santayana to Lewis, via Williams.[18]

This observation might strike you as implausible, but I consider it worth pursuing, if only as an indication that Williams's way of accommodating universals was in the air, at least at Harvard in the 1950s and early 1960s. As a preliminary, consider Williams's ontological picture (as sketched in his 1953 and 1959). A preliminary word of warning. Do not be misled by Williams's reference to '*qualia*' in the quoted passage. It is no part of Williams's thesis that *qualia* are mental (see the passage from Lewis quoted below).

The world whole, I take at least as a working hypothesis, absolutely all there is, is a four-dimensional plenum of *qualia* in relations, eternally actual through and through. Its fundamental pattern, which all other structure presupposes, is that of whole and part: the Big It is not merely infinitely divisible, or virtually infinitely, but infinitely *divided* in the sense that it is the sum of countless actual parts, countlessly including, overlapping, and excluding one another, each part and each whole as genuinely real and individual, in the cardinal logical and ontological respects, as any whole which includes it, right up to the World All, and as any part which is

18 More perspicuously: from Locke and Hume to Santayana to Williams to Lewis.

included in it, right down to the ultimate indivisibles which have no proper parts, if such there be. Each of the parts thus intrinsically individuated, identical with itself and distinct from everything else, each *thing*, is related to each other part or thing, and to itself, in two further fundamental ways, by location, that is, the distances and directions which compose the four-dimensional spread, and by resemblance (with the proviso, in default of a better inclusive word, that 'resemblances' covers both likeness and unlikeness).

Now among the 'parts' of the world, and of lesser things, to which I have referred, are the ordinarily recognized things and parts, like automobiles, their wheels, the tires on the wheels, the atoms in the tires, and so forth. But there are others, equally actual and individual, which in ordinary life we seldom notice or name: the top half of the automobile, for example, or the sum of a thread on a bolt in it with the Queen of Sheba. Since the Queen is not a contemporary of the bolt, this entails that an individual whole may have parts widely dispersed not only in space but in time, which reminds us, in turn, that since the objects we are dividing,—i.e., whose dividedness we are observing—are four-dimensional solids, we may and often must distinguish temporal parts as well as spatial ones. We acknowledge, for example, the October 1959 span of, say, a horse's total being as well as the horse's neck for his life long. Now, all these parts and individuals, even the gerrymanders thus so oddly spread or sliced in space and time, are not only real but *concrete*. One thing, if not the main thing which this means is that, however discontinuous the place-time, or 'plime', which just contains such an object, the object exhausts or is the whole content of it. (1959, 2–3)

The ultimate parts are spatio-temporally distributed tropes, particular qualities. What of space–time itself? Is space–time a trope? One interesting possibility, a possibility not addressed by Williams, is that space–time subserves the role of substance. Tropes, then, would not be qualities distributed *in* space–time, but qualities *of* space–time, modes. Such a view would belong to the substance–attribute tradition. Williams's official view is apparently quite different. Objects are sums of tropes. Our everyday and scientific concepts honour a tiny minority of these: tables, trees, persons, but not 'the sum of a thread on a bolt in [an automobile] with the Queen of Sheba'. This is the deep story about reality, a story that nowadays might be classified as reductionist, or eliminativist, or perhaps conventionalist, labels Williams would certainly have rejected. Trees and the like are perfectly real. Our tree concept encompasses some mereological sums

and excludes others. The sums—all of them—are perfectly real. A paltry few satisfy our various concepts, most do not. The whole affair is deeply contingent.

On this conception, there are no connections between distinct existences, where the existences are the tropes: Santayana's 'radical contingency' (Santayana 1930, 110). What you do find, particularly as you descend toward a fundamental level, are repeatable patterns, exactly resembling concurrences of tropes peopling (or perhaps 'characterizing') distinct regions of space–time, distinct 'plimes'. These are the bases of our everyday universals. But, as Williams would insist, truthmakers for statements concerning universals will be particular through and through: individual tropes, trope sums, and the Big It, the sum of all the tropes.

The similarity between this conception of the universe and Lewis's is striking. Here is Lewis's oft-cited account of 'Humean supervenience'.

Humean supervenience is named in honor of the great denier of necessary connections. It is the doctrine that all there is to the world is a vast mosaic of local matters of particular fact, just one little thing after another. (But it is no part of the thesis that these local matters are mental.) We have geometry: a system of external relations of spatio-temporal distance between points. Maybe points of spacetime itself, maybe point-sized bits of matter or aether or fields, maybe both. And at these points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point to be instantiated. [In a footnote, Lewis mentions Williams's tropes as candidate qualities.] For short: we have an arrangement of qualities. And that is all. There is no difference without difference in the arrangement of qualities. All else supervenes on that. (1986a, ix–x)

All that is missing in Williams is the modal element Lewis provides by introducing alternative, 'possible' worlds.[19] What is much less obvious is the extent to which the core ideas can be found in Santayana, or at least in Williams's account of Santayana's ontology. That account is

---

19 Although no 'modal realist', Williams invokes possible worlds in a number of places; see, for instance, his 1953, 3, 7, 8; 1966, 74, 78, 80. In thinking about Lewis on modality, it is tempting to imagine that Lewis's worlds serve as truthmakers for modal claims, but this is not his view. Modal truths are made true by intrinsic features of our world. I shall not defend this contention here, but see Lewis 1986b, 22: 'it is the character of our world that makes the counterfactual true . . . the other worlds provide a frame of reference whereby we can characterize our world'. See also 192–8. Daniel Nolan and Denis Robinson helped me appreciate this point.

worth looking at as much for what it tells us about Williams—and perhaps Lewis—as for what it tells us about Santayana.

5.10 Williams's Santayana

Nowadays, philosophers who read David Lewis are unlikely to read George Santayana. One barrier is Santayana's literary style, which suggests a kind of overwrought, nineteenth-century European approach to philosophy, temperamentally at odds with the kind of unflinching dispassionate realism nowadays associated with honest ontology. According to Williams, however, Santayana stands out among his generation of American philosophers as profoundly realist and particularist, a successor to Locke, not Hegel.

George Santayana was master of a solemn showmanship which is to blame for his being taken for a more trivial and alien philosopher than he was, both by those who like the trivial and the alien and by those who do not. . . . The longer one peers, however, through the opalescent glass with which his style conceals the man, the more one discovers of a more solid philosophical citizen. (1954, 31–2)

Once you move beyond what would strike the contemporary philosophical temperament as literary affectations, you are positioned to appreciate Santayana's contribution to serious metaphysics. Williams puts it this way:

However long men may pay their respects to the Santayana who is the elegant and perceptive essayist, it is in his ontology that he proves himself a philosopher of the grand mold, a philosopher, if I may borrow a flourish from Mr. Wyndham Lewis, 'absolute, displayed, and regardant, in the Chief, the Pale, and the Quarter Fess'. (1954, 32–3)

Williams finds in Santayana a philosopher bent on wresting American philosophy from the clutches of the Hegelians and returning it to its roots in traditional, scientifically informed, level-headed empiricist thought.

When Santayana went back [to] the teachings of his American contemporaries it was not, as he sometimes pretended, to recover a secret of the ancients or some subtlety of the Spanish schools, but to consummate better

than the English and the Americans a principal theme in our own heritage. For in an important respect the nineteenth-century spokesmen catalogued as typically American, and so irksome in that capacity to Santayana, represented a freakish and foreign episode. James, Dewey, Royce, Thoreau, Whitman, even Emerson, native enough by birth and animus, drenched themselves to distraction in the imported liquors of German idealism, of French spiritualism, and even of Indian mysticism. Their transcendentalism and 'egotism', as Santayana called it, mistaking their thoughts for the masters and makers of reality, taking, as he said, the turnips and cabbages of their own kitchen gardens for the signs of the zodiac,—these were academic borrowings from Europe and Asia. For two and a half centuries before them, our countrymen had been taught in a profoundly different philosophy of realism, that of dissenting Christianity and Cartesian science, the common sense of Locke, of the Scottish philosophy, and of the plain citizen who had to contend with Indians and the weather. What Santayana did, in effect, was to restore the older and deeper strain, disciplined and criticized with all the mordant ingenuity which he inherited from the later British empiricists, and sweetened, perhaps, with a certain specious Iberian nostalgia. As Locke's philosophy was to its Scholastic prototype, reformed by a more scientific logic and plain speaking, so Santayana's is to Locke's. (Williams 1954, 34)

This strain in Santayana's thought emerges clearly in his conception of universals. Santayana, as noted above, distinguishes 'essences' from 'existents', but Santayana's essences are not the essences beloved of the Aristotelians. Existences are particular entities; essences are Williams's universals, 'generized' ways existents are.

To draw the line between essence and existence is no more than to notice the difference between *what* a thing is and *that* it is, between its character or kind and its occurrence as a case of the kind, between the rosiness of the rose and the fact that it is *there*. Language and logic as well as metaphysics have always found the contrast well nigh indispensable, from Socrates to Carnap, but for the scientific purposes of modern philosophy it has been spoiled and discredited by the metaphysical marvels which have been associated with it. (Williams 1954, 36)

A thing's 'essential properties' are not a privileged subset of its properties, distinguishable from its inessential, merely accidental properties, they are just the properties it has. Indeed *it*—any it at all—might be

no more than a sum of properties, a sum of tropes (in Williams's sense), each as 'essential' to its identity as any of the rest.[20]

Plato, for example, taught that the essences are eternal prototypes, an ethereal aristocracy of ideals, laid up in heaven as patterns to be distantly imitated by the things that exist here below. Aristotle taught that the essence of a thing is its inner principle, its ontological soul which preserves and improves it, explains it, and makes it what it is, in contrast with the rest of its properties, which are merely accidental. Both of those philosophers, and most of their more official heirs for two thousand years, agreed that true essences are relatively few, that they are accessible only to reason, that they are good, that they are what things ought to be, that they act on things by luring or egging them forward, and that in so far as things fall short they not merely are wrongly constituted, they are *un*constituted, intrinsically deficient and vague. (Williams 1954, 36)

'This', Williams concludes, 'is the metaphysics of the dog show: that there is a sublime essential Poodlehood, for example, to which every poodle aspires and deserves a blue ribbon if he achieves it better than other poodles' (Williams 1954, 36).

In contrast, Santayana's essences are ways things are. These ways are universals in precisely Williams' sense: they are lodged in a universe of particulars, they represent 'no addition of being'. If anything is or could be any way at all, you have the possibility of some other thing being 'the same' way, where sameness is the sameness of similarity, not identity.

Universals—Santayana's tropes—are powerless. A law of nature is a trope, which Santayana describes as a 'skeleton' that

exists, in the sense that by hypothesis this law prevails now in this region; but it is no skeleton in its operative function, like that of an animal; it is no rigid substance within the soft substance of events; it is only a trope which the thrifty mind selects from the tangle of relations which hold those facts in a mesh of existence. (Santayana 1930, 112)

Essences generally, and laws in particular, do not direct traffic, do not oversee, or govern, or steer the course of nature, but merely express recurring patterns among similar existences. The traditional doctrine of essences,

---

20 I hesitate here because it is not obvious to me that Santayana could be read as a bundle theorist.

though congenial to the human heart and endowed with that *prima facie* plausibility which is the one chief organ of Peripatetic thought, this 'moralistic physics', as Santayana called it . . . was worse than useless to the scientist and a millstone round the neck of the ethics it was supposed to salvage. (Williams 1954, 36)

We have come full circle. Santayana's essences, Santayana's tropes, Williams's universals, are grounded in the utter particularity of existents, Williams's tropes. In neither case are universals meant to belong to an other-worldly Platonic realm. Nor do they assert themselves in the affairs of existents. Santayana locates them in a *this-*worldly 'Platonic realm' (1930, 113). Given a particular, you have a way that particular is or could be, a way that would be 'manifested' by any similar particular. What more could you want?

CHAPTER 6

Causing

The Great majority of efficient natural causes are simultaneous with their effects, and the sequence in time of the latter is due only to the fact that the cause cannot achieve its complete effect in one moment. But in the moment in which the effect first comes to be, it is invariably simultaneous with the causality of its cause . . . If I view as a cause a ball which impresses a hollow as it lies on a stuffed cushion, the cause is simultaneous with the effect. But I still distinguish the two through the time-relation of their dynamical connection. For if I lay the ball on the cushion, a hollow follows upon the previous flat smooth shape; but if (for any reason) there previously exists a hollow in the cushion, a leaden ball does not follow upon it.[1]

Nature comes in package deals. (Martin 2008, 3)

6.1 The Received View

Philosophers today commonly regard causation as a relation among particular events. According to what I shall call the received view, one event, the cause, brings about a distinct event, the effect. The relation is asymmetrical (effects follow causes), nonreflexive (no event can cause itself), and transitive (if A causes B, and B causes C, A causes C). Humeans think that a particular A causes a particular B just in case there is a true universal generalization of the form 'If an A-type event occurs, a B-type event occurs.' Others hold that A's causing B is a matter of there being a law of nature governing A- and B-type events, a law to the effect that A-type events necessitate B-type events. Such

---

1 Kant (1787, A203, B248–9). Stephen Mumford called my attention to this passage.

laws entail, but are not entailed by, true generalizations. Laws of nature are contingent in the sense that there could be universes including *A*- and *B*-type events, but, owing to differences in the laws, the *A*'s do not cause the *B*'s. Still other philosophers invest objects involved in causal sequences with powers: to be an *A* is, at least in part, to possess the power to cause a *B*; to be a *B* is, at least in part, to possess the power to be caused by an *A*. On such a conception, laws of nature 'supervene' on the powers, at least in the sense that truthmakers for statements about laws are powers.

Each of these ways of conceiving causation could be thought to reflect an aspect of the phenomenon, but each is, in its own way, ill-conceived. I believe our understanding of causation would benefit from a shift of attention from causal sequences and laws, to instances of causation: *causings*.

6.2   The Causal Nexus

Consider a simple example. You stir a spoonful of salt into a glass of tap water, and the salt dissolves. Focus on the salt's dissolving. The received view of causation might lead you to think that the water and the salt are related as agent and patient: the water, or maybe the water's enveloping the salt, causes the salt to dissolve. Perhaps the water possesses an 'active power' to dissolve salt, and salt, a complementary 'passive power' to be dissolved by water.[2] But look more closely at what happens when you stir salt into a glass of water. Certain chemical features of the salt *interact with* certain chemical features of the water (Ingthorsson 2002). This interaction is, or appears to be, continuous, not sequential; it is, or appears to be, symmetrical. Both the salt and the water work in concert to yield a certain result: the salt's being dissolved in the water. Think of this interacting as a causing.

There is an event sequence here: your stirring the salt into the water, and the water's becoming saline. But it would be odd, or at

---

least overly instrumentalistic, to regard *this* sequence as the causal nexus. Rather, you have a continuous process whereby the water and salt work together to yield an outcome, salt water.

One way to understand such cases is to imagine that salt and water possess *reciprocal* powers or dispositions.[3] The salt's dissolving is a *mutual manifestation* of these dispositions. The result is something with new powers, new dispositions capable of further mutual manifestations with further reciprocal disposition partners. We philosophers are trained to think of causation sequentially: a cause occurs, followed by an effect. In the case under consideration, the effect—the salt's being dissolved in the water—is the outcome of a causal process, a causing, that is itself symmetrical. Might this process be identified with the cause? Doing so would move us closer to the truth, but I do not think we are quite there yet.

Consider another example. You take two playing cards and prop them up against one another so they stand upright on the table. The cards—with the help of the table—are mutually supporting: they remain upright. The cards work together with the table (and the gravitational field) to produce this result. But their working together and the result are not sequential. The effect obtains so long, and only so long, as the causing does. The cards' remaining upright is a continuous mutual manifestation of reciprocal powers possessed by the cards and the table.

Examples of this kind do not fit comfortably with the received model, yet, arguably, they are by far the most common species of causal interaction. We depend for our existence on stable structures that we inhabit, move about in and on, and deploy. We count on our environment's maintaining a high level of stability. Stability requires massive cooperation, the mutual manifestation of countless reciprocal powers to hold things together, to preserve the status quo.[4] Their holding together is an outcome, but one that temporally coincides with their manifesting themselves as they do. The causing here is a reciprocal, symmetrical, continuous affair.

---

2 Locke has 'active' and 'passive' powers, Sydney Shoemaker (1980, 1998, 2007) speaks of 'forward-' and 'backward-looking' 'conditional powers'. Davidson (1973, 64) advances a position consistent with the one defended here and in § 4.10 above.

3 See chapter 4, above. For background, see Martin 1993, 2008, and Martin's contribution to Crane 1996; Heil 2003a. Here, as elsewhere, I use 'power' and 'disposition' interchangeably.

4 Another reason to doubt the thesis that an object's continued existence over time is a matter of temporal parts of the object causing their successors. Most objects require the presence of endless favourable conditions to persist from moment to moment. See § 3.2 above.

What of familiar causal sequences commonly deployed to motivate the event–causation model? One billiard ball approaches another, stationary, billiard ball. The balls collide. The second billiard ball moves off in a particular way. This sequence can be described as one in which one event, the first billiard ball's striking the second, causes a distinct event, the second billiard ball's rolling across the table. Such a description is unperspicuous. The trajectory of *both* balls is altered. Think now about the causing in such cases. When the first billiard ball makes contact with the second, *both* balls compress, then decompress. This process is, or appears to be, continuous, symmetrical, reciprocal (Huemer and Kovitz 2003). Its outcome is a change in the velocity of each ball. This outcome resembles the outcome in the salt water case. It results from a causing. In the playing card case, the outcome is simultaneous with the causing.

So we have causings and outcomes. Let outcomes be effects. Then effects can, but need not, follow causings. We are making progress. What of *causes*? Well, what *are* the causes in each of the cases we have considered? In the salt water case, perhaps the cause is your stirring a spoonful of salt into the water. In the playing card case, it might be your placing the cards together in an upright position of mutual support. In the billiard ball case, the cause might be the first billiard ball's coming into contact with the second. In each of these situations, what you might regard as a cause is something that makes it the case that powers of the objects involved come to be manifested in a particular way. Often this involves bringing them together. Considered ontologically, causes and effects take a back seat to causings. Causing is where the action is.

6.3 Powers, Dispositions

Several points here are worth noting. First, causings are typically mutual manifestings of *many* reciprocal powers. We are inclined to omit mention of most of these in our descriptions, relegating them to the status of 'background' conditions. This practice is unobjectionable in cases in which our interest lies in singling out particular components of causings in assessing responsibility for various outcomes (the spark, not the surrounding oxygen, caused the explosion),

or in determining what you would need to add to achieve a particular outcome (the spark, not the oxygen). But this way of talking cuts no ontological ice. Any outcome, even the status quo, turns on the cooperation of multitudinous reciprocal powers.

Second, the status quo is not a matter of powers waiting docilely to be manifested. The status quo is itself a mutual manifesting of countless powers, indeed many of the same powers that manifest themselves in new ways with the advent of new reciprocal partners.

Third, to the extent that it might be correct to think of causings as mutual manifestings of reciprocal powers, it is important to see that one and the same power is capable of manifesting itself differently with different kinds of reciprocal partner. Failure to appreciate this point, has led to confusion in recent discussions of powers.

Consider an electron's power to repel fellow electrons. Suppose this power resides in the electron's charge. The very same power would lead the electron to attract positrons, to trigger a Geiger counter, and to behave as it does in the company of neutrons and protons. One unhappy consequence of denying that powers could manifest themselves in different ways with different reciprocal partners is that this would render most powers undetectable. If an electron's negative charge is *solely* a power to repel other electrons and nothing more, we could have no evidence that it possessed this power, at least not if our having evidence requires our being causally connected in the right way with whatever our evidence is evidence for.

A ball's sphericality endows it with a power to roll. But it is also in virtue of being spherical that the ball has the power to make a concave, circular impression in a cushion, the power to reflect light so as to look spherical, the power to feel spherical to the touch. Talk of single- and multi-track dispositions or powers is confused from the outset. Powers, quite generally, are 'multi-track', if this means that they would manifest themselves differently with different kinds of reciprocal partner. If you start with the thought that the diversity we find in the universe stems from varying combinations of a small number of different kinds of fundamental entity, then you will want powers to be capable of diverse manifestations with diverse reciprocal partners.

Philosophers who have been attracted to an ontology of dispositions or powers, sometimes characterize powers as features of objects that manifest themselves in a particular way given a particular kind of 'trigger' or 'stimulus'. Alexander Bird (2007), for instance, takes a disposition, D, to be characterizable by reference to a manifestation, M, resulting from D's being stimulated by S.

Note, first, how poorly this way of thinking about dispositions fits with our examples, with, for instance, salt's dissolving in water. You have salt, water, and the salt's dissolving. Where do you locate D? In the salt? In the water? And where is S? Is S the salt, the water, or something else?

Second, note that dispositionality is being characterized by reference to causation. But it is easy to think that this has matters backwards: you can explicate causation in terms of the mutual manifesting of dispositions or powers.

Third, note that by individuating dispositions Bird's way, you are bound to over-count. Suppose one and the same power manifests itself one way with one kind of reciprocal partner and another way with a different kind of partner. This, I think, is the norm, although it altogether eludes the individuative scheme implicit in the idea that a disposition can be pinned down in the manner suggested by Bird.

So where are we? We have causes, causings, and effects. Both causes and effects are *themselves* causings. This, I think, is a more perspicuous way of rendering the idea, implicit in the received view, that causes are themselves effects of prior causes, effects themselves causes of further effects. Both conceptions settle on a single ontological category: events, on the received view, causings, on the alternative view under consideration. Maybe causings are really events or events really causings. Well and good. In that case, the view under consideration could be taken as a suggestion for improvement on the received view.

6.4 Liabilities of the Received View

Think of some of the traditional liabilities associated with the received view. Suppose you thought that causal sequences were relations among distinct events, the cause preceding the effect.

Must the causing event be completed before the effect occurs? Suppose event A causes event B. Must A expire before the onset of B? That seems unpromising. How could an event that has run its course *do* anything? Perhaps A and B temporally overlap. Consider that portion of A that occurs prior to the onset of B. This portion of A has run its course, so how could *it* play a role in causing B? It appears that only the portion of A that overlaps with B, A_2, actually causes B. So, really, it's not that A causes B, but that some portion of A, A_2, causes B. But unless A_2 and B are temporally coextensive, there will be a portion of B, B_2, that occurs after A_2 has run its course. How could A_2 be responsible for B_2?

These difficulties will not strike you as difficulties if you think of causation abstractly: the A's cause the B's because there is a *law* to the effect that the A's cause the B's; *what's the problem*? But when you move from pencil and paper, where stipulation rules, to serious ontology, the worries creep in. Note that these worries do not arise if you think of causation as I have recommended, if you think of causation in terms of causings.

On the conception of causation under consideration, causation is understood in terms of causings, and causings are mutual manifestings of powers. This might arouse suspicion. Is the recommended conception of causation dependent on a hokey ontological picture of a universe populated by mysterious powers?

A conception of causation as causing and an ontology of powers do, as befitting sensible ontological theses generally, enjoy mutual support. Nevertheless, what I have said about causing could doubtless be rephrased so as to fit less adventurous ontologies.

I have recommended that we replace the received conception of causation as a relation among successive events with a conception of causation as causings. One consequence of replacing the received picture is a lessening of the appeal of thinking of causation in terms of causal *chains*. If our interest is in the ontology of causation and not simply in prediction and control, we are better off thinking of causal *networks*, what Martin (1993, 2008) calls *power nets*. Power nets are evolving, massively cooperative ventures among the powers, constellations of causings, mutual manifestings among intertwined reciprocal partners. Talk of causal chains, no less than talk of 'background

conditions', is perspectival. Such talk is explanatorily or instrumentally useful, perhaps, but ontologically ill advised.

Is this all hopelessly metaphorical? I suppose it is metaphorical. But if it is, it is no more metaphorical than more familiar talk of causal chains and sequences. Our thoughts about the universe are guided by metaphors. A metaphor apt in one context can mislead in another. The trick is to understand a metaphor for what it is and limit its application accordingly. What I am after here is a conception of the ontology of causation, one that does not take on at the outset substantive, albeit implicit, ontological baggage. The idea that causation is a relation among successive events did not fall from the sky. It is the accepted view, the received view, only because we philosophers have colluded in according it this status.

6.5 Causing and Indeterminacy

In scientific circles, causation was long regarded as wholly deterministic. You can see why such a view might have ontological appeal. Suppose the A's cause the B's—but only *sometimes*. What distinguishes situations in which an A causes a B from those in which it does not? In such cases it would be natural to suppose that, when an A causes a B, this is not the whole story. A must have had help. Really, A, together with C, causes B. That is why you can have A's without B's. C operates as a 'hidden variable'.

But suppose you have reason to think that there is no C: A's sometimes cause B's and sometimes not. Then the mechanism of causation looks mysterious. What is going on when an A *does* cause a B? Saying that the A's necessitate the B's—but only sometimes— appears paradoxical. If it is only *sometimes*, in what sense do the A's *ever* necessitate the B's? Talk of A's 'raising the probability' of B's, sensible, perhaps, when our interest is confined to prediction or explanation, appears ontologically dodgy. Suppose the A's *do* increase the probability of the B's; what happens on an occasion when a particular A 'produces' a particular B? The notion of production in play is entirely mysterious, occult.

All this will sound excessively naïve and aprioristic to philosophers attuned to contemporary physics. The success of the quantum theory

has led physicists to the thought that the universe includes non-deterministic processes at the fundamental level. There is no saying precisely when a radium atom will decay. When a radium atom decays, its doing so can be entirely spontaneous, entirely independent of goings-on outside the atom. What are the implications of such cases for our conception of causation?

First, it appears that a spontaneous happening can have momentous repercussions. Unmoved movers, once regarded as impossible, turn out to be ubiquitous. You might be tempted to describe cases in which this happens as cases in which a causal chain is initiated. I have suggested that we abandon the chain metaphor, and take up a picture of causation as a network of reciprocal causings. Quantum indeterminacy means that some of these causings will include spontaneous elements.

Second, although causal networks need not be deterministic, this does not mean that *causings* are indeterministic. A radium atom will decay only with a certain probability. This introduces a measure of indeterminacy into a system that includes the atom. What is not introduced is a special kind of causing: probabilistic or indeterministic causing. The probabilistic locus lies in the atom, not in its causing whatever it does when it does. Whether various reciprocal powers are on hand at a given time can be probabilistic, but, given the powers at t, causings at t are thoroughly deterministic.

The point can be illustrated by means of an analogy borrowed from Martin. An election can have a disjunction of outcomes, but not a disjunctive outcome. Unless one candidate *is* elected, there is no election.

If there is not such a selection or production of one disjunct rather than another between the candidates, there is no election nor a production of a disjunct at all. The explanation of the success of a candidate or production of a disjunct, short of magic, would be incomplete, and the links in this disjunctive linkage would not take us to the result, but only to a pre-result determination of disjunctivities. (Martin 1997, 219; 2008, 72)

The label 'probabilistic causation' is thus misleading. What you have is not a weakened or tentative kind of causing, but a probability of there being a causing. The probability might stem from oscillating features of objects such as radium atoms. The oscillations are uncaused

and probabilistic. I take this to be central to physical theory as it now stands. Indeterminacy is introduced into the causal network, not in the form of probabilistic causings, but in the form of spontaneity. Spontaneity coupled with deterministic causings yield an indeterministic universe.

6.6   Absences, Preventers, Antidotes, Blockers, Inhibitors, Finks

British seamen contracted scurvy owing to a lack of ascorbic acid, vitamin C, which is required for the synthesis of collagen, a protein essential to the support of living tissue. The disease is brought on by a lack, an absence. But how does this square with the picture of causing advocated here? Descartes denied that absences could be causes: nothing comes from nothing. David Lewis accepts absences as causes (Lewis 2000, 2004; see also Schaffer 2004). Who is right?

We certainly speak as though absences are causes: for lack of a nail a shoe was lost, for lack of a shoe a horse was lost, for lack of a horse, a battle was lost. . . . If our interest is in giving an analysis of the causal concept, we shall want to accommodate such locutions. If, in contrast, our interest is in giving an account of the nature of causation, its ontology, we shall want to say more, we shall want to say something about the truthmakers for judgements that invoke absences as causes. And this requires our saying something about the ontology of absences.

Do absences exist? Well, the universe would be very different if it differed with respect to its absences. The universe appears to be an arrangement of absences and non-absences. Borders, edges, and boundaries are interfaces between absences and non-absences. Absences are real, but it would be a mistake to think of absences as shadowy entities—properties or states of the universe existing alongside the non-absences. Think of a pothole in the highway. The pothole results from the road's surface taking on a particular shape. Suppose an automobile tire encounters the pothole and a blowout ensues. It seems natural to say that the pothole caused the blowout. But if this is true, what is the truthmaker? The blowout is a mutual manifestation of features of an automobile tire and features of the

road—in this case the way a particular portion of the road's surface is configured. A tire's smoothly rolling over the road is a mutual manifestation of features of the tire and the road's smoothness. Minus the smoothness you have a different sort of manifestation.

What of scurvy and the lack of vitamin C? A living body's healthy condition is a mutual manifestation of myriad finely tuned reciprocal disposition partners. When one of these is missing, you can have a different sort of manifestation, just as you have a different sort of manifestation when you remove one of the cards from a pair of propped-up playing cards. Here, as elsewhere, what you have is not an absence's stepping in and producing a particular kind of effect, but a different collection of reciprocal powers yielding a different kind of manifestation. An absence is not an entity, not a something with properties providing it with distinctive powers. But certain kinds of manifestation require appropriately propertied somethings as reciprocal partners. When these are missing, the result is a different kind of manifestation.

Here is David Lewis on the deadly void, a potentially lethal parcel of pure, unadulterated nothingness.

The void is deadly. If you were cast into a void, it would cause you to die in just a few minutes. It would suck the air from your lungs. It would boil your blood. It would drain the warmth from your body. And it would inflate enclosures in your body until they burst. (2004, 277)

Lewis's idea, apparently, is that the void includes an impressive array of causal powers. But *does* it?[5]

What I've said is literally true, yet it may be misleading. When the void sucks away the air, it does not exert an attractive force on the air. It is not like a magnet sucking up iron filings. Rather, the air molecules collide and exert repulsive forces on one another; these forces constitute a pressure that, if unresisted, causes the air to expand and disperse; the void exerts no force to resist the pressure; and that is why the air departs from the lungs.

Likewise, when the void boils the blood, there is no flow of energy from the void into the blood. It isn't like a stove boiling a kettle of water. The blood is already warm enough to boil, if its vapor pressure is unresisted; the void exerts no counterpressure; and so the boiling goes unprevented.

---

5 The quotations from Lewis that follow echo Martin 1996.

Likewise, when the void drains your warmth, what happens is that your thermal energy, left to itself, tends to dissipate; and the void provides no influx of energy to replace the departing heat.

And when the void inflates enclosures, again what happens is that the enclosed fluids exert pressure and the void exerts no counterpressure. So nothing prevents the outward pressure from doing damage.

In short, you are kept alive by forces and flows of energy that come from the objects that surround you. If, instead of objects, you were surrounded by a void, these life-sustaining forces and flows would cease. Without them you would soon die. This is how the void causes death. It is deadly not because it exerts forces and supplies energy, but because it doesn't. (2004, 177)

Lewis's account of the truthmakers for the assertion that the void would cause you to die fits nicely with what I have said about the manifesting of dispositions. If you think of the manifestation of dispositions as the mutual manifestation of reciprocal disposition partners, allowing that the selfsame disposition could manifest itself differently with different reciprocal partners, you can allow, as well, that a disposition would manifest itself differently in the absence of a given partner. In the absence of another card, a playing card would topple over. It is not however, that the absence exerts a force on the toppling card. Rather a card's remaining upright is a mutual manifestation requiring the right compliment of reciprocal partners.

Lewis's description of the effects of the void can be read as a restatement of this point. But you might worry that the description introduces a further difficulty. Lewis says that it is the absence of a 'preventer' that causes enclosures in your body to burst. Your belt prevents your trousers from falling down. Vitamin C prevents scurvy. But what is a preventer?

Talk of preventers, antidotes, blockers, inhibitors, and the like is perspectival in the way talk of *the* cause of a fire is perspectival. Consider a clock that keeps time incorrectly because the lubricant coating one of the gears has evaporated. You can say that the clock *malfunctions* because of a lack of lubricant. But what is going on is in fact just what *ought* to go on in a mechanism with the characteristics this one has. As Descartes remarks (*Meditation* VI, see § 9.9 below), describing the clock as malfunctioning is merely to note that it is not doing what we want, or intend, or manufacture it to do. In fact the clock is doing precisely what a device of this sort *ought* to do given its

current dispositional makeup. Its state is a manifestation of this dispositional makeup. Its makeup lacks a component, the lubricant, the addition of which would yield a different sort of manifestation.

Descartes makes the same point in discussing a healthy body. If vitamin C is not on the scene, bodily states manifest themselves differently than they would in concert with vitamin C. You can say that vitamin C prevents scurvy, but you could just as well regard vitamin C as a facilitator of a healthy body. This point can be generalized to antidotes, blockers, and inhibitors. Your swallowing a poison has a particular sort of effect on your system. This effect is a mutual manifestation of dispositions of your gut and dispositions present in the poison. If you subsequently swallow an antidote, the antidote blocks or inhibits this manifestation. But this is simply a matter of the dispositional system that includes dispositions of your gut, dispositions of the poison, and dispositions of the antidote yielding a different sort of mutual manifestation. Blocking and in-hibiting, like preventing, are a matter of dispositions manifesting themselves with various reciprocal partners as their natures dictate.

What of finks? Finks were unleashed on the philosophical com-munity by Martin in a paper aimed at undermining conditional analyses of dispositions (Martin 1994). Suppose you thought that 'the wire is live' expressed a conditional judgement: if the wire were touched by a conductor, then electrical current would flow from the wire to the conductor. But now imagine a device, an 'electro-fink', constructed so as to shut off current to the wire in the event that the wire comes into contact with a conductor. The wire is live, but the conditional is false.

Much has been written on finks and finkishness. The important point here, however, is that nothing is a fink *by nature*.[6] An ordinary circuit-breaker could be seen as a kind of electro-fink. When you look at finkish systems dispassionately, when you consider them as physical systems rather than as counter-examples or potential counter-examples to philosophical theses, they manifest themselves exactly as they ought to manifest themselves given their dispositional composition. Some of these ways of manifesting themselves frustrate our purposes, but they invariably do just what they ought to do. They

---

6 You might put this by saying that finks are not 'natural kinds'.

have no choice. Describing a system as finkish, or describing something as an antidote, an inhibitor, a preventer, a blocker—or, for that matter, as a facilitator—is to deploy what Descartes calls 'extrinsic denominators'. Nothing is, of itself, any of these things.

Where does this leave us? In the first place, although absences are not entities with properties or powers, when a given property is required for a particular kind of manifestation, its absence can result in a different kind of manifestation. It is not that the absence produces, or is causally involved in the production of, the manifestation. Rather properties on the scene yield a different, possibly unwelcome, kind of manifestation. Second, it is a mistake of a fundamental sort to imagine that, when a disposition fails to manifest itself in a particular way owing to the absence of a particular sort of reciprocal disposition partner, it is not manifesting itself at all. This mistake is abetted by talk of preventers, antidotes, and the like. Such talk has its place alongside talk of *the* cause of an effect in everyday discourse with its assorted aims and provocations. But this kind of talk has no place in serious ontology.[7]

6.7 Is Causality Fundamental?

Before moving on, let me step back for a moment and consider whether causal relations are utterly fundamental or whether truthmakers for causal claims might turn out to be non-causal ways the universe is. This might seem an odd endeavour. How could causation *fail* to be fundamental? Isn't causation a pillar of fundamental physics?

In answering this question, everything turns on what is meant by 'causation'. Thus, Bertrand Russell, in an oft-cited passage, remarks that

physics has ceased to look for causes.... in fact, there are no such things. The law of causality... like much that passes muster among philosophers,

is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm.[8]

You hear this sentiment echoed today, most often by philosophers intent on shocking tradition-bound colleagues by declaring that causation is a myth that has no place in serious science. What of Russell? Is Russell claiming that physics, or science generally, has no place for causation? I doubt it. Russell's target, like mine in this chapter, is a particular conception of causation, a conception built around a 'law of causality' according to which causation is an asymmetrical relation among distinct events, an instance of an exceptionless regularity. If this A causes this B, then whenever an A-type event occurs, a B-type event follows. Russell argues that, owing to inevitable complexities and wide-ranging interactions among objects studied by scientists, the chances of discovering significant regularities of this kind are vanishingly small. So, Russell rejects what many philosophers—Armstrong, for instance—regard as an important consequence of an adequate account of causation, the idea that causation is backed by laws that entail true, exceptionless universal generalizations. This is not a rejection of causation, however, but at most a rejection of a prominent empiricist analysis of causation, a rejection I am happy to second.

A deeper point lurks in the background, however. In discussing causation, I have been tacitly invoking a broadly corpuscularist picture of the universe. Distinct substances enter into causal relations owing to their properties. Suppose that the substances are not corpuscular, however. Suppose the substances are the fields, or the fields together with space–time. Ordinary objects, as well as the particles, would be modes, ways the fields are. Causal interaction would be preserved in so far as interaction occurs among the fields. Causation would survive in the form of interactions among the fundamental things.

Suppose, however, that the fields are unified: there is but a single, pervasive field. In that case truthmakers for claims about objects, including the particles, would be ways the field is. What of truthmakers for claims concerning causal interactions among the particles?

---

7 This, in fact, is close to the line advanced in Lewis 2004. In offering a counterfactual analysis of causation, Lewis takes himself to be engaging in conceptual analysis of a sort made popular in the twentieth century, a linguistic endeavour, not in providing an account of the nature of causation, an ontological enterprise. Moving from the former to the latter is to move from a consideration of truths expressible in a particular idiom to a consideration of truthmakers for these truths.

8 Russell, 1912, 1. See Chakravartty 2005, and 2007, chap. 4, for discussion.

If the universe were a single, unified field, or a single field pervading space–time, or, for that matter, a Spinozistic unified One, truth-makers for causal claims would be non-causal ways the universe is. The field, or the One, would evolve in ways that would be describable in terms of particle interactions, collisions among billiard balls, salt's dissolving in water, and all the rest. But the deep story would be non-causal.

If the field or the One played the substance role, it would be worse than misleading to imagine that *it*, the field or the One, caused its states, ways it is. The relation of a substance to its modes is not like the relation of internal states of your body to your breaking out into a rash. Your body is a complex thing made up of many complex things in constant interaction with one another and with the surrounding environment. But the field or the One is a unified simple with no parts to interact, and no other substances with which to interact. The evolution of such a substance over time would amount to an expression of its nature. In the absence of any other substance, this evolution would not be an effect of a cause. Ordinary perceived change, ordinary causal interaction would resemble ripples arising in a pond and moving across its surface, but uncaused by the wind or changes in the pond's constituents or surroundings.

Just as it would be a mistake to think of the properties of an electron as being caused by the electron, so it would be a mistake to think of the evolving properties of the unified field or the One as being caused by the field or the One. An electron's properties are ways it is, modes, expressions of its nature. If these properties change, the source of the change is either spontaneous or the result of an interaction with something outside the electron. In the case of the unified field or the One, there is no outside, there is only the spontaneous expression of the nature of the one substance.

I admit that these remarks are excessively speculative, but my aim here is only to leave open the possibility that we could be led by fundamental physics to allow that causal truths could have non-causal truthmakers. How likely this is to happen is anyone's guess. Were we led in this direction, we would have vindicated Spinoza, not, perhaps, as he would have wished, by means of an a priori, demonstrative argument, but empirically, defeasibly. Let it be noted, however, that we would have done so within the purview of a substance–mode ontology.

6.8   The Horse Before the Cart

What I have called the received conception of causation—an asymmetric relation among distinct events, the cause preceding the effect—has application, perhaps, in cases in which our interest is in explanation, manipulation, or the assessment of responsibility, natural or moral. In this regard it reflects the kinds of perspectivalism associated with talk of 'the' cause of a particular event. The mistake is to allow these epistemological concerns to call the shots ontologically. In so doing, you saddle ontology with the philosophical baggage of verificationism and instrumentalism.

My suggestion is that a dispassionate look at the ontology of causation takes us to causings, the mutual manifestings of reciprocal powers or dispositions. Causings are symmetrical, deterministic, and continuous. Causing is a cooperative endeavour in which outcomes depend on the mutual manifestings of reciprocal powers or dispositions. Such a picture leaves room for indeterminacy, but not for indeterministic causing. Indeterminacies are grounded in spontaneity, not in an attenuated kind of causing.

I propose this picture in an effort to get clear on the ontology of causation. Recent philosophy has seen a disturbing tendency to massage the notion of causation in the service of one or another provincial interest. Think of current debates over mental causation. Philosophers start with a conception of the mental and the physical and seek a conception of causation consistent with this picture that allows for mental–physical causal interaction. One popular view is that, given that the mental is dependent on, but distinct from, the physical, and given that the mental can have physical effects, we should rest content with the idea that all there is to causation is counterfactual dependence.[9]

This is the cart before the horse. If you are serious about mental causation, you will want your ontology of the mental and its relation

---

9 Imagine Descartes responding to Princess Elizabeth by noting that (1) science is committed to mental–physical causal interaction, so (2) if a conception of causation is inconsistent with this possibility, it should be jettisoned and replaced with one—counterfactual dependence, perhaps—that secures the desired result. The question is not *whether* the mental affects the physical, but *how* this could be possible given assumptions about the mental and the physical that are in play.

to the physical to mesh with a viable ontology of causation. Onto-logical problems are solved by getting clear on the ontology, not by shifting the goalposts.[10]

10 After completing much of this chapter, I discovered Anjan Chakravartty's 2005 and 2007, chap. 4, in which a related line of argument is advanced and defended in a way more sensitive to issues in the philosophy of science.

CHAPTER 7
Relations

It appears that a single thing, which must be imagined as some sort of interval (*intervallum*) existing between two things, cannot exist in extramental reality, but only in the intellect. [This appears to be the case] not only because nature does not produce such intervals, but also because a medium or interval of this sort does not appear to be in either of the two things [it relates] as in a subject, but rather *between* them where it is clear that there is nothing which can serve as its subject.

(Peter Auriçol, *Scriptum super Primum Sententiarum*, fols 318 v a–b; quoted in Brower 2009)

Socrates is similar to Plato by the very fact that Socrates is white and Plato is white . . . Yet, despite this, the intellect can express these many absolute things by means of concepts in diverse ways: in one way, by means of an absolute concept, as when one says simply 'Socrates is white' or 'Plato is white'; in a second way, by means of a relative concept, as when one says 'Socrates is similar to Plato with respect to whiteness'.

(William of Ockham, Quodl. VI, q. 25 in *Opera Theologica* ix, 679.)

7.1 Towards an Ontology of Relations

Paris is north of Marseilles, Simmias is taller than Socrates, Edward VII is the son of Victoria, exposure to the sun causes sunburn. These assertions are, or could be, true. But, given that they are true, what makes them true? The short answer is that 'Paris is north of Marseilles' is made true by Paris's being north of Marseilles, Paris's standing in the to-the-north-of relation to Marseilles; 'Simias is taller than Socrates' is true in virtue of Simias' being taller than Socrates,

Simmias' standing in the taller-than relation to Socrates; 'Edward VII is the son of Victoria' is made true by Edward's being the son of Victoria, Edward's bearing the son-of relation to Victoria; and 'exposure to the sun causes sunburn' is made true by sunburn's standing in the causal relation to the sun. Truthmakers for claims about relations include objects standing in relations.

If you think this is right, you will want to go on to say something about the *ontology* of relations. Although for much of philosophy's history every philosopher took a stand on the status of relations, the topic is not one that has been much addressed since the early years of the twentieth century when, thanks largely to Bradley and the idealists, relations briefly occupied centre stage. Relations are taken for granted. Indeed, Russell (1903, chap 26) established to the satisfaction of most philosophers that relational terms are ineliminable. Without a relational vocabulary, there would be truths about the universe you would be unable to express. This stands to reason if relations are ontologically fundamental. So here you have the makings of a proof for the reality of relations: representations of the universe include ineliminably relational elements; these relational elements require relational truthmakers.

Reasoning of this sort lies behind a large and influential body of work rooted in mainstream twentieth-century metaphysics. It is, for instance, a central feature of Quine's criterion of 'ontological commitment'. You are committed to the existence of entities ineliminably 'quantified over' in your best attested theories. Those theories are going to include mention of relations. Indeed, if the structural realists are right, the whole of science is most perspicuously expressed by means of a purely relational vocabulary (see, for instance, Dipert 1997, and Ladyman 2007.) If that were so, if you could represent the universe in a wholly relational vocabulary, then it would stand to reason that the universe is likely to be a wholly relational affair, its nonrelational features belonging to the realm of appearances.

You might react to this picture by noting that relations are most naturally thought to require *relata*, and *relata* might be thought to be prior to relations: relations need something to relate. One worry is that this kind of thought smacks of the kind of naïve a priorism that those favouring the primacy of relations find so annoying. Science can be seen to yield a relational picture of the universe. Philosophers have no

business telling scientists how the universe must or must not be. We ought rather to take up a cheerfully modest attitude toward scientific practice. History is littered with fashionable philosophical doctrines decisively overturned by science. When philosophy and science cross swords, only the fool sides with the philosopher. Science rules.

7.2 The Truthmaker Gap

Although the topic of relations is fascinating in its own right and far too little discussed in contemporary philosophy, I believe there are important lessons to be learned from ways different philosophers have grappled with the ontology of relations, lessons that extend to ontology generally. My guiding thought is two-pronged. First, it is a bad idea to assume from the outset that truthmakers for relational truths must be relations. It is easy to show that many relational truths, truths expressible only in a relational vocabulary, are made true by nonrelational features of the universe. Second, however, this in no way impugns the integrity of relational truths: relational truths are, after all, *true*. The philosophical mistake is to imagine that, if relations failed to be ontologically fundamental, you would have to 'deny the existence' of relations, imagining that it is false that anything stands in any relation to anything.[1]

This kind of mistake is endemic to linguisticized metaphysics. It results from setting the bar for realism impossibly high. Suppose you thought that statements of the form *a* is *F*—the tomato is red, for instance—were, if true, made true by some object, *a*'s, possessing some property, *F*. In a way, this is innocent enough. In ordinary speech, to say that something has a property, the property of being red, is just to say that it is red. But this kind of relaxed property talk leaves the ontology entirely open: what it is about the tomato in virtue of which it is true to say that it is red. (For that matter, it leaves open what it is about the universe in virtue of which it is true to say that this is a tomato.)

---

1 Jeffrey Brower (2009), in discussing Albert the Great's conception of relations, notes that 'Albert aligns himself with a medieval tradition of rejecting the view that for every distinct type of concept there is a distinct type of entity', a kind of correspondence that would hold only 'if our conceptual framework displayed an exact isomorphism to the structure of the universe. But it does not.' I am happy to align myself with Albert and his medieval tradition.

When you move directly from ontologically opaque reflections on objects and properties to substantive ontological conclusions, the result is shallow ontology. The Quinean maxim that our ontological commitments are revealed by noting what we quantify over in our best theories leaves untouched the question of what the universe might be like were those theories true. Only in fundamental physics do we approach anything like a one–one predicate–property mapping. Even there, matters are largely unsettled. Consider the quantum theory. Most physicists accept the quantum theory as true, but there is little agreement as to what the universe must be like if it *is* true. *This* is what Feynman means when he says, 'nobody understands quantum mechanics' (1965, 129). Taking properties with full ontological seriousness requires something more than sliding directly from predicates to properties, even when the predicates have the imprimatur of a going science.

Let me illustrate what I have in mind by reference to Descartes and Locke.[2] According to Descartes, the physical universe consists of a single extended substance: space itself. What you might consider to be objects—particles, trees, billiard balls, planets—are, in fact, local 'thickenings' of space, modes of extension, ways space is. Note that on such a view, motion—a billiard ball's rolling across the table, for instance—is wavelike, akin to the motion of your cursor across your computer screen.

Although Locke himself flirted with the idea that there was a single, extended substance, Locke ultimately differs from Descartes in embracing a corpuscularian conception of the physical universe (see Locke 1706). There are many corpuscles. Dynamic, interactive, *arrangements* of corpuscles serve as truthmakers for claims about trees, billiard balls, planets. Again, and to a first approximation, ordinary objects turn out to be modes, ways the corpuscles are arranged.

For my purposes, the important point here is that, for both Descartes and Locke, assertions about objects of the sort you would find in ordinary discourse and in the various sciences, are made true by

dynamic arrangements of the fundamental thing or things. It would never have occurred to Descartes or Locke to imagine that this casts doubt on trees, billiard balls, or planets. Assertions about such things can be, and often are, perfectly true. The deep story about what makes them true is what we seek in pursuing fundamental physics. That deep story is not something to be discovered via conceptual analysis or by asking what we quantify over in theories we accept. (I shall have more to say on this topic in chapter 8.)

I believe that Descartes and Locke are entirely representative of philosophical attitudes commonly accepted by philosophers from Plato, through the medievals, and into the Enlightenment. Somehow, in the foggy interval between Kant and Carnap, the attitude was bred out of philosophy. For millennia, philosophers operated with an implicit conception of truthmaking, a conception that remained unarticulated only because it was part of the very fabric of philosophy. Not until the late twentieth century did philosophers find it necessary to advance explicit conceptions of truthmaking. Although invocations of a truthmaking principle are nowadays regarded by some with deep suspicion, I consider the notion of truthmaking to be no more contentious than the notion of truth. What is contentious are particular philosophical *conceptions* of truthmaking.

I have said that a notion of truthmaking was long an implicit factor in philosophical reasoning. In fact it has always been implicit in *any* serious reasoning about the universe, including scientific reasoning. When you have a theory or equation that seems truly to describe the universe, it is natural to ask what it is about the universe in virtue of which the theory or equation *is* true. This is what is so unsettling about the quantum theory. It is hard to say what the truthmakers *could* be. Instrumentalism is unsatisfying precisely because it finesses truthmaking. Instrumentalism represents a kind of failure of nerve in the face of difficulties thrown up by reality.

7.3 Relations: An Opinionated History

The very familiarity of the truthmaking idea among philosophers who have puzzled over relations meant that its influence was felt rather than consciously invoked. One unhappy result was that

---

2 A reminder: I do not pretend to offer scholarly treatments of Descartes, Locke, or any of the other philosophers invoked here. My aim is not to explore subtleties of their views, but only to indicate a pervasive conception of the relation truths bear to the universe.

philosophers were bound to talk past one another. The history of philosophical views on relations provides a nice illustration of what I have in mind. My suggestion is that many philosophers from the medieval period onward zeroed in on a general conception of the ontology of relations, but struggled to articulate that conception. The result is an apparent proliferation of views, many of which might best be seen as attempts to articulate more or less a single thought. If you recast these views in terms of truthmakers—truthmakers for relational truths—their differences begin to blur.

Aristotle set the tone for ontologically serious discussions of relations by asking whether relations were substances or accidents.[3] The trouble is, as Aristotle and countless philosophers after Aristotle recognized, relations appear to be neither. Relations evidently depend on their relata. The taller-than relation that obtains when Simmias is taller than Socrates requires both Simmias and Socrates. In this regard relations resemble modes or accidents. But the identity of an accident depends on its bearer, and what exactly *bears* a relation among distinct relata? The bearer is not Simmias; the relation could cease to exist while Simmias is unaffected. Nor is the bearer Socrates. Could the bearer be *both* Simmias and Socrates? But how could an accident belong to distinct entities? Leibniz puts the matter this way:

The ratio or proportion of two lines L and M can be conceived in three ways: as a ratio of the greater L to the smaller M; as a ratio of the smaller M to the greater L; and lastly as something abstracted from both of them, that is to say as the ratio between L and M, without considering which is the anterior and which is the posterior, which the subject, which the object. In the first way of considering them, L the greater is the subject; in the second, M the smaller is the subject of this accident which philosophers call relation. But which will be the subject of the third way of considering them? We cannot say that the two, L and M together, are the subject of such an accident, for in that case we should have an accident in two subjects, with one leg in one and the other leg in the other, which is contrary to the notion of an accident. (1715, Fifth Paper, §47)

Take Simmias' being taller than Socrates. *Where* is this relation? What is its *bearer*? It seems neither to be in Simmias, nor in Socrates. If you

think of the relation as subsisting somehow *between* Simmias and Socrates, you are thinking of it as a kind of shadowy substance. The idea is difficult to credit. Substances, unlike accidents, are not dependent entities. Simmias' being taller than Socrates, however, depends on Simmias and Socrates.

These comments provide a hint as to why exactly you might find the ontology of relations deeply puzzling. To appreciate what philosophers as different as Peter Auriol, Ockham, and Leibniz see as the difficulty, think back to the account of substance set out in chapters 2 and 3 (see especially, § 2.4). Substances are bearers of properties, properties are ways substances are, modes. If this is right, it evidently follows that properties must be properties of simple substances: no property bearer, no substance, could have parts that are themselves substances.

I expect many readers to resist this thought. We are used to considering complex objects as bearing properties. My claim, however, was that you can account for thoughts about, descriptions of, ordinary objects without supposing that such objects themselves bear properties. Truthmakers for judgements about complex things and their characteristics—their properties in the relaxed sense—are dynamic, interactive arrangements of the fundamental things.

But why have qualms about complexes' bearing properties? Recall the discussion in chapter 2. If you start with the idea that substances are property bearers and properties are ways substances are, then it is not easy to see how a complex made up of substances could be a suitable property bearer. The complex is not itself a substance, it is made up of substances. As such, it seems not to be an apt bearer of properties. True, a complex can be various ways, but this is a matter of how its constituents are arranged.

If you can feel the pull of this thought, you are in a position to see why so many philosophers have had qualms about relations. Relations evidently depend on their relata. In this regard they resemble properties, modes. But then what bears a relation? Not the relata: relations hold *between* relata. If you locate relations between relata, however, they appear substance-like.

Relations, so regarded, are neither fish nor fowl, neither substances nor modes. You can elect to regard relations as *sui generis*, fundamental entities of a peculiar sort required to account for our best theories,

---

3 For present purposes, I use 'accident' and 'mode' interchangeably to denote particular ways substances are.

or you can try to find a way to accommodate judgements about relations, relational thought and talk, in a way that does not make relations ontologically fundamental.

I believe this pattern of reasoning occurs again and again in the epochs between Aristotle and Kant. The upshot is that, on the whole, philosophers saw relations as ontologically problematical. Leibniz speaks for many when he says

We are bound to say that the relation in this third way of considering it is indeed outside the subjects; but that being neither substance nor accident, it must be a purely ideal thing, the consideration of which is none the less useful. (1715, Fifth Paper, §47)

One reading of this passage turns Leibniz into a kind of projectivist about relations: relations are (mere) creatures of reason, the result of a mental comparison. This appears to have been the view of Peter Auriol quoted at the outset of the chapter. I shall suggest an alternative interpretation, but first, some comments on conceptions of relations that arose in response to Aristotelian worries concerning their ontology. Here is a partial taxonomy.

(1) Relations are ontologically fundamental; truthmakers for relational judgements are relations.
(2) Relations are creatures of reason, mental comparisons.
(3) Relations are 'reducible to' nonrelational features of relata.
(4) Truthmakers for relational judgements are nonrelational features of the universe.

I take (1) to be the default position today. Relational predicates are ineliminable, so, the thought goes, neither elimination nor reduction is in the cards. In contrast, I believe that most philosophers from Aristotle through Kant held one of the three remaining positions— and for good reason.

All of this is historically controversial, I admit. There are endless subtleties and distinctions to add to the mix. Suppose, however, there is something to this way of thinking about the history of philosophical conceptions of relations. What I am leading up to is the suggestion that philosophers who have defended versions of (2) and (3) on the basis of ontological considerations, were, often enough, struggling to say what is most perspicuously expressed by (4). They recognized,

albeit only implicitly, that truthmakers for relational judgements could not be distinctively relational *entities*. The difficulty is to express this thought without misleading, to express it without implying that relational judgements are invariably false, spurious, or otherwise defective.

Leibniz, in the passage quoted above, describes relations as 'ideal beings'. This sounds like anti-realism. He then goes on to note, however, that consideration of relations 'is none the less useful'. This suggests that relational judgements get something about the universe right, indeed that they are often true.

Now suppose you held something like (4), suppose you thought that truthmakers for relational judgements were, at bottom, nonrelational features of the universe. You might try to express this as Leibniz does, or by denying that relations exist, or by attempting to 'reduce' relations to something nonrelational. In that case, many proponents of (2) and (3) might best be understood as advancing towards (4). Options (2) and (3) appear attractive only as alternatives to (1). By the same token, the attractiveness of (1) owes much to the unattractiveness of (2) and (3). If this is right, we are moving in the direction of convergence on (4).[4]

7.4 The Ontology of Relations

The discussion thus far has been conducted at an inordinately high level of abstraction. I have suggested that you could have true relational judgements even though relations are not ontologically fundamental, even though truthmakers for those judgements are nonrelational features of the universe. The question now is, why should anyone think *that*? I have set out considerations originating with Aristotle that I claimed motivated generations of philosophers interested in fundamental ontology. Those worries began with the idea that relations, if they are ontologically fundamental, must be either substances or

---

4 Option (2) threatens to turn relations into mere appearances, and (3) brings with it a commitment to the kind of quixotic reductionist programme characteristic of an earlier era, and still pursued in some quarters today.

accidents. But perhaps relations are neither. Were relations onto-logically fundamental that is precisely what you would expect.

Here is how I see the dialectic. We need truthmakers for relational truths. A kind of historical consensus converges on the idea that truthmakers for such truths are nonrelational features of the universe. But how is this supposed to *work*? Given the ineliminability of relational predicates, isn't the most obvious option to suppose that truthmakers for relational truths are *relations*? If, for some reason, you dislike this possibility, the onus is on you to show how nonrelational features of the universe *could* make true every relational truth. My thought is that we should accept relations as ontologically fundamen-tal only if we have no plausible alternative. My reasons are just those that motivated all those philosophers who have found the ontology of relations deeply puzzling.[5]

Let me begin by setting aside one strategy for de-ontologizing relations. Suppose it is true that Simmias is taller than Socrates. Might the truthmaker for this truth be Simmias' possessing the property of being-taller-than-Socrates? Of course, if Simmias possesses this property, Socrates must, of necessity, possess the complementary property of being-shorter-than-Simmias, a somewhat puzzling neces-sity requiring a correlation of intrinsic features of distinct objects.[6] The idea is that you could replace relations with *relational properties*, properties possessed by objects individually.

Talk of relational properties, however, appears to be little more than a dodge, a thinly disguised way of talking about relations (Russell 1903; Moore 1919). Ascribing to Simmias the property of being-taller-than-Socrates is an oblique way of saying that Simmias is taller than Socrates. We need more than linguistic sleight of hand to dispense with relations as truthmakers for relational truths.

Consider six's being greater than five. It appears to be of the essence of six and five that they stand in this relation. Relations of this kind, so-called *internal relations*, were characterized by Moore (1919) as relations essential to their relata. Suppose *a* and *b* are related *R*-wise. If *R* is an internal relation, *a* and *b* could not fail so to be

---

5 Keith Campbell is one such philosopher; see his 1990, especially chapter 5; see also Parsons 2009.
6 A philosophical analogue of quantum entanglement! Another puzzling feature of the view is that it posits properties that by their nature involve other, distinct substances. How could an intrinsic property of something involve, of necessity, something else, some distinct particular thing?

related; you could not have *a* and *b* without having *R*. Otherwise *R* is external.

(R_i) If *R* internally relates *a* and *b*, then, if you have *a* and *b*, you *thereby* have *a*'s bearing *R* to *b*.

Internal relations were traditionally described as 'founded' on non-relational, intrinsic features of their relata. You could put this by saying that nonrelational features of relata are truthmakers for truths concerning internal relations. In creating all the objects, God *thereby* creates all the internal relations among those objects. Internal rela-tions are 'no addition of being'.

The identity relation seems clearly internal. If you have *a*, you have *a*'s being self-identical. Similarity or resemblance, too, looks internal, as does dissimilarity. To make it the case that *a* is similar (or dissimilar) to *b*, all God needs to do is create *a* and *b* *as they are*.

Might *all* relations turn out to be internal, might *all* relations be founded, might truthmakers for *all* relational judgements be nonrela-tional features of relata? The prospect seems hopeless. Consider Simmias' being taller than Socrates. Were the taller-than relation internal, it would have to be the case that Simmias and Socrates could not fail to stand in this relation: if you have Simmias and Socrates, you thereby have Simmias' being taller than Socrates. But that seems clearly false: Simmias could have failed to be taller than Socrates. Come to think of it, the same is true of similarity. Suppose Simmias and Thaetetus are similar. Surely they could have *failed* to be similar. To make it the case that Simmias and Thaetetus are similar, God needs to do more than create Simmias and Thaetetus: God needs, in addition, to *make them similar*.

Making Simmias and Thaetetus similar, however, is a matter of giving Simmias and Thaetetus similar *properties*. God might, for instance make both Simmias and Thaetetus pale. Now we have Simmias' being pale and Thaetetus' being pale. This—Simmias' being pale and Thaetetus' being pale—is enough to make them similar colourwise. The similarity relation is founded on intrinsic, nonrelational features of the relata.

Now consider Simmias' being taller than Socrates. God could create Simmias and Socrates without making the one taller than the other. To make Simmias taller than Socrates, God needs to do

something more: God needs to give Simmias a definite height and Socrates a definite height. Suppose God makes Simmias six feet tall and Socrates five feet tall; God endows Simmias and Socrates with these intrinsic features. God has thereby made Simmias taller than Socrates. The taller-than relation, like the similarity relation, is founded on nonrelational features of the relata.

You can now see how this conception of internal relations might be broadened. Internal relations are relations founded on nonrelational features of their relata; if you have the relata *as they are*, you thereby have the relation. Does this mean that internal relations are 'reducible to' nonrelational features of the relata or that internal relations do not really exist? I believe that these would be decidedly misleading ways of putting the point. Rather, relational truths, where the relations are internal, are made true by nonrelational features of the universe. You can have relational truths even if relations— relational *entities*—are absent from the fundamental ontology. My thought is that this, or something close to it, is what most philosophers who have sought to 'eliminate' relations or to reduce relations to or identify them with nonrelational features of the universe have had in mind.

Suppose I am on the right track. You could have relational truths without an ontology that includes relations, when the relations in question are internal. Internal relations are ontologically recessive. But what of truths concerning *external* relations? These, surely, require relational truthmakers.

Consider Paris's being north of Marseilles or the sun's causing sunburn. Spatial and causal relations seem ineliminably *relational*. You might try finessing spatial relations by supposing that Paris possesses the property of being-north-of-Marseilles, while Marseilles possesses the property of being-south-of-Paris. I have suggested already that the move from relations to relational properties is a kind of linguistic evasion. It looks as though the truthmaker for Paris's possessing the property of being north of Marseilles is Paris's being north of Marseilles, Paris's standing in the to-the-north-of relation to Marseilles.

Notice, however, that certain assumptions are required for spatial relations to count as external. Consider a Cartesian conception of space as a single extended substance. Particular objects would be

'thickenings' of particular regions of space. Objects, in this respect, would resemble patterns of illumination on your computer monitor. The motion of an object would resemble the motion of your cursor across your monitor or the motion of lights around a theatre marquee. The deep story about the motion of a billiard ball across the table would be that, at the fundamental level, nothing moves. Better, truthmakers for claims about moving objects—people, automobiles, footballs—would be non-moving features of the universe. If this sounds odd, think of the truthmaker for a judgement—a *true* judgement—that your cursor moves from one region of the screen to another.

One consequence of this Cartesian view would be that the identity of particular things is bound up with their spatial locations. God would not create objects, then locate them in space as you might locate pieces on a chessboard. In creating the objects, God would thicken regions of space. A particular thickening could not have existed in some other location any more than a particular freckle could have existed elsewhere on your skin. If you have the objects— the thickenings—as they are, you have them located as they are. Spatial relations, on such a view, are internal relations.[7]

The same result would follow if objects turned out to be disturbances or fluctuations in fields. Fields pervade space and, in this case, would play the substance role. If an object—a particle, say—were a fluctuation of a particular sort, it would be a mode, a way a field is, a way a field is in a particular spatial region. Its identity would be bound up with its location, and so for other particles to which it was spatially related. In creating *these* particles, God would create them located as they are. If you have the particles—the fluctuations—you have them situated as they are in space. Their spatial relations are internal.

Perhaps the only conceptions of space that make spatial relations external are 'container' conceptions, according to which space is a diaphanous medium within which objects are situated and move about. I shall return to this thought presently, but for the moment I want merely to use it to make the point that the idea that spatial relations are external incorporates a substantive conception of space and objects in space.

---

7 If you consider space and time to be unified, then these observations would extend smoothly to temporal relations.

What of causal relations? As noted in chapter 6, causality is most often taken to be a relation among distinct events, the causing event preceding its effect. Humean conceptions of causation and conceptions such as Armstrong's that appeal to contingent causal laws, depict causal relations as paradigmatically external.

I have already provided reasons to doubt the prevailing conception of causation as a relation among successive events. Start with the idea that causation is grounded in the dispositionalities of properties possessed by substances. A causing is a mutual manifesting of an ensemble of dispositions. The resulting manifestation could be regarded as the effect. If a property's identity is bound up with dispositionalities it would confer on its possessors, and if causal relations are the manifesting of powers, then causal relations would appear to be a species of internal relation (see Campbell 1990, 117–33). This impression is reinforced, if you think of manifestings of powers—causings—as cooperative, symmetrical affairs.

My suggestion in chapter 6 was that the received paradigm of causation—one billiard ball's striking a second and causing the second ball to roll in a particular way—be replaced with a pair of playing cards propped against one another so as to remain upright. Here you have a truly mutual manifestation of powers or dispositions possessed by the cards and the surface on which they are standing. This kind of mutual support can be found everywhere, even in dynamical settings such as the solar system. When a spoonful of salt dissolves in a glass of water, powers of the salt and the water manifest themselves mutually. And what of the billiard balls? No causing occurs until the first billiard ball comes into contact with the second. When that happens there is a mutual manifestation—a compression and decompression of each ball in a particular way for a temporally extended period—the upshot being an alteration in each ball's velocity. Notice, however, that *both* balls' velocities are altered, the causing is fully mutual, fully symmetrical.

Once you start thinking of causal relations in this way, the idea that causal relations are internal might begin to shed some of its initial implausibility. We have an evolving network of interrelated causings. The network includes indeterministic—spontaneous—elements, perhaps, but the causings are as internally related as interlocking pieces of a puzzle.

7.5 Take Relations . . . *Please*

Yes, I know, this is all much too quick. But you have to start somewhere. My thought is that we avail ourselves of a rich philosophical tradition that could be seen as encouraging us to regard relations as ontologically non-fundamental. Suppose you go with this tradition; suppose you ask whether relations might all be internal. Were that so, you could accept that truthmakers for relational truths are nonrelational ways the universe is.

One advantage of such a view is that it would account for a perennial philosophical perplexity over the nature of relations. Relations evidently depend metaphysically on their relata, so could not be substances. It is hard to see how a relation 'between' substances could be a mode or accident, however. We seem pushed to choose between founding relations on nonrelational features of relata, or adding them to the ground-level ontology. But how would this second option *work*? What exactly would we be imagining when we imagine relations 'out there' connecting their relata? The strangeness of this possibility will not strike you if you are content to posit an entity answering to each predicate. But this is just to return to the kind of linguisticized metaphysics that we ought, by now, to have outgrown.

Suppose you had reason to think that all relations *but one* were internal? Suppose, for instance, your conception of the universe made spatial relations, and only spatial relations, both external and fundamental. This might be a consequence of accepting a 'container' view of space, according to which objects stand in spatial relations to one another in virtue of occupying determinate spatial regions. In that case, you would have some reason to reconsider your conception of the universe. Faced with a choice between a conception according to which all relations are internal, and one requiring external spatial relations, you would have some reason to prefer the former, some reason to prefer Descartes, Leibniz, and Spinoza to Locke.

Is this metaphysics exerting illicit pressure on science? I do not think so. Metaphysics constrains the space of possible empirical theories. It does so, not in a vacuum, however, it does so in concert with the sciences. Fundamental ontological categories provide a taxonomy of reality, but these categories did not fall from the sky.

They arose from, and are coloured by, our continual give and take with the universe as we encounter it. I find the thought that relations might be ontologically fundamental disquieting, but I admit that could change were the universe shown to require them. Until then, I shall hold out for a universe with nonrelational truthmakers for irreducible relational truths.

If you disagree, if you are keen to have relations, then the ball is in your court. It is no good simply announcing the existence of relations alongside substances and properties. It is up to you to provide a tangible ontological story. From where I stand, the prospects do not look promising.

CHAPTER 8
Truthmaking

Sinbad himself may have fallen by good-luck on a true description, and wrong reason sometimes lands poor mortals in right conclusions: starting a long way off the true point, and proceeding by loops and zigzags, we now and then arrive at just where we ought to be.

(George Eliot, *Middlemarch*)

True and false are attributes of speech, not of things. And where speech is not, there is neither truth nor falsehood.

(Hobbes, *Leviathan*, Bk 1, chap. 4)

8.1 Why You Should Care

Truthmaking has made numerous appearances in previous chapters. The time has come to gather together strands of the discussion in a way that makes it clear what truthmaking encompasses and why the notion of truthmaking is central to ontological enquiry. The lack of concern with—even disdain for—talk of truthmaking in much contemporary philosophy is almost certainly due to a widespread unquestioning acceptance of Quine's criterion of ontological commitment: you are ontologically committed to whatever you ineliminably 'quantify over' in your best theories (see Quine 1948). If you cannot paraphrase or analyse talk of trees into talk of collections of particles or perturbations in the quantum field, then you are committed to the existence of trees as distinct from—something other than, something *in addition to*—collections of particles or perturbations in the quantum field.

Whatever its epistemological merits, this ostensibly innocent precept has proved disastrous for ontology. Its very familiarity makes it

difficult to appreciate its radical nature. Perhaps more than anything else, acceptance of Quine's dictum is responsible for the current fashion for fusing metaphysics and philosophy of language.

A prominent theme in the preceding chapters is the thought that it is a mistake to imagine that we are in a position to extract ontological lessons from ways we talk about the universe. The persistent tendency to conflate predicates and properties is just one example of what I have in mind. If it is true that this is a tree or that there is at least one tree in the quad, this must be because some object possesses the property of being a tree. The property of being a tree is a *property*, so the property of being a tree must possess the property of being a property. Everything is self-identical in virtue of possessing the property of being self-identical. What could be more obvious!

The *Book of Common Prayer* speaks approvingly of God as a being 'whose property is always to have mercy'. My suggestion is that this is little more than a lyrically inspired way of saying that God is always merciful. This kind of relaxed Episcopalian use of 'property' is completely unobjectionable when all that is intended in ascribing a property to an object is to register that something is true of that object. The mistake is to begin with such relaxed property ascriptions, then proceed to draw substantive metaphysical conclusions from them.

When a predicate, 'F', applies truly to some object, a, this might be because a possesses a property, the property of being F. But it could also be the case that 'F' applies truly to a because 'F' is a determinable predicate satisfied by a's possession of a property answering to a determinate of that determinable predicate. An electron, for instance, has mass by virtue of having some determinate mass, m. Or 'F' might truly characterize a particular dynamic, interactive arrangement of substances. The substances possess properties, but the arrangement does not. Nevertheless, the arrangement as it is serves as truthmaker for this application of 'F'.

You might wonder how predicates that *do* designate properties are to be distinguished from those that, while applying truly under the right circumstances, do *not* designate properties. As the electron example suggests, 'determinateness' is one factor. Properties are 'maximally determinate'. The philosophical distinction between 'determinables' and 'determinates' is not a distinction among kinds

of *property*, but a distinction among kinds of *predicate*. A particular tomato is a determinate shade of red, R. Its being this determinate shade of red is truthmaker for assorted determinable truths: 'The tomato is red', 'The tomato is coloured', 'The tomato has a property'. The tomato does not have, in addition to the property of being some determinate shade of red, R, the property of being red, the property of being coloured, and the property of having a property. Rather, it is true that the tomato is this determinate shade of red, that it is red, that it is coloured, and that it has a property by virtue of its being this determinate shade of red: one truthmaker, R, many predicates, many truths.

An astute reader will have noticed that I could not really think that the tomato's being a determinate shade of red is a matter of the tomato's possessing a property. On the conception of substance and property in play, properties must belong to substances and substances must be simple; substances lack substantial (proper) parts. On such a conception the only predicates that designate genuine properties are those characterizing the fundamental thing or things.

This means that the question whether a given predicate designates a property is at bottom an empirical question. The properties are going to belong exclusively to the fundamental things. You start by treating tomatoes as substances and characteristics of tomatoes as properties. In so doing you treat tomatoes as unified simples, 'substances-by-courtesy', quasi-substances. As you move away from this preliminary conception, you revise your assessment as to what the real properties are.[1]

Pretend for a minute that the universe is granular, made up arrangements of 'insensible' indivisible particles possessing various properties. When you arrange particular kinds of particle in the right way, you get a tomato (and all the many parts we commonly regard as making up a tomato). Truths about the tomato, including truths about its colour, shape, mass, density, and the like are made true by this arrangement. This does not mean that the tomato is merely a collection of particles—or even a collection of particles 'arranged tomatowise'. No, the particles must be of the right kinds and organized in the right way (and perhaps have the right kinds of

---

1 See Campbell 1990, chapter 6, for extended discussion of this point.

history and bear the right kinds of relation to other particles). Particles making up the tomato can come and go—within limits. Judgements about the tomato are made true by this evanescent, dynamic, interactive arrangement. Because tomatoes are not proper substances, they lack hard-edged identity and persistence conditions.

Wait! I have characterized the tomato as an arrangement of particles. But *which* arrangement? Looked at closely, the arrangement is not dense and continuous, but cloud-like. The arrangement includes a penumbra of particles. Which of these belong to the tomato, which do not? Many candidate arrangements would fill the bill.

The difficulty, if it is a difficulty, is not confined to tomatoes in granular universes. Suppose, for instance that the universe were continuous, suppose the tomato were a local thickening of space–time. The thickening gradually abates around the edges. Where do you draw the line? Where does the tomato end and empty space begin?

The idea that, unless we can answer such questions, tomatoes are in doubt, is unpromising. When you point to a tomato, when you pick one from the vine, when you slice a tomato in the course of preparing a salad, the object you handle might be cloud-like. There is no saying where the tomato leaves off. But this makes no difference to our ability to discuss, handle, and dispose of tomatoes. Our *talk* of tomatoes is designed to get us *to* tomatoes. What you find when you look closely at tomatoes might be cloud-like, it might resemble a swarm of bees, or a bruise, or a stain spreading on the tablecloth. There is no reason to think that application conditions for the tomato predicate could be spelled out by reference to particles or thickenings of space–time. The tomato predicate gets us to tomatoes, then physics takes over.

Although I have tried to be careful in discussing these matters, I am aware that it is easy to miss the point. Again, pretend that the universe is granular. I do not say that a tomato is merely a collection of particles. Any tomato is a *particular* collection, a *particular* dynamic arrangement of *particular* kinds of interacting particle, some of which occupy a penumbra of particles, each of which has a *particular* trajectory through space–time. If you do not believe me, pick up a tomato and inspect it closely. Look at the tomato under a scanning electron microscope. As you look more closely, you begin to lose the tomato and find the particles.

Does this mean that tomatoes are mere appearances? Are particles duly arranged mere appearances? You picked up the tomato and looked at it closely. Did you pick up a mere appearance? You *start* with the tomato, one you picked from the vine and placed in a basket. You learn that the tomato is cloud-like, you learn that the truthmaker for talk of the tomato resembles a swarm of bees. Some of the outlier bees follow the swarm from the vine to your basket, some do not. The question, which bees belong to the swarm, which particles belong to the arrangement, is not a question about tomatoes, it is not a question that calls for an answer so long as your interest lies in the tomato.

To reiterate: there is no suggestion here that you could analyse or translate talk of tomatoes into talk of particle arrangements, or that you could express application conditions for 'is a tomato', or express all the truths about tomatoes, using a vocabulary drawn exclusively from fundamental physics. More particularly, I am not arguing that you could forgo talk of tomatoes in favour of talk of arrangements of particles, or fields, or anything else. The route from ordinary talk of tomatoes to truthmakers for that talk is not linguistic.

Linguisticism aside, our access to the universe is not in general mediated by language. You can talk about tomatoes, describe them, and discourse on their nature. But you can also plant and harvest tomatoes, ingest them, step on them, heave them at politicians, buy and sell them, slice them up for salads. You can teach someone what a tomato is by providing a definition, but you can also do so by handing over a tomato for closer inspection.

I point all this out simply to make it clear that the truthmaking relation, whatever it is, is not an analytic, conceptual, linguistic, or quasi-linguistic relation. You could be, and probably are, perfectly adept at talking or thinking about tomatoes, without having any very clear conception at all of the deep story about tomatoes: what it takes to answer to judgements about tomatoes. You might, for instance, think that tomatoes are ultimately arrangements of particles, but be utterly deceived. Tomatoes might be fluctuations of the quantum field, wrinkles in space–time, jiggling super strings, or something as yet undreamt of by modern physics. The deep story is the story of the truthmakers for assertions about tomatoes.

8.2 Truthmaking as Entailment

Truth is a relation between a truth bearer and a truthmaker. But what sort of relation?

In discussing relations, I have followed tradition, distinguishing internal and external relations. Internal relations, such as the taller-than relation, hold in virtue of intrinsic features of their relata. If you have Simmias' being six feet tall and Socrates' being five feet tall, you *thereby* have Simmias' being taller than Socrates. Truthmakers for the judgement that Simmias is taller than Socrates are nonrelational features of Simmias and Socrates. Relations that are not internal would be external. Familiar examples of external relations include spatial relations, temporal relations, and causal relations. Socrates' and Simmias' being two miles apart, for instance, apparently depends on something more than intrinsic, nonrelational features of Simmias and Socrates. You could have Simmias and Socrates *just as they are*, without its being the case that they are two miles apart.

At least this is how it *seems*. In chapter 7 I offered reasons for thinking—or at any rate hoping—that all relations are internal, that is, that truthmakers for every relational truth are nonrelational features of the universe. Were *that* so, relations would not be ontologically fundamental; there would be relational truths, but no relational truthmakers.

For purposes of the present discussion, I propose to set aside this possibility. Whatever the standing of external relations, I follow Armstrong (2004, 9) in regarding the truthmaking relation as an internal relation. If you have a truth bearer—a judgement that *p*, for instance—and you have its being the case that *p*, you have truthmaking, you have the judgement's being true, you have the judgement's being made true by its being the case that *p*. (What its being the case that *p is*, how the universe must be for it to be the case that *p*, is completely open. I shall return to this point in § 8.4.)

This conception of truthmaking differs from conceptions that appeal to entailment or necessitation. According to John Bigelow and David Lewis, 'truth supervenes on being' (Bigelow 1988; Lewis 2001). Once you have the universe as it is, you *thereby* have all the

truths. Expressed differently: the universe, or ways the universe is, entail or necessitate the truths. Thus, according to John Bigelow,

> whenever something is true, there must be something whose existence entails that it is true. The 'making' in 'making true' is essentially logical entailment . . . Suppose there to be something which is proposed as a truthmaker for some truth. And suppose it is admitted that the existence of that thing does not entail the truth in question. This means that it is logically possible for that thing still to exist, even if what is actually true had not been true. In the actual world, *a* exists and *A* is true, say; but in some other possible world *a* might still exist, even though *A* is not true. There must surely be some difference between these two possible worlds! So there must be something in one of these worlds which is lacking in the other, and which accounts for this difference in truth . . . . If something is true, then there must be, that is to say there must exist, something which makes the actual world different from how it would have been if this had not been true.[2]

All this seems reasonable enough, even innocuous, until you pause to consider how it is supposed to *work*. Let Socrates' being five feet tall be the truthmaker for 'Socrates is five feet tall'. How could Socrates' being five feet tall logically entail *anything*? Both Socrates and Socrates' being five feet tall evidently belong to the wrong category. Entailment is a *semantic* relation standardly defined by reference to truth: if *p* entails *q*, then *p* could not be true if *q* is false; the truth of *p necessitates q*'s truth. Both *p* and *q* are items possessing truth values. But truthmaking is a paradigmatic *cross-categorial* relation, a relation between a truth—a *bearer* of truth, a judgement—and something, some way the universe is, that makes it true.

Indeed, Bigelow notes that entailment is 'a relation between propositions'. Thus truthmaking is not a matter of 'an object [entailing] a truth; rather, it requires that the proposition that an object exists entails the truth in question' (1988, 126). So talk of objects, facts, or states of affairs entailing the truth of judgements is just an abbreviated way of saying that *representations* of those objects, facts, or states of affairs entail the truth of the judgements in question. In that case, there would be no mystery in the idea that the truthmakers entail truths. A representation—Bigelow would say a proposition—that the

---

2 Bigelow (1988, 125–6); see also Mulligan, Simons, and Smith (1984); Fox (1987); Armstrong (1997, chap. 8). The discussion of truthmaking as entailment here is based on Heil (2003a, chap. 7).

fact that Socrates is five feet tall obtains entails that 'Socrates is five feet tall' is true.[3] The difficulty is that, to perform this function, the proposition doing the entailing would need to be true! And this brings us back to our starting point: what relation must obtain between the pertinent proposition and whatever makes it true in order for the proposition to *be* true?

An account of truthmaking that introduces propositions as intermediaries between truthmakers and truth bearers appears unpromising. Mediating propositions, it seems, would *themselves* require truthmakers. If these mediating propositions are made true by virtue of being logically entailed by further mediating propositions, nothing has been explained. If, in contrast, their being made true requires no further mediation, you have a truthmaking relation that does not involve entailment. And if you are entitled to invoke such a relation in this case, why not invoke it at the outset? The relation between propositions introduced to entail truths and the relation between these propositions and whatever it is about the universe that makes *them* true, is the very same relation.

Appealing to propositions as ingredients in an account of truthmaking brings with it at least three problems:

(1) the problem of providing an account of the ontology of propositions consistent with the theoretical role propositions are meant to play;

(2) the problem of connecting propositions to judgements taken to 'express' them;

(3) the problem of spelling out the relation propositions bear to whatever it is that answers to them.

Problem (3) is just the truthmaker problem all over again, the very problem propositions were introduced to help solve. Problems (1) and (2) are, however, no less challenging. As in the case of properties, talk of propositions can be innocuous. When propositions are invoked in the course of advancing substantive philosophical theses, however, their credentials deserve special scrutiny (see § 12.5

---

3 Propositions afford excellent examples of the kinds of dodgy entity posited by philosophers to solve various problems. I discuss propositions in chapter 12. For the moment, I shall play along and frame the discussion in terms of propositions.

below). In any case, propositions have no place in a serious discussion of the ontology of truthmaking.

8.3 'Truth Supervenes on Being'

Entailment appears unhelpful in explicating truthmaking. What of 'supervenience', what of Bigelow's 'truth supervenes on being' slogan? In one sense, the slogan is unobjectionable. The truths cannot change without a change in the truthmakers. But this leaves open the nature of the intended supervenience relation (see Heil 1998). Suppose supervenience is understood as some form of *necessitation*: if the *A*'s supervene on the *B*'s, the *A*'s are necessitated by the *B*'s. Notice that unless more is said about the posited necessitation relation, its introduction sheds little light on the nature of truthmaking. In any case, attempts to explicate truthmaking in terms of necessitation of whatever sort are, I believe, fundamentally ill-conceived.

Start by reflecting on the Tarskian truth schema

(T) '*S*' is true if and only if *p*

So, if 'Snow is white' is true, snow is white; and if snow is white, 'Snow is white' is true. Is the previous sentence trivially true? Many philosophers have thought so. But consider the second of these conditionals, the conditional going from right to left. Could it be that snow's being white necessitates an *entity*, a sentence, presumably in the form of an utterance or an inscription? There are endless ways the universe is, was, and will be for which there are no corresponding utterances or inscriptions. You might object: snow's being white necessitates, not the *sentence* 'Snow is white', but the sentence's *truth*. For this to be so, however, you need *both* the sentence (meaning what it does) and snow's being white. Given *these*, you have the sentence's being true.

This point, which seems obvious, is easily obscured by philosophers' talk of propositions. If propositions are *abstracta*, they seem ontologically innocent. They take up no room, they are never in the way, they are whatever we say they are. So it appears harmless to allow a one–one correspondence between ways the universe is and propositions. Even if you are unashamed to play this game, however,

you might still want to see truthmaking as just the sort of relation proposed in this chapter—or so I hope.

The deeper point is that talk of the necessitation of truths by truthmakers suggests an asymmetry that is not in fact in evidence. To get the truthmaking relation, to get something's being made true, you need *both* a truthmaker, the universe's being some way, *and* something made true, a truth bearer, a judgement, a representation of the universe's being that way. If you have these, you have truth, you have the truthmaking relation. Now it becomes obvious that the truthmaking relation is an internal relation: if you have the relata, you thereby have the relation.

If this is right, it is easy to see why it is a mistake to characterize truthmaking in terms of entailment, or supervenience, or necessitation. Simmias' being taller than Socrates is not entailed or necessitated by Socrates' being five feet tall, Simmias' being taller than Socrates does not supervene on Socrates' being five feet tall. If there is no temptation to say such things about the taller-than relation, there should be no temptation to say the analogous things about truthmaking. If God makes it the case that snow is white *and* that there is a judgement that snow is white, God has thereby made it the case that the judgement is true, that it is made true by snow's being white.[4]

So truthmaking is an internal relation: if you have the relata—a truth bearer, a judgement that *p*, and it's being the case that *p*—you have the relation. If you have both of these, you could say that truth 'supervenes' in the way you could say that Simmias' being taller than Socrates supervenes on Simmias' being six feet tall and Socrates' being five feet tall. But this is to read supervenience as something like the truthmaking relation, in which case it is unilluminating to invoke supervenience in explicating truthmaking.

Really, appeals to supervenience to explicate truthmaking result in a *loss* of information. The philosophical notion of supervenience is a purely modal notion.[5] When the *A*'s supervene on the *B*'s this could be because

---

(1) the *A*'s *are* the *B*'s;
(2) the *A*'s are made up of the *B*'s;
(3) the *A*'s are determinates of determinable *B*'s;
(4) the *B*'s causally suffice for the *A*'s;
(5) the *A*'s and *B*'s both causally depend on the *C*'s.

To this list we could now add the truthmaking relation as understood above. But then it is clear that appealing to supervenience to *explicate* truthmaking is a bad idea.

8.4 Truth Bearers

In regarding truthmaking as an internal relation, I am adopting the view that truthmaking is 'founded' on intrinsic features of the relata. But what are the relata? David Armstrong (1997, chap. 8; 2004) has argued that truthmakers are states of affairs in his proprietary sense of the term: substances instantiating universals at times, Socrates' being warm on Tuesday, for instance. Others have argued that truthmakers are tropes (Mulligan, Simons, and Smith 1984). As I have characterized truthmaking, however, your acceptance of the need for truthmakers does not, by itself, commit you to any definite ontology (see Parsons 1999). What it does commit you to is there being truthmakers of *some* kind for truths concerning the universe. (I leave open for the moment the question whether every truth requires a truthmaker.) Perhaps you could say that truthmakers for judgements about the universe are particular ways the universe is. This is at least consistent with the ontological picture developed thus far.

How are we to think of truth *bearers*? I have talked casually about judgements, utterances, assertions, statements, sentences, inscriptions, and representations. I have recommended avoiding talk of propositions. If you are attracted to propositions, then explicating truth bearers is a breeze. Propositions are intrinsically meaningful entities, a proposition's meaning is essential to it. Propositions are tailor made to serve as truth bearers.

Once you let go of propositions, matters are trickier. Suppose you utter the sentence, 'Snow is white'. The sentence, we say, is true, and it is made true by snow's being white. But now what, exactly, *bears*

---

[4] A useful exercise here would be to pause and reflect for a moment on what is required in order for there to be snow and snow to be white.
[5] Here I am indebted to Brian McLaughlin; see also Kim 1984, 1990.

this truth? The sentence? Your utterance of the sentence? What you had in mind in uttering the sentence? What you meant in uttering it? Whatever truth bearers are, they must be meaningful. Their being meaningful equips them to be bearers of truth—or falsehood. The question as to what exactly bears meaning—assertions, sentences, utterances of sentences, agent's intentions in uttering sentences—is not one I intend to pursue here. Nor shall I offer an account of what is required for an act or entity to mean what it means. I propose, instead, to use the neutral term 'judgement' to designate whatever it is that is meaningful hence truth apt. 'Judgement' is conveniently ambiguous as between an action and an entity or state.

Given a judgement that p, and the universe's being a particular way, the p-way, the judgement is true: made true by the universe's being as it is. Given the universe as it is and the judgements as they are, you thereby have the judgements' being true (or false).

Now the question arises, given the judgement that p, *what* state of the universe is relevant? If I judge truly that there is a jacaranda in the quad, *what* makes this judgement true? I have argued that truthmakers are ways the fundamental things are. But quads and jacarandas are not to be counted among the fundamental things. So what *connects* the fundamental truthmakers for my judgement with my judgement?

A long tradition of correspondence theories of truth is replete with attempts to tie features of truth bearers to 'corresponding' features of the universe. But I have repeatedly insisted that you cannot read off the truthmakers from syntactic features of our judgements. We are all adept at uttering truths without the slightest inkling as to the deep story about their truthmakers. In learning to talk about the universe we master application conditions for various terms. These application conditions take us *to* the universe without, in most cases, revealing very much at all about the nature of whatever it is about the universe that satisfies those application conditions. We learn application conditions for 'is a jacaranda', for instance, without thereby being in a position to say much about whatever it is that answers to these application conditions. Perhaps the tree in the quad is a dynamic collection of particles; perhaps it is a fluctuation of the quantum field; perhaps it is a local thickening of space.

Maybe this is all you need. Maybe all you need is the idea that our thoughts about the universe, entangled as they are with extensive

non-verbal interactions with our surroundings, suffice to circumscribe ways the universe is, leaving open the deep story about those ways. Thoughts are directed onto a universe the ontology of which remains largely hidden to thinkers. We manage well enough by picking out salient similarities and differences; but this does not require a grasp of the *bases* of these similarities and differences.

Consider your judging that this tomato is red, and pretend that application conditions for 'is a tomato' and 'is red' are satisfied. Your judgement concerns a particular portion of the universe, a particular object's being a particular way. You could be wholly ignorant of the deep story concerning that object, the tomato, and its features, the truthmakers for your judgement. You no more need to have access to that story in order to entertain thoughts about the tomato than you do in order to pick it up, slice it, and toss it in a salad. Indeed, parallels between thinking about an object and interacting with it non-verbally—picking it up for instance—are instructive. In each case your intercourse with the object depends only on your *getting to it*, not on your appreciating *what* precisely it is.

The point of these remarks is simply that the idea that truthmaking is an internal relation between a judgement and some way the universe is, does not require a mapping of syntactic features of the judgement, or some regimented version of the judgement, and features of whatever it is about the universe that makes the judgement true. The judgement, or the judgement in concert with multifarious non-linguistic interactions with the universe, takes us to the universe, takes us to some way the universe is. If application conditions for the judgement are satisfied, then the way the universe is makes the judgement true. In making the judgement, you could be entirely ignorant of what that way is. You can know that 'Snow is white' is true while remaining wholly ignorant of the nature of snow or snow's whiteness.

8.5 Truthmaking and Quantification

We are accustomed to accepting judgements such as

(1) there is a tree in the quad
(2) the tree in the quad is a jacaranda
(3) trees are plants

to be true. When you represent these in first-order logic, you 'quantify over' trees. It is most unlikely that we could analyse away talk of trees, or translate claims about trees into the language of fundamental physics. This suggests to many philosophers that we are 'committed to the existence of trees'.

Philosophers who take this line are impressed by Quine's aforementioned criterion of ontological commitment (Quine 1948). But what is it to be committed to the existence of something? If being committed to the existence of trees means accepting judgements such as (1)–(3) as true, the thesis is ontologically innocent. Ontology enters only when you turn your attention to the nature of the truthmakers. Knowing what truths you accept is one thing; knowing what it is about the universe in virtue of which those truths are true is another matter altogether. All this leaves the ontology of trees entirely open; the nature of truthmakers for judgements about trees might turn out to be almost anything. And this is consistent with my insistence that there is no a priori or analytic way to get from these judgements and what they mean to an accounting of what makes them true.

But now we seem faced with a puzzle. If there is an epistemological or conceptual gap between truths and truthmakers, how could the relation between the two be internal?

Earlier I suggested that the application conditions at work in our judgements get us to the universe, then the universe takes over. Suppose you judge, 'This is a stick' as you pick up a stick on the path in front of you. If the object satisfies the application conditions for talk of sticks, then what you have said is true. What makes it true is what you would find were you to get at the deep story about this object. *Which* object? The one you are holding.

You could say, then, that what makes it true that snow is white is snow's being white: the application conditions for 'Snow is white' are in fact satisfied. This reveals nothing about the nature of snow, however. For that you must turn to science. If you think that snow's being white is what makes 'Snow is white' true, snow's being white is the truthmaker for 'Snow is white'. Now you turn to chemistry to tell you what snow and snow's whiteness amount to, the deep story about snow. The deep story eventually brings in fundamental physics. Indeed I have argued that the job description

of fundamental physics includes informing us as to the real nature of truthmakers for judgements about the universe.

I have been revelling in talk of the 'nature of the truthmakers', but our interest is only rarely in such matters. We are—in the various sciences and in ordinary life—content to describe the universe truly without asking for an accounting of what precisely it is about the universe in virtue of which those truths are true. Why then concern ourselves about the deep story, the nature of the truthmakers?

Why indeed? The concern arises in philosophy only when our topic is, as it is here, fundamental ontology. How are we to think of the truthmakers? Does the universe exhibit ontological 'levels' that supply 'higher-level' truthmakers for the various sciences? I have denied that this is so, and I shall, in the next chapter, remind you of reasons that have been in play throughout the previous chapters that support this denial. We have descriptive, taxonomic levels, perhaps, and levels of organization, but no levels of reality.

Let me try to put all this in a slightly different way. Suppose it is true that a particular tomato is red. It seems right to say that the truthmaker for this judgement is the tomato's being red. And for the tomato to be red, there must be a tomato. So isn't Quine right, after all: you *can* discover your ontological commitments by considering what you quantify over.

My concern is to block the following sort of move. Suppose that you cannot paraphrase away talk of tomatoes, suppose you cannot analyse assertions about tomatoes in terms borrowed from fundamental physics. Does this mean that you are committed to the existence of tomatoes *in addition to* the existence of the fundamental things related as they are in space and time? Must you say, for instance, that there are particles related to one another just as they are *and* there are tomatoes? *This* is the move I hope to discourage. Talk of the existence of tomatoes is ontologically innocent when your aim is to get clear on the truths, when your interest is in *tomatoes*. But when you engage in ontology, when your goal is to get clear on the fundamental truthmakers, it is ill advised.

All this suggests an unwelcome distinction between fundamental truthmakers and truthmakers that are non-fundamental. In fact, the fundamental truthmakers *are* the truthmakers. We expect these to be revealed by fundamental physics. These make true every truth

concerning the universe, whether or not these truths are thinkable or expressible in terms borrowed from physics. This is not something that follows from the sciences, and in that sense it is not an empirical claim. It is a substantive ontological thesis that I believe finds support both in scientific practice and in serious ontological engagement.

8.6 Truth and Existence

When you ask about the existence of things, when you ask whether there are trees, or planets, or cassowaries, for instance, you are asking, in effect, whether judgements concerning such things are true. This is the nugget of truth in Quine's criterion of ontological commitment. Answering such questions is often—not always: consider logic and mathematics—an empirical endeavour. But if deciding whether trees, planets, or cassowaries exist is a matter of deciding whether to accept, as true, judgements about trees, planets, and cassowaries, then talk of what exists is, at bottom, an oblique way of talking of accepted truths.

Serious ontological judgements, in contrast, look beyond the accepted truths to being itself, *being qua being*. Less dramatically, ontology concerns what D. C. Williams called the *elements* of being.

Return to the tomato. When you examine a tomato closely, you realize that it is made up of assorted interrelated parts variously organized. When you examine these parts, you find that they, too, have parts. This pushes you towards a conception of particles of the kind discussed in fundamental physics. Further empirical investigation might reveal that this granular picture is misleading: what you regard as distinct particles might turn out to be perturbations in fields, or eddies in space–time.

Fundamental physics is fundamental in the sense that it gives us a general, all-encompassing accounting of truthmakers for truths that have truthmakers. Truths of physics are not 'more true' than truths emanating from everyday experience and the various sciences. In that regard, truths of physics are in no way privileged. Still, in providing an account of the truthmakers, fundamental physics provides a glimpse of what lies at the basis of all the worldly truths: the buck stops with fundamental physics. If the truths of physics entail every such truth, however, this is not a matter of deducibility, not a matter

of 'higher-level' terms or theories being deducible from, or replaceable by, physics.

To this point our investigation would have relied exclusively on the sciences. Philosophical reflection takes over when you ask about existence or being generally. In previous chapters I have defended the thesis that being comprises propertied substances: truthmakers for truths about the universe will be substances duly propertied. You might amend this characterization: duly propertied and *duly arranged*. This reflects the idea that the universe is not merely a collection of substances, but a definite *arrangement* of substances: this yields substances, properties, *and relations*. I accept the amendment with the proviso that relations could turn out to be ontologically recessive: truthmakers for judgements concerning relations could turn out to be nonrelational ways the universe is.

My quarrel with Quine does not concern what exists. It concerns the province of ontology. Think of Quine's criterion of ontological commitment as a kind of syntactic regimentation of accepted truths. Think of ontology as attempting a general accounting of truthmakers for these truths. The accounting will be sensitive to what the sciences tell us, more particularly, what fundamental physics tells us. If the sciences are concerned with the truths, ontology is concerned with the deep story about truthmakers for these truths.

Let me try to summarize the preceding remarks in a way some readers might find helpful. Suppose you ask, do tomatoes exist? Planets? Electrons? Yes, certainly, barring mass hallucination, we take it to be true that there are tomatoes, planets, electrons. All these things exist.

Does this watch exist? Do the watch's parts exist? Yes and yes. There is a watch *and* there are its parts duly arranged. What would be misleading would be to insist that *in addition* to the parts organized as they are there is a watch. I have put this in terms of truthmaking: the truthmaker for 'there is a watch' is the collection of parts (duly organized, with their various interactions, histories, &c). So there is a watch and there are its parts, but not these parts *plus* the watch.

Note that the fact, if it is a fact, that the watch could survive the replacement of parts, is irrelevant here. The truthmaker for 'this is a watch' is going to be some particular collection. If the watch could survive the demise of the collection, this means only that the identity

and persistence conditions for watches and for collections of parts differ. What counts as 'the same' collection over time differs from what counts as 'the same' watch.

Considerations of this kind lead philosophers to conclude that 'composition is not identity': the watch cannot be identified with a particular collection of parts. But this is not what I am suggesting. My contention, rather, is that, when it is true here and now that this is a watch (or that there is a watch here), the truthmaker for this judgement is this very collection (duly arranged, &c). For a watchmaker to make it the case that there is a watch here, the watchmaker need only take these parts and put them together in the right way.

The story about the watch is meant to generalize. God creates everything, supplying truthmakers for all the truths requiring truthmakers, by creating the fundamental things propertied as they are, and giving them their trajectories through space–time.[6] In so doing, God thereby creates everything there is: the tomatoes, the planets, the electrons, and, yes, the Beethoven concertos.

The important point here is that you can know that it is true that there is a tomato here—and, if you like, know that tomatoes exist—while remaining in the dark as to the nature of the truthmakers. This is why there is little or no prospect of expressing precise, sharp-edged identity conditions for tomatoes, planets, or watches (electrons are another matter) in terms of the fundamental truthmakers and their relations.

8.7 A Truthmaker for Every Truth?

Some philosophers regard the notion of truthmaking as hopeless. To such philosophers I have little to say. The truthmaking idea strikes me as utterly fundamental, something we all grasp immediately in grasping the notion of truth, a notion we are unlikely to be in a position to explicate in simpler, more transparent terms. Other philosophers are attracted to the thought that *every* truth, or at least every truth that deserves to be taken seriously, must have a truthmaker. Although it

6 This way of putting it leaves open the possibility that identity conditions for the fundamental things include their trajectory through space–time, and it allows for spontaneous occurrences.

goes without saying that I am on shaky ground here, I prefer a more modest conception of truthmaking. Before saying what I have in mind, let me suggest one way to think about the philosophical utility of appeals to truthmaking.

When you encounter philosophers defending metaphysical theses, it is often useful to ask those philosophers to say something about what truthmakers for their theses might be, what it is about the universe that might make those theses true. Consider a philosopher who, in the course of developing some philosophical thesis, advances various 'modal' claims, for instance, claims about what might have been the case. You might fairly ask what would make such claims true. Nowadays philosophers commonly respond by invoking alternative universes: *possible worlds*. To say that Napoleon might have won at Waterloo, is to say that there is a 'possible world' in which Napoleon, or a surrogate Napoleon, prevails. To say that, had Wellington died in Portugal in 1811, Napoleon would have remained Emperor, is to say that in 'nearby worlds' in which Wellington (or a general very like Wellington) dies, Napoleon (or his doppelgänger) remains Emperor.

Adverting to possible worlds is one way to respond to requests concerning truthmakers for such judgements. Sadly, many philosophers who find talk of possible worlds indispensable are also quick to tell us that such talk is merely figurative. It is not that these other worlds are really *out there*. Rather, appeals to alternative universes in explicating modal judgements provide us with a convenient way of *understanding* these judgements. But if these other universes are merely useful fictions, we are left wondering what the truthmakers for all those modal judgements might be. Perhaps the truthmakers are intrinsic features of our universe, *the* universe. But if they are, *which* features? Appealing to make-believe universes, possible worlds, looks like a delaying tactic, a philosophical evasion.

It is worth pointing out that David Lewis, the source of the present-day infatuation with possible worlds, believed in their existence as real, concrete entities, universes on a par with our universe. If you are so inclined then, in the interest of ontological candour, you might want to describe the worlds as I have described them above: *alternative* universes. Labelling those other universes 'possible'

is misleading to the extent that it suggests that the other universes are somehow *merely* possible, less real, 'less actual' than *our* universe.[7]

In my experience philosophers who talk possible worlds talk while shrugging off the existence of universes other than our own, typically feel a twinge of guilt when asked for replacement truthmakers. This is as it should be. The need for truthmakers, patent in the empirical sciences, is one important aspect of the need for candid ontology. It is, for instance, the difficulty in imagining what truthmakers for normative judgements might be that leads many philosophers to find 'moral realism' unappealing. Realism evidently requires mind-independent truthmakers. In the normative case it is not obvious what these could be. This is especially so if the earlier discussion of properties is on the right track: the only properties are properties of the fundamental things. If moral realism is taken to require distinctively moral *properties*, realism looks hopeless. If, in contrast, realism requires only moral truths and mind-independent truthmakers, matters might not be quite so bleak.

You could, then, see the demand for truthmakers as a theoretical presumption, not as the expression of an iron-clad 'truthmaker principle'. The demand resembles invocations of Ockham's Razor meant to keep us ontologically honest and on our toes. Indeed, it would be difficult to understand the popularity of 'abstract entities'—numbers, sets, propositions—in the absence of an implicit acknowledgement of the importance of truthmakers. We are thought to need numbers, for instance, not merely to supply a subject matter for mathematical judgements, but to supply truthmakers for such judgements as well.

I shall have more to say about *abstracta* in due course. For the present, I merely note that the idea that we need *abstracta*—numbers, for instance, or sets—is largely a byproduct of the Quinean doctrine of ontological commitment. We quantify over numbers and sets, so we are taken to be committed to the existence of such entities. Contingent occupants of our universe appear not to be entities of the right *kind*. So the entities to which we are committed must reside

elsewhere, in a spectral realm of incorruptible objects untainted by contingency.

8.8 Logic and Mathematics

Stewart Shapiro provides an example of the kind of reasoning I have in mind.

According to model-theoretic semantics, the singular terms of a mathematical language denote objects, and the variables range over a domain-of-discourse. Thus, mathematical *objects*—numbers, functions, sets, and the like—exist. This is what I call *realism in ontology*. A popular and closely related theme is the Quinean dictum that one's ontology consists of the range of the bound variables in properly regimented discourse. The slogan is 'to be is to be the value of a bound variable'. (Shapiro 1997, 4–5)

A running theme in this and in previous chapters has been that we have no good reason to follow Quine, and every reason not to. Looking at what you quantify over reveals, at most, *truths* to which you are committed. What the ontology is, what the truthmakers are for these truths, is another matter, one to be tackled, if at all, only in the pursuit of fundamental physics. This, it appears, leaves no room for *abstracta*, in so far as *abstracta* are taken to be items existing of necessity and apart from space and time. Although mathematics permeates physics, few philosophers, and fewer physicists, suppose that physicists will one day stumble across numbers or sets amongst the fundamental things.

But then where does this leave mathematics? Mathematical judgements—and for the moment I am restricting the discussion to judgements of pure mathematics, not its applications—are frequently true, but what, if not numbers, makes those judgements true? One possibility is that mathematical truths are made true by 'structures', relationships among objects in the universe.[8] Such relationships would hold whatever the nature of the objects. Thus truthmakers

---

7 Lewis himself seems not to have regarded the worlds or universes as truthmakers for modal judgements. (See § 5.9 above.) On Lewis's view, the truthmakers for such judgements were intrinsic features of our universe, *the* universe, features grounding relations to other universes. Still, similarity relations require relata, so those other universes would need to be real.

8 See, for instance, Resnik 1997 and Shapiro 1997. See also Bigelow 1988 and Armstrong 1997 (chap. 11). The positions defended by Resnik, Shapiro, Bigelow, and Armstrong differ in important ways and go well beyond my informal characterization here. My interest is not in showing that structuralism and its variants are wrong, but in suggesting that a simpler, more satisfying account of mathematical (and logical) truths is available.

for mathematical judgements would be found in any imaginable universe. (Or any universe of sufficient cardinality, any universe with a sufficient number of objects.) This would account for our confidence that truths of mathematics will hold *however* the universe might turn out to be. Whatever universe God elects to create would be a universe with the kinds of structural feature required for mathematical truths—provided only that the universe were populated by enough entities.

This is an interesting thought, but I confess that I do not understand it. I do not understand how structures *could* serve as truthmakers for mathematical truths, and I do not understand why anyone would think that such structures, whatever they are, must be present in any universe, must, so to speak *mark reality*. This pathetic admission of failure is not meant as an argument against such views. Rather, I shall offer what strikes me as an independently plausible alternative: mathematical truths do not require truthmakers. Mathematical truths are not true in virtue of *any* way the universe is, the universe's being some way rather than some other way does not make them true. Mathematical truths hold, *whatever* ways the universe is.

The idea is not that mathematical predicates are analysable in a manner that eliminates reference to numbers. Nor is this an 'error theory': it is not that the semantics of mathematics commits us to non-existent truthmakers so that mathematical judgements are literally false. Rather, mathematical truths, as Ross Cameron puts it, belong to a class of truths that 'make no demands on reality' (Cameron 2010). When God creates the universe, God must do many things: God must create particular substances, arrange them, and endow them with properties. But God does not, *in addition*, need to make the universe such that everything is self-identical or that 7 + 5 = 12.

Cameron is drawing on and extending the work of Agustín Rayo, who defends what he calls 'trivialism'.[9]

Trivialism is the view that true sentences of pure mathematics have trivial truth-conditions (and that false sentences of pure mathematics have trivial falsity-conditions). According to the trivialist, nothing is required of the world in order for the truth-conditions of a mathematical truth to be

satisfied: there is no intelligible possibility that the universe would need to steer clear of in order to cooperate with the demands of mathematical truth. This means, in particular, that there is no need to go to the world to check whether any requirements have been met in order to determine whether a given mathematical truth is true. So once one gets clear about the sentence's truth-conditions—clear enough to know that they are trivial—one has done all that needs to be done to establish the sentence's truth. (Rayo 2010, 240)

'Trivialism', the label, has unfortunate connotations. To say that a judgement's or sentence's truth-conditions are trivial in Rayo's distinctive sense is not to say that they are obvious or immediately cognizable: 'getting clear about the truth-conditions of a given mathematical sentence can be highly non-trivial. So determining whether the sentence is true is not, in general, a trivial affair' (Rayo 2010, 240–1). A judgement's truth-conditions are trivial 'if any scenario in which the truth-conditions fail to be satisfied would be unintelligible' (Rayo 2010, 240).

These points hold as well for truths of logic. Logical truths 'make no demands on the universe'. In creating the universe, God needs to do nothing to ensure that truths of logic hold. However the universe is or might be, they could not fail to hold.[10] God does not create the universe, then provide it with a mathematical structure or stir in *abstracta*, numbers and sets.

Your immediate reaction to trivialism might be to wonder how mathematics could have any point if it traffics in trivialities. We all depend on mathematics in describing and explaining worldly goings-on in science and in everyday life. How could this possibly work if mathematics is trivially true? Mathematical equations guide the construction implements we use to negotiate the universe and explore its nature. We *depend* on mathematics in a million ways.

If you are suspicious that trivialism throws the baby out with the bath water, that trivialism makes it utterly mysterious how mathematics applies in the universe, you should be no less suspicious of the idea that the subject matter of mathematics is a realm of *abstracta*. Why should truths about such a realm have application to denizens of space

---

9 Rayo 2010. An early stab at trivialism can be found in the *Tractatus*: 'All the propositions of logic say the same thing, to wit nothing' (Wittgenstein 1921 § 5.43).

10 I leave aside the question whether this implies that there is 'one true logic'. See Rayo 2010 for discussion.

and time? If you thought the Demiurge used mathematical archetypes as patterns for worldly states of affairs, a conception of this kind *might* make sense. But few philosophers who regard themselves as Platonists about mathematical entities, are prepared to think as Plato recommends in the *Timaeus*.

So how might trivialism make sense of the central role of mathematics in scientific descriptions and explanations of goings-on in the universe?

Start with truths of logic. As Rayo notes, 'in learning a logical truth one increases one's ability to distinguish between intelligible and unintelligible scenarios, and therefore one's ability to use old information in new ways' (Rayo 2010, 256). Imagine that you have acquired seventeen apples, and that you know this because you have counted them. You are now in a position to do a number of things. 'You are able to determine whether you got short-changed at the market, or whether there are enough apples for your recipe. You are also able to answer questions of the form "How many apples?" (Rayo 2010, 256). Suppose that, in addition, you have, and know you have, twenty-nine pears. You are now positioned to perform a number of further tasks.

Perhaps you are able to combine these two cognitive accomplishments in the service of a single task. You might, for instance, be in a position to determine whether there are more apples than pears. But other tasks might elude you. Say you know that every relevant piece of fruit is an apple or a pear, and that no piece of fruit is both an apple and a pear. Then you have all the information you need to answer questions of the form 'How many pieces of fruit?'. But you may still not be in a position to use the information at you disposal for that particular task, at least not immediately. (Rayo 2010, 256)

What you need is knowledge of a particular logical truth, namely that if there are exactly seventeen apples and twenty-nine pears, and if nothing is both an apple and a pear, then there are forty-six pieces of fruit.

In performing the relevant computation, do you acquire novel information about the world? It is tempting to say that you do since you will learn there are forty-six pieces of fruit. But . . . the right thing to say is that you don't. For you already knew that every piece of fruit is an apple or a pear (but not both) and that there are seventeen apples and twenty-nine pears, and part of

what it is for that to be the case is that there be forty six pieces of fruit. In carrying out the computation, your cognitive accomplishment consists not in the acquisition of new information, but in the ability to deploy old information in new ways. Before you carry out the computation you are unsure about whether a scenario in which there are, say, thirty-six pieces of fruit could be genuinely intelligible while respecting the information you already had about apples and pears. What the computation reveals is that it is not. You have increased your ability to distinguish between intelligible and unintelligible scenarios, and this gives you the ability to see how to answer questions of the form 'How many pieces of fruit?' in light of the information you had at your disposal all along. (Rayo 2010, 256-7)

All this could be expressed in the formula

$$\exists_{17}! x(\text{Apple}(x)) \land \exists_{29}! x(\text{Pear}(x)) \land \neg \exists x(\text{Apple}(x) \land \text{Pear}(x)) \supset$$
$$\exists_{46}! x(\text{Apple}(x)\text{Pear}(x))$$

(In English: 'If there are exactly seventeen apples and there are exactly twenty-nine pears and nothing is both an apple and a pear, then there are exactly forty-six things that are either apples or pears.')

So far as truthmaking is concerned, a trivialist sees no deep difference between logic and mathematics. Neither purports to gives us descriptions of the universe, although both prove to be indispensable in our providing illuminating descriptions of the universe. Assertions of logical truths and pure mathematics alike have trivial truth conditions, neither 'places any demands' on the universe. 'The difference is simply that the language of mathematics enjoys expressive resources that the language of logic lacks' (Rayo 2010, 258). Thus, in accepting

$$17 + 29 = 46$$

you acquire a *general* ability,

the ability to rule out as unintelligible a scenario in which there are seventeen *F*s, twenty-nine *G*s and anything other than forty-six *F*s-or-*G*s (provided no *F*s are *G*s). And of course an improved ability to sort out the intelligible from the unintelligible is important because it gives us an improved ability to transfer insights from one context to another. (Rayo 2010, 258-9)

This brief discussion does not begin to do justice to Rayo's full account of logic and mathematics or Cameron's extension of the

thesis to necessary truths generally. I mean it only as a sketch of what appears to be a most promising way of reconciling the character of certain kinds of necessary truth and the natural thought that truths require truthmakers. What I have sought to do here is to indicate what truthmaking amounts to when truths are made true by ways the universe is, why it is philosophically sensible to expect that substantive truths have truthmakers, and how it might make sense to set aside judgements expressing necessary truths.

8.9 Non–Trivial Necessities

Cameron holds that necessary truths quite generally are characterizable by the fact that they 'make no demands' on the universe, their truth in no way depends on how the universe is. Although I am prepared to agree with Cameron, it is worth noting that there might be two subtly different ways in which a particular truth could be said to 'make no demand on the universe'. Truths of logic and mathematics provide examples of the first of these.

Suppose, however, that there were general constraints or limits on being itself. One such constraint might be that any universe must be a universe of substances (property bearers) and properties (ways substances are). Were this so, it would seem to qualify as a necessary truth on Cameron's view. A universe of propertied substances is not simply one from among various options. *Any* universe must be a universe of propertied substances. For there to be something rather than nothing, there must be something that is some way or other. But my contention that substance and property are ontologically fundamental evidently falls short of the kind of epistemic standing accorded truths of logic and mathematics. You could agree that, if true, my ontological precept is necessarily true, while remaining agnostic as to its truth. In fact, I think it possible, even likely, that reality is constrained in still more definite ways. Reality—*any* reality, being *qua* being—might need to possess certain features that appear patently contingent to us mere mortals.

Imagine a stone mason setting out to build a cathedral. Although the mason could erect many kinds of building, the kinds of building the mason could erect are constrained by the nature of

the available materials: given *these* materials, some kinds of building would be out of the question. Now imagine God's setting out to create the universe, perhaps *ex nihilo*. God's creation might be analogously constrained: there are universes God could not create, just as there are walls a mason could not build using stones and mortar. Various fundamental features of the universe—the gravitational constant, for instance, or the electron's mass—might strike you as deeply contingent. But maybe they are not contingent. Maybe the universe is fine-tuned as it is, not because things happened to fall out as they did during and immediately after the Big Bang, or because God so ordained it, but because God or the Big Bang had no choice.[11]

Some readers will be sceptical. My contention, however, is not that this is how it is or must be, but only that this is how it *could* well be. The idea that modal status—and, in particular, what the contingencies are—can simply be eye-balled is without merit. Being itself, reality, could have a texture, a nature that allows for variation, but not for indefinite or unconstrained variation. If this is a genuine possibility, then it appears that there could be truths about reality *qua* reality that hold of necessity, yet are made true by reality's possession of particular characteristics. I mention this possibility, not merely because I regard it as genuine. I mention it because it strikes me as overwhelmingly *likely*. If electrons and gravitational fields have definite natures, why not reality itself? And if reality has a nature, if this makes sense, then reality grounds the deepest necessities of all.

Would such necessities 'make no demands on the universe'? No and Yes. Keeping with the theological metaphor, it is not that God must do something definite to see that the universe satisfies them: *whatever* God does will satisfy them. Still, the necessities, if apt, would reflect substantive features of being itself, its *grain*. It is not that reality could have been different in various ways, or that there might be different possible realities. It is, rather, that there could be substantive constraints on ways of being. If these are trivialities, their being trivialities is far from trivial.

---

11 If so, this would no more be a limitation on God than God's inability to create round squares or make the sum of 7 and 5 something other than 12 would be.

8.10 Truthmaking and Serious Ontology

I had better not leave the discussion at this level of abstraction, but end by reasserting the chapter's central focus: truthmaking is an internal relation that holds between a truth bearer—a judgement or representation—and a truthmaker—some way the universe is. Our recognizing truths *as* truths requires our recognizing that application conditions for particular judgements are satisfied. It is no part of our doing this that we grasp the nature of, the deep story concerning, the truthmakers. Just as you can identify tomatoes, pick them up, sort them, and juggle them without any understanding at all as to their nature—without any understanding concerning what it is about the universe in virtue of which they answer to judgements pertaining to tomatoes—so you can entertain thoughts of such things without a grasp of what they are.

One way to put this point would be to say that you can appreciate the truth *conditions* for judgements about the universe without appreciating, without even having a capacity to appreciate, the nature of the truth*makers*. Failure to recognize this simple idea has plagued recent philosophy (Dyke 2008). In particular, it has led to inflated ontologies that are thought, quite mistakenly, to be required by the various sciences. This is the topic to which I shall now turn.

CHAPTER 9
Reduction, Kinds, and Essences

Physics develops the taxonomy of its subject-matter which best suits its purposes: the formulation of exceptionless laws which are basic. . . . But this is not the only taxonomy which may be required if the purposes of science in general are to be served: e.g., if we are to state such true, counterfactual supporting generalizations as there are to state. So, there are special sciences, with their specialized taxonomies, in the business of stating some of these generalizations. If science is to be unified, then all such taxonomies must apply to the same things. If physics is to be basic science, then each of these things had better be a physical thing. But it is not further required that the taxonomies which the special sciences employ must themselves reduce to the taxonomy of physics. It is not required, and it is probably not true.

(Fodor 1974, 114)

All the great Business of *Genera* and *Species*, and their *Essences*, amounts to no more but this, That Men making abstract *Ideas*, and settling them in their Minds with names annexed to them, do thereby enable themselves to consider Things, and discourse of them, as it were in bundles, for the easier and readier improvement, and communication of their Knowledge, which would advance but slowly, were their Words and Thoughts confined only to Particulars.

(Locke 1690, iii, iii, 20)

I hope that readers will gradually get used to my principles and recognize their truth, before they notice they destroy the principles of Aristotle.

(Descartes, Letter to Mersenne, 28 January 1641)

9.1 Big Pictures

The foregoing chapters have focused on topics in fundamental ontology. I have sketched a Big Picture that I consider has considerable independent plausibility. I would like to think that readers shocked by such a pronouncement have inhaled too deeply of the fumes of linguisticized metaphysics. The position developed here is, as I see it, continuous with those advanced by Descartes, Locke, Spinoza, my early modern and Enlightenment heroes. So, although I am prepared to admit defeat, I am much less prepared to concede that the view I defend is a philosophical outlier. But, enough special pleading: back to the Big Picture.

The nurturing and fine-tuning of Big Pictures is a traditional philosophical endeavour. The question is, how might the adequacy of such endeavours be measured? One way to gauge a Big Picture would be to determine how it fares in particular cases: Big Pictures are corroborated in their applications, their success in opening the way to sensible solutions to otherwise intractable difficulties. A Big Picture affords a perspective from which various problems become salient while others recede into the background, a perspective that constrains the range of possible resolutions of problems it exposes.

To the extent that this evokes visions of a realm of philosophically untainted data, it is misleading. True, Big Pictures are sometimes accepted or rejected in response to longstanding purely philosophical puzzles. Think of Descartes's solution to worries about the status of minds in the material realm, or Kant's response to puzzles about the applicability of mathematics in the spatio-temporal universe of experience. Philosophical puzzles, however, are invariably puzzles in light of some other, some alternative, Big Picture. Appealing to one Big Picture as a solution to a particular puzzle amounts to an implicit comparison of Big Pictures.

Under the circumstances, Wittgenstein's contention that philosophical problems are largely of our own making looms large. We philosophers are, most of the time, occupied with local puzzles. We face a philosophical difficulty and find our way past it by concocting a theory or, more likely, adding an epicycle to a received theory. In this way we gradually build up our houses of cards, shaky edifices propped

up by armies of graduate students aiming to become assistant professors and assistant professors angling for tenure. Whatever our Big Picture, it resides in the shadows, largely unacknowledged, but no less influential for that. Our Big Picture is a comfortable pair of spectacles we forget we are wearing.

Still, we *are* philosophers and, as such, we are bound to reflect from time to time on what we professionally take for granted. When we do so, we are in a position to appreciate the extent to which an implicit Big Picture can colour our conception of what the important problems are and what constitutes the space of permissible solutions. In this book, I have endeavoured to place in relief some of the most prominent assumptions definitive of the Big Picture underlying the enterprise of analytical philosophy as it has evolved during recent decades. My strategy has been twofold.

First, I have offered a specific ontological scheme against which you can measure your own ontological preferences. If nothing else, this encourages us all to acknowledge those preferences and bring them into the foreground where they might lose their innocence and their aura of inevitability.

Second, I have tried to call attention to particular assumptions by making them explicit—'assembling reminders'. This chapter implements both strategies. I begin with a discussion of what has come to be called 'non-reductive physicalism', a doctrine originally invoked in the philosophy of mind in response to attempts to identify states of mind with physical states, most particularly brain states. Once the tenets of non-reductive physicalism are spelled out, however, it is clear that the doctrine has sweeping implications and applications beyond psychology to the 'special sciences' generally.

Despite its early successes, philosophers have discovered numerous difficulties with non-reductive physicalist solutions to problems the doctrine was put together to address. In response, some philosophers advocate embracing its contrary, *reductive* physicalism, a species of traditional materialism (Kim 1989, 2005), some have endorsed species of dualism (Chalmers 1996). Many more philosophers, however, have persisted with the project, preferring to tinker with what they regard as ancillary matters—the nature of causation, for instance—hoping to salvage what, to an outsider, would appear to be a sinking ship (Gibbons 2006; Bennett 2007). The interesting point is the

extent to which all parties to this debate regard the Big Picture underlying it as wholly innocent and all but unavoidable. I hope to convince you that it is neither.

9.2   Non-Reductive Physicalism: Background

If nothing else, the advent of non-reductive physicalism in the 1970s affords a tidy case study of the role played by Big Pictures in shaping the philosophical landscape. Non-reductivists regard their background assumptions as entirely innocent philosophically. Indeed, non-reductivist metaphysics has become so deeply identified with a prevailing conception of the sciences that philosophers under its sway imagine that calling it into question is tantamount to calling the sciences into question—and who are philosophers to challenge the sciences?

If you back up, however, and look dispassionately at motives behind arguments for non-reductive physicalism, you can see how a single vital insight has been subverted by ontological manoeuvers that are far from innocent. Here, as in earlier chapters, Quine is the villain. If you think that ontology can be extracted from our ways of talking about the universe, you are susceptible to mistaking truths about our ways to thinking about or describing the universe for truths about the universe. Thus the linguisticization of metaphysics.

I have, more than once, called attention to the tendency to slide from arguments concerning *predicates* to conclusions concerning *properties*. Hilary Putnam's retitling his 'Psychological Predicates' (1967) as 'The Nature of Mental States' when it was republished in his collected papers (1975) provides a nice illustration of what I have in mind. Predicates and categories are one thing, ways the universe is are something else altogether. The thought that, if mental predicates are not analytically decomposable into physical predicates, mental predicates must designate a family of non-physical properties if they designate anything, is characteristic of a pattern of thought that continues to be broadly influential, even—or especially—among philosophers who would, if pressed, disavow it.

In the sections that follow, I discuss arguments advanced by three stars in the non-reductivist firmament, Donald Davidson, Jerry Fodor, and Richard Boyd. My preliminary conclusion will be that,

while interesting and important, the arguments fall well short of establishing the metaphysical conclusions they are most often taken to establish. I am confident that Davidson never meant to defend such conclusions. I am less sure of Fodor and Boyd. In Fodor's case, it is tempting to speculate that, after providing cogent reasons to doubt that the special sciences could be reduced to fundamental physics, Fodor began to believe enthusiastic but misguided admirers, who persisted in reading arguments concerning scientific categories and taxonomies as metaphysical arguments with momentous ontological implications. The rest is history, a desultory history that we should do well to put behind us.

9.3   Davidson: First Impressions

In 'Mental Events', maybe the most influential paper of the period under discussion, Davidson argued that, although mental *types* could not be reduced to or identified with physical types, every particular mental event, every *token* mental event was in fact identical with some token physical event (Davidson, 1970). Many of Davidson's readers took him to be defending a thesis according to which

(1)   every mental *token*—every particular mental *state* or *event*—is identical with some physical token, some particular physical state or event; however,

(2)   mental *types*—and here the assumption was that types are *properties*, and properties *universals*—are distinct from physical types, physical properties.

Putting these together meant reading Davidson as defending the idea that mental events are events possessing mental properties, physical events are events with physical properties. One and the same event could be both mental and physical by virtue of possessing both a mental property and a physical property. Thus, Jaegwon Kim, an influential expositor of Davidson's position:

Imagine a Davidsonian universe of events: all of these events are physical, and some of them are also mental. That is to say, all events have physical

properties, and some of them have mental properties as well. Such is Davidson's celebrated 'anomalous monism'.[1]

This way of understanding Davidson evolved eventually into a characterization of non-reductive physicalism generally. Mental properties were taken to be dependent on, indeed, necessitated by, but distinct from, physical properties. Davidson had spoken of *supervenience*, arguing that the mental supervenes on the physical. As Davidson put it, 'This might be taken to mean that there cannot be two events alike in all physical respects but differing in some mental respect, or that an object cannot alter in some mental respect without altering in some physical respect' (Davidson 1970, 215).

Appeals to supervenience are often traced to G. E. Moore, although Moore did not use the term. Prior to its invocation by Davidson, supervenience had been deployed most prominently by R. M. Hare (1952, chap. 5) in explicating the relation between normative, 'prescriptive', judgements and natural, 'descriptive', judgements. According to Hare, paintings could differ aesthetically, only if they differed in some natural, non-aesthetic way; agents who differed morally, must differ in some non-moral, natural respect. Although Hare was a moral 'prescriptivist', not a moral realist, not a believer in mind-independent moral properties, this facet of his appeal to supervenience was either ignored or downplayed. It is worth noting, then, that whatever Hare meant by supervenience, it would not have been a relation among 'natural properties' of agents or works of art and 'normative (moral or aesthetic) properties' (see Hare 1952, 94).

Considerable philosophical brainpower, always a scarce commodity, was expended in the effort to plumb the metaphysical nature of the supervenience relation.[2] As standardly defined, supervenience is a 'modal' notion. That is, it tells us that, where the A's and B's are 'families' of properties, if the A's supervene on the B's, the A's and B's co-vary, without specifying the basis of the covariation. The A's and B's could covary because the A's *are* the B's, for instance, or because

---

1   Kim 1989, 269. Earlier (§ 2.4) I discussed the implausibility of supposing that events as Kim understands them—substances' possessing properties at times—could *themselves* possess properties, a serious problem for Davidson were Kim's gloss on Davidson apt.

2   See § 8.3 above. See Heil 1992, chap. 3 for a dreary blow-by-blow account, and Heil 1998 for a follow-up. Nuanced modal characterizations of supervenience were originally developed by Kim in the early 1980s; see the essays collected in Kim 1993.

the A's are made up of the B's, because the A's are caused by the B's. The trouble is, none of these relations is a good fit for the non-reductive physicalist's conception of the relation of the mental and the physical.

Philosophers' infatuation with supervenience eventually subsided, not because it had become clear what it might mean for mental properties to supervene on physical properties, but owing to a lack of appreciable progress on the matter. Obsession with supervenience was replaced by a morbid fascination with 'realization', When the mental supervenes on the physical, it was thought, this is because mental properties or states have physical 'realizers'.

In retrospect, it is hard to see the move from talk of supervenience to talk of realization as progress. The realization relation has remained contentious and elusive. In some quarters the attitude has seemed to be, 'look, we *know* that mental properties (or maybe states) are realized by physical properties (or states); our inability to produce an account of a respectable realization relation no more calls this fact into question than our inability to provide a satisfactory account of the causal relation calls causal relations into question'. The non-reductivists' Big Picture at work.

A more pressing difficulty facing philosophers convinced that mental properties are physically realized is a difficulty usually pegged to Davidson. Suppose particular mental events, 'mental tokens', are identical with particular physical events, 'physical tokens'. If mental properties or types are distinct from, even though dependent on, physical properties or types, you can always ask, when a mental event causes a physical event, whether it does so in virtue of its mental properties or in virtue of its physical properties. Remember, an event is mental if it has a mental property, physical if it has a physical property, and every mental event, every event with a mental property, is a physical event, an event with a physical property. The physical property in question is the realizer of the mental property. (Or maybe it is the resulting physical state that realizes the mental state. Mere details.) But now the question arises, when a mental event has a physical effect, does it do so *qua* mental or *qua* physical?

A tomato is red and spherical. When it rolls, it rolls because it is spherical, not because it is red, even though it *is* red. Maybe mental events are like this. Maybe mental events, events with mental

properties, can have physical effects, but not because they are mental. And in fact, given Davidson's commitment to the two-pronged thesis that causation requires subsumption under 'strict law', and that strict laws are the exclusive province of fundamental physics, it would seem that, whenever a mental event has a physical effect, it has that effect by virtue of its physical properties, not its mental properties: the mental must be causally impotent, real, perhaps, but epiphenomenal! Maybe the mental *piggybacks* on the physical, going along for the ride, but contributing nothing to the effort.

Philosophers offering accounts of the realization relation have hoped to provide a way of insinuating mental properties into the causal nexus but in a way that maintains their distinctness from their physical realizers.[3] Without going into detail, I shall note only that the results have not been encouraging. My contention, however, is that non-reductive physicalism in general, and ever more complex accounts of the realization relation in particular, are beside the point. Both issue from a fundamental mistake that, once exposed, obviates the need for discussion of the details. In seeing what has gone wrong, you can start to appreciate one of the beneficial philosophical roles of truthmaking.

9.4   Davidson: Beyond First Impressions

Let us look more carefully at the doctrine advanced by Davidson in 'Mental Events'. Earlier I noted that Davidson is most often depicted as defending the thesis that mental events are identical with physical events, but mental properties, although dependent on physical properties, are not identifiable with physical properties. An event is mental by virtue of possessing a mental property, physical by virtue of possessing a physical property. Thus, when a mental event has a physical effect, one question is whether the event has this effect by virtue of being mental ('*qua* mental'), by virtue of having a mental property, or by virtue of being physical ('*qua* physical').

3 See, for instance, Shoemaker 2001, 2007, 2011; Pereboom 2002; Wilson 2011. For discussion, see Heil 2003b, 2011.

The matter is pressing because Davidson insists that 'events related as cause and effect fall under strict deterministic laws' *and* that 'there are no strict deterministic laws on the basis of which mental events can be predicted and explained', no 'strict' psychological or psycho-physical laws. But then how could an event '*qua* mental' ever be causally efficacious? As many, many philosophers have noted, it looks as though, by Davidson's own lights, an event's possessing a physical property *pre-empts* the contribution of any supervenient or realized mental property. Davidson notoriously waffles on the point.

This would be worrying indeed were it not for the fact that none of the previous two paragraphs has much to do with Davidson's considered view. In 'Mental Events', Davidson goes to great lengths to say, and say repeatedly, that events are mental or physical 'only as described'. The mental–physical distinction, for Davidson, is not ontologically deep. The distinction is one of conception only, not what philosophers of earlier eras would have called a *real distinction*. I could say much more about this, but I shall resist the temptation except to note that it is a mistake of a fundamental sort to describe Davidson as a physicalist, non-reductive or otherwise. Yes, according to Davidson, anything, any particular object, state, or event, that can be described using a psychological vocabulary could be described, as well, in a vocabulary borrowed from fundamental physics.[4] But this in no way impugns the psychological description. *That* description could be perfectly true. If you want to know the deep story about the truthmakers, however, you will need to consult physics.

You might put this by saying that there are countless ways to describe or 'taxonomize' the universe. These various ways are not replaceable by equivalent descriptions framed in terms at home in fundamental physics. In one sense physics is the 'science of every-thing': any particular occurrence has an explanation expressible the vocabulary of fundamental physics. Even so, truths, including causal truths, expressible in the fundamental physical vocabulary are not the only truths. Biological descriptions, psychological descriptions, sociological descriptions, meteorological and geological

4 Davidson *also* notes that, because any physical event (his example is 'the collision of two stars in distance space') could be referred to via a mental description, every physical event is a mental event, a symmetry endorsed by Spinoza (Davidson 1970, 211).

descriptions, can be no less true, no less accurate, than descriptions found in fundamental physics. In this regard physics is in no respect privileged.

According to Davidson, then, when a mental event (your intending to raise your hand, for instance) causes a physical event (your hand's moving in a certain way), a particular event answering to a mental description causes a particular event answering to a physical description. *Both* events satisfy descriptions in the vocabulary of fundamental physics, both fall under fundamental 'strict exceptionless laws'. These provide the deep story about the truthmakers for the judgement that your intention caused your hand to move as it did. (Typically, in cases of this kind, we would be happy to accept *that* there is a deep story.) There is no question of physical properties preempting mental properties. Remember, for Davidson, the mental–physical distinction is not a real distinction. An event is mental or physical 'only as described', so asking whether one event caused another *qua* physical or *qua* mental makes no sense. One event causes another (assuming events are causes!) owing to its nature, not owing to the way you might elect to describe it.

Davidson's 'anomalousness of the mental' is not a doctrine pertaining to mental properties and curious relations these bear to non-mental, physical properties. The 'anomalousness of the mental' stems from the evident fact that there is no systematic mapping between mental categories and categories of fundamental physics, something made much of by Fodor in his defence of the 'autonomy' of the special sciences (discussed below).

I regard Davidson as an ally, but my belief is that I could just have readily cited any of my early modern and Enlightenment heroes—Locke, Descartes, Priestley, Leibniz, Spinoza—and scores of philosophers at least since Plato. All would have rejected the idea, had it ever occurred to them, that ontology could be recovered directly from ways we represent the universe. It is one thing to know the truths, quite another matter to comprehend the truthmakers. The hierarchical picture implicit in the idea that, for every kind of truth, you need a distinct kind of truthmaker, remains attractive only so long as you imagine that different, often incommensurable, ways of describing the universe must, if they are legitimate, correspond to distinct orders

or levels of being. Give up that idea, and you are free to abandon the ungainly ontology that accompanies it.[5]

9.5 Fodor's Anti-Reductionism

Earlier, I suggested that non-reductive physicalism takes an important insight and redirects it ontologically. The result is an unstable, deeply implausible metaphysical doctrine. Proponents of non-reductive physicalism have saddled themselves with gratuitous metaphysical baggage not required by their anti-reductionist convictions. Worse, the unsavoury metaphysics threatens to undermine what *is* plausible about the anti-reductionist world view. Our best hope of making sense of the irreducibility of the special sciences requires jettisoning the ontology embraced by non-reductive physicalists and their opponents alike.

Take Jerry Fodor, who, together with Davidson, is often regarded as an archetypal non-reductive physicalist. Fodor's 'Special Sciences (or the Disunity of Science as a Working Hypothesis)' is widely cited as a definitive early statement of the doctrine. Fodor argues that reductionists begin with an acceptance of 'token physicalism', a perfectly reasonable starting point, but err in thinking that token physicalism implies 'type physicalism'.

Token physicalism endorses 'the generality of physics *vis-à-vis* the special sciences; roughly, the view that all events which fall under the laws of any science fall under the laws of physics' (Fodor 1974, 97). 'Token physicalism is weaker than what might be called 'type physicalism', the doctrine, roughly, that every property mentioned in the laws of any science is a physical property, where a physical property is one mentioned in the laws of physics (1974, 100). (Here Fodor is assuming that properties are reflections of predicates, so physical properties correspond to predicates occurring in statements of fundamental physical laws.)

Reductionists, Fodor contends, falsely suppose that token physicalism mandates a systematic mapping of predicates belonging to the special sciences onto predicates figuring in the laws of fundamental physics.

---

5 Galen Strawson has argued these points forcefully in a number of places. See Strawson 2006, 2008a, and papers collected in Strawson 2008b.

Reductivism is the conjunction of token physicalism with the assumption that there are natural kind predicates in an ideally completed physics which correspond to each natural kind predicate in an ideally completed special science.... [But] the truth of reductivism cannot be inferred from the assumption that token physicalism is true. (1974, 100)

'Reductivism' and a commitment to type physicalism do not follow from the acceptance of the generality of physics and token physicalism, however.

I ... want to argue that reductivism is too strong a constraint upon the unity of science, but that the relatively weaker doctrine will preserve the desired consequences of reductivism: token physicalism, the generality of physics, and its basic position among the sciences. (1974, 101)

My suggestion is that in these passages Fodor is most charitably read as upholding the same general line as Davidson, or rather Davidson rightly construed. True, Fodor goes on to argue for the possibility of a science of psychology that includes psychological laws, and Davidson is often interpreted as an opponent of such a possibility. My interest here, however, concerns only an ontological picture that is apparently consistent with both.

Fodor is standardly portrayed as holding that each of the special sciences concerns itself with an autonomous domain of kinds or causally significant properties. Some domains depend upon others, the whole edifice forming a hierarchy of levels of being. Fundamental physics concerns the basic level, but, although higher levels depend upon lower levels and ultimately on the basic physical level, the levels include irreducible properties and 'non-strict', *ceteris paribus* laws governing the behaviour of objects possessing those properties. Higher-level domains are dependent on lower-level domains but autonomous with respect to those domains.

Whatever Fodor himself might think, his anti-reductionist arguments provide little support for such a position. In fact, to the extent that he regards himself as a non-reductive physicalist, Fodor would have done himself a manifest disservice in taking on board any such picture. Fodor's explicit concerns are categories and taxonomies deployed in the special sciences. Special science categories, he argues, are not replaceable by categories belonging to the enterprise of fundamental physics. There is, for instance, no analysing or

translating psychological or biological predicates into predicates at home in physics.[6] This is clear in the passage from Fodor quoted at the outset of this chapter. The special sciences are autonomous and irreducible in the sense that taxonomies they deploy do not map smoothly onto taxonomies deployed by 'lower-level' sciences. These points extend beyond psychology to the special sciences generally. But all this is compatible with the idea that truthmakers for judgements couched in a vocabulary borrowed from the special sciences answer as well to judgements expressed in the vocabulary of fundamental physics.

One lesson to be learned here is that anti-reductionism need not, and in fact ought not, to be seen as a thesis concerning properties or families of properties, but as a claim about predicates, or categories, or taxonomies. The anti-reductionist denies that every truth about the universe, every 'lawlike generalization', is expressible in the vocabulary of fundamental physics. Suppose this is so. The question is, to what sort of ontology does this commit us? My answer is that it commits us to no particular ontology. It certainly does *not* commit us to a hierarchical ontology of states and properties, distinct levels of being. On the contrary, it is wholly consistent with the ontology advanced in earlier chapters of this book.

9.6 Boyd on Reduction and Natural Kinds

Richard Boyd, another prominent source of inspiration for non-reductive physicalists, defends an account of 'natural kinds' according to which kinds are 'homeostatic property clusters', 'contingently co-occurring families of properties' (see, for instance, Boyd 1980, 1999). As both Fodor and Boyd emphasize, sciences such as psychology and biology recognize kinds not recognized by chemistry or physics (and so for all the special sciences). This affords the special sciences with a measure of taxonomic and explanatory autonomy and tells against reductivist attempts to replace the special sciences with 'more basic' sciences. For purposes of this discussion, the important question is

---

6 Indeed, biological and psychological predicates differ in their extension, what they encompass when applied to the universe, from predicates at home in fundamental physics.

whether considerations of the sort invoked by Boyd yield an onto-logical, as distinct from a taxonomic, hierarchy.

Boyd begins by distinguishing 'natural' from purely 'nominal' kinds.

It is a philosophical truism that the philosophical theory of *natural* kinds is about how classificatory schemes come to contribute to the epistemic reliability of inductive and explanatory practices. . . . The theory of natural kinds is about how schemes of classification contribute to the formulation and identification of projectable hypotheses. . . . The naturalness of natural kinds consists in their aptness for induction and explanation; . . . definitions of natural kinds are reflections of the properties of their members that contribute to that aptness. (Boyd 1999, 410)

Predicates designating natural kinds are 'projectable', applicable to an open-ended range of new cases. Explanations featuring such predicates do more than merely summarize discoveries or extract instances from established generalizations. Natural kind terms enable us to say not only what instances of those kinds *do* but what they *would* do.

Natural kinds are to be distinguished from merely 'nominal kinds', kinds with 'purely conventional specifications of membership con-ditions' (407). The wholly conventional nature of nominal kinds stands in sharp contrast to 'a posteriori definitions of natural kinds . . . such as natural definitions of chemical kinds by molecular formulas (e.g. "water = H_2O")' (Boyd 1999, 407). In contrast, genuine natural kind terms designate Boyd's 'homeostatic property clusters,' families of properties that co-vary, exhibit stability in the face of change, and contribute to the capacities of their possessors thereby affecting what they do or would do. In the special sciences, in biology, for instance, such clusters resist precise characterization. Not every property asso-ciated with a given cluster need be present for an object to be an instance of the pertinent kind. Members of kinds exhibit family resemblances.

Property clusters—the kinds—are discoverable empirically, not by a priori analysis of terms used to express them. In this regard they differ from nominal kinds. 'Bachelor' expresses a purely nominal kind knowable simply by knowing the meaning or application conditions of 'bachelor'. Natural kind terms are meant to map 'the causal structure of the world' (409). Nominal kind terms incorporate merely

conventional designations. Nominal kinds are not genuine kinds at all, not kinds 'out there' doing causal work. The extension of a nominal kind term—'mascot of the 95th Rifles'—could be empty, the term could turn out to designate nothing at all. In contrast, a natural kind term—'water', 'gold', 'tiger'—is subject to 'accommo-dation', its significance adjusted in light of ongoing empirical investigation.

This is merely a rough outline of Boyd's intricate view, but it is sufficient for purposes of the present discussion. If you deflate Boyd's use of 'property', so that to say that something has the property of being salty, for instance, is just to say that 'is salty' applies truly to it, then it is patent that Boyd's conception of kinds, and the anti-reductionist line it supports are wholly compatible with the meta-physical picture I have been championing. The universe is as described in fundamental physics. But the universe, the very same universe, can also be described and 'taxonomized' in indefinitely many ways, ways incommensurable with what would be found in physics. Some of these ways latch onto important, but imperfect, causal or, more accurately, dispositional similarities. Systems of these similarities make up the subject matter of the special sciences.

The trick here is to distinguish broadly epistemological points concerning explanation and ontological theses taken to be mandated by those points. Thus Boyd's sharp distinction between kind terms that are 'pure constructs' (the pretenders) and those subject to 'accommodation' (terms expressing genuine natural kinds) could be doubted. You might wonder how a term deployed in substantive judgements about the universe could fail to be subject to 'accommo-dation', how it could be a 'pure construct'. In general, terms we invent to describe the universe are invented in response to our detection of important qualitative and dispositional similarities, and their use reflects adjustment in light of our continuous engagement with the universe. It is hard not to regard *all* kind terms as 'nominal' in this less artificial sense. We devise terms to reflect those divisions in the universe that exhibit what are, for us, striking, projectable, dis-positionally significant similarities.

Philosophers sometimes make a point of describing science as in the business of 'carving reality at the joints', a phrase traceable to Plato. But reality exhibits endless joints. The task of science is not

to find 'the' joints, but to circumscribe *significant* joints, joints that figure most prominently in our commerce with the universe. This is why scientific domains are typically irreducible. The sciences are distinguished by their distinctive taxonomies, by the joints these taxonomies commemorate. The limiting case is fundamental physics, which endeavours to locate joints that speak for themselves: joints without joints.

None of this calls for ontological drama. The thought that, corresponding to every distinct, but explanatorily serviceable taxonomy is a distinct level of reality is entirely gratuitous. It is not supported by anti-reductionist arguments, and it is not required to illuminate or explain scientific practice. Far less are the sciences themselves committed to an ontology of levels. You can fairly speak of levels of description, taxonomic levels, levels of explanation, levels of complexity and organization, but it would be a mistake of a fundamental sort to imagine that such talk requires us to posit levels of being.

9.7 Similarity in the Special Sciences

The suggestion on the table is that 'kind terms' in the special sciences pick out important qualitative and dispositional similarities. These similarities are genuine and perfectly objective. This conjures a picture of physics marking off the fundamental, fine-grained similarities while sciences at 'higher levels' circumscribe ever broader, more coarse-grained similarities. This way of putting it is misleading to the extent that it fails to reflect the extent to which categories in the special sciences can *cross*-categorize, the extent to which special science taxonomies can be *orthogonal* to one another and to taxonomies associated with physics and chemistry. If the sciences are hierarchical, the hierarchy is not one of neatly nested categories, but a messy hierarchy of the kind depicted in Dupré (1993; see also Cartwright 1999).

Such a picture can be refined by considering what exactly is involved in the identification of similarities. I have already discussed similarity in (too much) detail in chapter 5. There, similarity was described as a species of internal relation: if you have the relata

(as they are), you thereby have the relation. In this regard it is important to distinguish similarity among *objects*—substances or complexes you treat as substances—and similarity among properties or characteristics of complexes. In general, objects are similar because their characteristics are similar.[7] These two tomatoes, for instance, are similar by virtue of having similar shapes, colours, and sizes. But the shapes, colours, and sizes are similar, or not, *tout court*.

When you identify similarities, you engage in *abstraction*, Locke's *partial consideration*. You consider the aforementioned tomatoes' respective colours, for instance, ignoring their shapes or sizes, or spatial locations; or you consider their shapes, ignoring their colours or sizes. A capacity for abstraction is central to our capacity to think about the universe systematically. The special sciences are distinguished, in large measure, by their abstractions. Biologists, for instance, engage in important abstractions in identifying eyes, reproductive systems, and myriad other 'functional' characteristics of organisms. In so doing, biologists 'abstract away' from physical differences that would be blindingly salient from the point of view of physics, or chemistry, or, for that matter, molecular biology.

Much more could, and probably ought to, be said here, but my aim is only to note yet again that none of this requires a commitment to hierarchies of being. We have the fundamental thing or things, duly propertied and arranged. These provide truthmakers for all the truths of the special sciences and everyday life, all the truths requiring truthmakers. From this assemblage stems nature's endless 'joints', junctures of potential abstraction for the various sciences.[8]

9.8 Essence

Boyd distinguishes natural and nominal kinds. Other philosophers prefer to speak of *essences*: natural kind terms encompass the essences of entities to which they apply. You might think of an essence as, in Locke's words,

---

7 One important exception occurs in counting. Two collections might be similar solely with respect to number.

8 I am grateful to Laura Franklin-Hall for discussion on issues broached in this section.

the being of any thing, whereby it is what it is. And thus the real internal, but generally in Substances, unknown Constitution of Things, whereon their discoverable Qualities depend, may be called their *Essence*. (1690, III, iii, 15)

This, at any rate, is what Locke calls an object's *real essence*, distinguishing it from an object's *nominal essence*, 'that complex *Idea* the word *Gold* stands for, let it be, for instance, a Body yellow, of a certain weight, malleable, fusible, and fixed' (1690, III, vi, 2).

Real essences are characterized by what have been traditionally called 'real definitions'.[9] Nominal essences are captured by nominal definitions. One and the same term can have both a real and a nominal definition. Your grasp of a nominal definition of 'horse', for instance, might be a matter of your mastering the application conditions for the term, a matter of your grasping the colloquial meaning of 'horse'. In contrast, mastery of the real definition of 'horse' would require knowing what it is that makes a horse, a *horse*. To a first approximation, nominal definitions are sought in dictionaries; real definitions are the product of scientific investigation. As the sciences advance, dictionaries adapt, so this way of dividing the territory can become blurred and, eventually, obliterated.

What is the ontology of essences? If Boyd's natural kinds are meant to be essences, then Boyd's essences would be 'homeostatic property clusters'. A creature is a horse by virtue of possessing properties definitive of horsehood. Boyd's idea is not that there is a finite 'cluster' or collection of properties possession of which is necessary and sufficient for horsehood. Rather, every horse partakes of a sufficient number of these properties, where sufficiency is deliberately left vague so as to capture shifting, empirically sensitive classificatory practices of biologists. In the case of a chemical stuff, such as gold, matters are simpler. Gold is an element. Chemistry provides us with a real definition of gold in terms of the nature of gold atoms.[10]

Boyd thinks of essences as clusters of properties, but a more traditional conception takes essences to be explicable in terms of

9 Kit Fine, who advances an influential conception of essence, has also revived discussion of real definition; see Fine 1991, 1994.

10 Although it is common to regard gold as being more or less 'pure', 14-carat gold is not a species of gold: it is an alloy comprising gold—'24-carat gold'—and some 'base' metal.

'substantial forms', 'substantial universals'.[11] Just as substances differ from properties, so substantial universals are taken to differ from non-substantial, property universals. A tomato is red, spherical, and has a particular mass. These are properties of the tomato. If you think of properties as universals with instances, then this tomato's redness, its sphericality, and its mass are instances of universals. The tomato is red by virtue of 'instantiating'—exemplifying an instance of—the universal red. But the tomato *itself* is an instance of another kind of universal, a substantial universal, being a tomato, or tomatohood. This universal is itself 'characterized' by all the properties essential to tomatoes, or perhaps all the properties essential to this species of tomato. The property-universals are ways the substantial universal is, just as the tomato's redness and sphericality are ways the particular tomato is. A particular tomato's possession of these properties is explained by its being a tomato, by its being an instance of tomatohood.

The example is meant to be illustrative only. If you agree with me that substances must be simple, then substantial universals would have as instances only simples. Tomatoes are not simples, so tomatoes do not qualify as substances, and features of tomatoes do not qualify as properties. But what are the implications of this conception of substances and properties for the notion of essences?

First, 'natural kinds' of interest in the special sciences are a consequence of the fact that our universe is characterized by multitudinous similarities and regularities. These can be perfect—electrons are perfectly similar, or nearly so—or imperfect, as in the case of tigers or tomatoes. Common observable similarities stem from arrangements of substances, Locke's 'unknown constitution of things'. Every individual arrangement is what it is. Why some kinds of arrangement obtain and not others is determined by the natures of the substances making them up, their qualities and dispositions, and relations they bear to other substances. This tomato is red, spherical, and whatever else it is, not because it is a tomato; it is a tomato because it exhibits the right characteristics, characteristics definitive of tomatoes. Could

11 Armstrong's (1997, 67–8) conception of kinds is similar to Boyd's. For Campbell (1990, 31), the kinds are families of resembling tropes. Brian Ellis (2001, 2002) and E. J. Lowe (2006) defend conceptions of essences that invoke substantial universals.

this tomato have failed to be a tomato? Perhaps its constituents—these very constituents—could have been differently arranged so as to exhibit characteristics definitive of apples, or doorknobs, or large intestines. The possibilities are constrained only by the natures of the constituents and their interactive capacities.

This is, in effect, to regard kinds and essences as exclusively 'nominal', not in Locke's sense that they pertain exclusively to 'superficial' observable features of objects, but in the sense that objects *count* as apples or doorknobs solely in virtue of answering to categories we devise in the course of interacting with the universe.

> The measure and boundary of each Sort, or Species, whereby it is constituted that particular Sort, and distinguished from others, is what we call its Essence, which is nothing but that abstract Idea to which the Name is annexed. (Locke 1690, III, vi, 2)

Please do not imagine that this is anti-realism about tomatoes, or doorknobs, or, for that matter, scientific kinds. Being a tomato is a matter of having the right sorts of characteristic. Arrangements responsible for these characteristics are perfectly real, perfectly mind-independent, as are observable similarities among arrangements. If you want to know why these arrangements, rather than various others, occur in abundance, why they are reproducible and manipulable in various ways, a bad answer would be that these are all tomatoes. A better answer would appeal to the sciences that explain why various kinds of bond form among various kinds of particle, why some kinds of combination are stable, some are not, why some combinations flourish, others do not.

This is not to 'reduce' biology to chemistry or physics, only to suggest that the answer to the question why complex objects exhibit the 'discoverable qualities' they do is going to depend, as Locke puts it, on their 'unknown constitution', not on their possession of a particular guiding essence. Biological descriptions and explanations trade on important similarities among 'discoverable qualities' reflected in biological categories painstakingly developed in the course of observation and experiment. Often, even typically, these categories will be orthogonal to categories proffered in fundamental physics. Reduction is not an option, but neither is an unmotivated 'hyper-realism' about kinds. You can accommodate the characteristics appealed to by

biologists in explaining the nature and behaviour of biological organisms without imagining that terms used to describe the organisms and their characteristics designate 'higher-level' biological entities, or that the norms implicit in our classificatory practices must answer to essences that enjoy a life of their own.

9.9 A Universe without Essences

Readers sympathetic to the ontological picture advanced in earlier chapters should be no less sympathetic to these deflationary remarks concerning essences. I see Descartes, Locke, and their Enlightenment successors as disdaining essences on both ontological and scientific grounds. Essences associated with substantial universals are ontologically dodgy. You might be willing to live with a measure of dodgyness were essences to have important application in the sciences. But—as Descartes and Locke would argue—appeals to essences are not merely scientifically pointless, they are positively pernicious in encouraging us to imagine that the fundamental things—the electrons and quarks, for instance—might be 'guided' by higher-level forms.

Why does a particular tomato have the characteristics it has? Essences are meant to provide an answer: because it is a tomato. Its having the distinctive tomato nature, its being a tomato, is supposed to explain its possession of certain characteristic features and its developmental trajectory. But this seems the wrong way round. It is not that the tomato has these characteristics because it is a tomato; it is a tomato because it has these characteristics.

Why does an electron possess unit negative charge? Is its having unit negative charge explained by its being an electron? Or is it that it is an electron—it counts as an electron—in part because it possesses unit negative charge? This is not to say that you could not ask why electrons are as they are, why they have the mass they have, why they have unit negative charge. The answer, if there is one, is to be found in fundamental physics. It is no help to appeal to electrons' essences or natures.

Might essences have *causal* roles? Might essences, for instance, *cause* objects' observable, secondary qualities? Might the observable features

of a tomato stem in part from the tomato's essence? Might its redness, sphericality, juiciness, and the like be brought about by its nature, by its being a tomato? Were that the case, essences would be explanatorily significant and *prima facie* ontologically respectable.

I have already argued that tomatoes and their observable characteristics are particular dynamic interactive collections of particular kinds of fundamental thing duly organized and situated. If you take these particles and arrange them in this way in these circumstances, you thereby have a truthmaker for tomato judgements, you thereby have a tomato.[12] The idea that objects' observable characteristics could be caused by properties of their constituents, essential or otherwise, is misguided. The sense in which properties of, and relations among, an object's parts are responsible for the object's observable characteristics is not causal. Imagine arranging three matchsticks so as to form a triangle. The matchsticks and their relations are responsible for, but do not cause, the triangular shape.

This is by no means to exclude the importance of extensive causal interactions among the parts of complex objects. Observable characteristics of such objects often depend on such interactions, but not because the observable characteristics are caused by those interactions. The tomato's redness and sphericality are non-causal consequences of causal interactions among the tomato's constituents. You can create something red and spherical, something with the tomato's various characteristics, by assembling the right constituents and ensuring that they interact in the right ways. The tomato, with its familiar characteristics, *is* these constituents interacting in these ways.

The kind of anti-essentialism I am advocating is of a piece with kinds of argument common in the early modern period and in the Enlightenment. Take Descartes. Descartes appears to be rejecting essences in the sixth *Meditation* in a passage alluded to earlier (§ 6.6).

A clock constructed with wheels and weights observes all the laws of nature just as closely when it is badly made and tells the wrong time as when it completely fulfills the wishes of the clockmaker. In the same way I might consider the body of a man as a kind of machine equipped with and made

---

12 Time for the usual disclaimer: assuming that the simple substances are particles. If they are fields, or space–time itself, the One, the tomato is a way the fields, space–time, or the One is, the tomato is a mode.

up of bones, nerves, muscles, veins, blood, and skin in such a way that, even if there were no mind in it, it would still perform all the same movements as it does now in those cases where movement is not under the control of the will or, consequently, of the mind. I can easily see that if such a body suffers from dropsy, for example, and is affected by the dryness of the throat which normally produces in the mind the sensation of thirst, the resulting condition of the nerves and other parts will dispose the body to take a drink, with the result that the disease will be aggravated. Yet this is just as natural as the body's being stimulated by a similar dryness of the throat to take a drink when there is no such illness and the drink is beneficial. Admittedly, when I consider the purpose of the clock, I may say that it is departing from its nature when it does not tell the right time; and similarly when I consider the mechanism of the human body, I may think that, in relation to the movements which normally occur in it, it too is deviating from its nature if the throat is dry at a time when drinking is not beneficial to its continued health. But I am well aware that 'nature' as I have just used it has a very different significance from 'nature' in the other sense. As I have just used it, 'nature' is simply a label which depends on my thought; it is quite extraneous to the things to which it is applied, and depends simply upon my comparison between the idea of a sick man and a badly made clock, and the idea of a healthy man and a well-made clock (Descartes 1641, 58–9)

Clocks and natural mechanisms do what they do owing to the character and organization of their constituents and relations these bear to goings on external to them. The idea that they are informed or guided by an essence or nature is the result of projecting our own perspectives and interests onto the universe. A particular clock or a particular human body could be said to have a nature, but this is just its natural make-up. It would be hard to improve upon Locke on this point.

I have often mentioned *real Essence*, distinct in Substances, from those *abstract Ideas* of them, which I call their *nominal Essence*. By this real Essence, I mean, the real constitution of any Thing, which is the foundation of all those Properties, that are combined in and constantly found to co-exist with the *nominal Essence*; that particular constitution which every Thing has within it self, without any relation to anything without it. But *Essence*, even in this sense, *relates to a Sort*, and supposes a *Species*: For being that real Constitution, on which the Properties depend, it necessarily presupposes a sort of Things, Properties belonging only to *Species*, and not to individuals; v.g. Supposing the nominal Essence of *Gold*, to be Body of such a peculiar

Colour and Weight, with Malleability and Fusibility, the real Essence is that Constitution of the parts of Matter on which these Qualities, and their Union, depend; and is also the foundation of its Solubility in *Aqua Regia*, and other Properties accompanying that complex *Idea*. Here are *Essences* and *Properties*, but all upon supposition of a Sort, or general abstract *Idea*, which is considered as immutable; but there is no individual parcel of Matter, to which any of these Qualities are so annexed, as to be *essential* to it, or inseparable from it. That which is essential, belongs to it as a Condition, whereby it is of this or that Sort: But take away the consideration of its being ranked under the name of some abstract Idea, and then there is nothing necessary to it, nothing inseparable from it. Indeed, as to the *real Essences* of Substances, we only suppose their Being, without precisely knowing what they are: But that which annexes them still to the *Species*, is the nominal Essence, of which they are the supposed foundation and cause.[13]

Everything is what it is, and in that sense, everything has a nature, a real essence. The mistake is to elevate this truism to the status of an ontological thesis concerning a special category of entity: essences. To do so is to saddle nature with a normative element yielding what D. C. Williams decried as 'the metaphysics of the dog show', to read our own interests into reality (Williams 1954; see § 5.10 above).

9.10   Historical Caveat

This chapter began with a discussion of anti-reductionist arguments and now concludes with a tentative discussion of essences. The running theme has been that both anti-reductionist and essentialist sentiments are grounded in considerations concerning how we categorize or taxonomize our universe. Confusion arises when these sentiments are read back into the universe and issue in unwarranted metaphysical pronouncements.

The mistake is to imagine that we are faced with a choice between accepting real essences, real divisions in the universe as a basis for our taxonomies, and 'merely conventional', 'projectivist', anti-realist conceptions of those taxonomies. To the extent that our taxonomies

are responsive to worldly vicissitudes, they reflect real divisions, real 'joints', they reflect 'real natures'. The choice is not between realism and objectivity in the sciences, on the one hand, and, on the other hand, the idea that the universe as described by science is a mere projection or construct. The choice is between ontologies that make sense of realism and objectivity in the sciences and, in addition, exhibit a suitable degree of internal integrity.

The kinds of mistake alluded to here, are, I consider, both characteristic of and unique to twentieth-century approaches to metaphysics. As soon as I say that, someone will produce examples of the kind of reasoning I regard as characteristic of twentieth-century analytic philosophy going back at least to the Presocratics. My aim throughout this endeavour is not primarily historical, however. If I am wrong, if the trends I decry are to be found throughout the history of philosophy, I would be saddened, even devastated, but philosophically unmoved. It is at least true that many significant figures—the likes of Locke and Descartes—are on my side. This does not provide independent support for my argument, but it does at least suggest that the position advanced here is not entirely unprecedented, not entirely crazy, not even *prima facie* crazy.

---

13 Locke 1690, III, vi, 6. Jonathan Bennett called my attention to this passage. For commentary, see Bennett 2001, § 197. Note that in the quoted passage Locke is using 'property' in the Episcopalian sense (§ 8.1 above).

CHAPTER 10

Mind and Mentality

It is the strangest anomaly of recent science that while an influential number of physicists, once supposed to be students of physical nature, are suggesting that only conscious experience exists, an equally influential number of psychologists, once supposed to be students of consciousness, have suggested that only physical nature exists. Either of these two contrary propositions seems paradoxical to the non-professional; but perhaps the greatest impatience was provoked by the psychologist's, both on the side of the behaviorists who defended it and on the side of the non-behaviorists who repudiated it.

(D. C. Williams 1934, 461/23)

I am rather inclined to think that, though the subject is beyond our comprehension at present, man does not consist of two principles, so essentially different from one another as *matter* and *spirit*, which are always described as having not one common property, by means of which they can affect or act upon each other; the one occupying space, and the other not only not occupying the least imaginable portion of space, but incapable of bearing relation to it; insomuch that, properly speaking, my mind is no more *in my body* than it is in the moon. I rather think that the whole man is of some *uniform composition*, and that the property of *perception*, as well as the other powers that are termed *mental*, is the result (whether necessary or not) of such an original structure as that of the brain. Consequently that the whole man becomes extinct at death, and that we have no hope of surviving the grave but what is derived from the scheme of revelation.

(Priestley 1775, xx)

10.1 Philosophical Puzzlement

Distinctively philosophical puzzles induce an atmosphere of hopelessness. Indeed any puzzle that steadfastly resists solution by ordinary or scientific means is on its way to becoming a philosophical puzzle—or, worse, a *merely* philosophical puzzle. Puzzles resist solution for many reasons. Three kinds of case come to mind.

First, there are puzzles arising from questions that seem answerable in principle but only after an investment of time and resources difficult or impossible to justify. Such questions are ones we know how to set about answering but lack the will or the funding to do so. In physics, for instance, questions about the deep nature of matter might require equipment too complex or expensive to build or operate. In the biological sciences, answers to some questions might require experimental procedures that would produce unacceptable levels of distress on the part of sentient creatures.

Second, there are puzzles we have only the foggiest idea how to go about resolving. How can the quantum theory and general relativity be reconciled? A solution to this puzzle does not await a refinement of instruments or faster, more powerful computing machines. It awaits the development of a unified theoretical apparatus concerning which there is at present little consensus.

Third, some puzzles resist resolution because they incorporate questions that arise in a way that undermines the prospects for any satisfactory answer. Most often this happens when a question originates within a conceptual framework that is itself the source of difficulty. When this is so, progress is sometimes possible when the framework, which initially seemed inevitable, is revealed to be optional. Recall Kant's treatment of the 'Antinomies of Pure Reason'. The antinomies arise, according to Kant, when we find ourselves obliged to choose between a pair of competing hypotheses, both of which stem from philosophical doctrines that could be doubted. As Kant sees it, the resistance of the antinomies to solution *itself* suggests that they stem from a mistake of this kind.

Wittgenstein regarded *every* philosophical question as a variant of this third category. Once you see through confusions underlying the formulation of a philosophical question, the question evaporates.

The fly escapes the fly-bottle, not by exercising cleverness or insight, but simply by reversing field, flying out of the bottle's open end.

Wittgenstein preached that philosophical riddles are rooted in linguistic befuddlement. We deploy our language unself-consciously and without undue puzzlement. But the very features of language that make it most adaptable can lead us astray when we move from settings in which we *use* language as it was meant to be used, to those in which we press beyond its intended application. Once we embark on philosophical reflection, forms of speech can render particular theses irresistible. The origin of these theses in our modes of judgement can make them especially difficult to single out for examination. Unexamined, they colour our thoughts about the universe without betraying their hand.

Some readers will already have entertained the suspicion that my regarding a substance–property ontology as inevitable is likely to be little more than a provincial reflection of my having inculcated a language organized around subject–predicate modes of expression. I admit this could be true. I prefer to think, however, that the subject–predicate form reflects a fundamental division in reality, and I stand by independent considerations offered in support of the correlativity of substances and properties.

Whether you agree with me on this, I hope you will at least agree that there is something to the idea that philosophical questions are, often enough, products of seemingly innocent background assumptions, the credentials of which might easily be questioned. You could think this without supposing, as Wittgenstein did, that it is all just linguistic confusion.

10.2 Philosophical Fashion

Failure to address what Collingwood (1940) called 'absolute presuppositions' is abetted by the increasing 'professionalization' of philosophy. Apprentice philosophers are rewarded for operating within prevailing frameworks. Reflection on the frameworks is regarded as, at best, optional. If your goal is to publish papers in the right journals, you will want to make eye-catching moves without disturbing the water more than absolutely necessary. Young scientists learn their

craft by working in the laboratories of senior scientists, imbibing aims and attitudes manifested by their seasoned colleagues. Young philosophers hone their skills by filling in the blanks. The cleverest, the most successful, are rewarded with appearances in the right journals and invitations to publish in edited collections aimed at specialist readers with shared tastes. We have become a culture of medieval monks copying and illuminating manuscripts in artful ways. Professional success is one thing, however, philosophical significance is another matter, altogether more challenging, more elusive. It has always been so; philosophical rock stars are abundant, genuine philosophers scarce.

If progress in philosophy is measured in the way progress towards tenure is measured, then philosophy has made enormous strides over the past fifty years. Journals have proliferated, and PhD programs turn out ever more polished products happy to contribute to backlogs of eminently publishable papers in those journals. But if progress is measured by the extent to which we philosophers have managed to excavate the pictures that drive our intuitions, matters are less encouraging.

I realize that going on in this way risks alienating readers who otherwise might be willing to give me the benefit of the doubt. I do not, however, regard myself as above the fray, a lonely champion of truth, a prophet crying out in the wilderness. I mean only to suggest that we might move ahead if only we relaxed a little, if only we were in less of a hurry, if only we were occasionally more open to alternative approaches to outstanding problems, approaches possibly out of line with prevailing fashion. We have the models: our eminent historical predecessors. The trick is to read them in something like their own terms, rather than through the lens of linguisticized metaphysics. The point is not that our predecessors, Descartes and Locke, for instance, had a pipeline to truth. The point, rather, is that living with them for a period throws our own preconceptions into relief. When that happens, we are in a position to see ourselves for what we are, reflected in the mirror of history.

So much for social narrative. In the sections that follow, and in the next chapter, I examine proposals aimed at opening up current discussions of consciousness and intelligent thought. I am less interested in defending a positive account than in convincing you, the

reader, that assumptions responsible for present-day perplexities over the place of the mind in the 'natural world' are at least *optional*: they did not fall from the sky, they are not forced on us by reality itself. Once you see that, a universe of otherwise invisible possibilities opens up before you. Hard philosophical problems are made harder by allegiances to doctrines that retain their force only so long as they are left unexamined.

10.3 The 'Mental' and the 'Physical'

The title of this section is borrowed from Herbert Feigl's 1958 essay defending an 'identity theory' according to which states of mind are to be identified with physical states, more particularly, states of the nervous system. Much of Feigl's long essay is concerned with spelling out philosophical presuppositions he regards as longstanding impediments to the identification of the mental and the physical, barriers to our recognition that whatever is designated by a mental term is, or could be, designated by some physical term.

Inspired by Feigl, let me mention two widely accepted background assumptions operative in much of what today constitutes mainstream philosophy of mind.

(1) Realism about the mental requires a commitment to distinctively mental states and properties.

(2) Mental states and properties are dependent on, but distinct from, physical states and properties.

These assumptions could be questioned. Indeed, I have tried to supply all the reasons you would need to reject both. 'Realism about tomatoes' requires, not that there be substances—tomatoes—possessing properties essential to, or definitive of, tomatoes, but that judgements about tomatoes are, often enough, true. What, in particular cases, answers to those judgements, the nature of their truthmakers, is at the present time anyone's guess. Similarly, realism about the mental does not require that there be mental substances or mental properties, but only that you accept judgements about mental goings-on as true and as being made true by ways the universe is. It could be that truthmakers for judgements about the mental include mental

substances, or at any rate substances possessing mental properties, but, not only is this extremely unlikely, it is a central source of confusion when you pose questions about the relation minds bear to bodies.

Suppose you followed Spinoza and Davidson in regarding the mental–physical distinction as one of conception only, not a real distinction.[1] There is no question of 'reducing' one conception to the other. There is no question of giving an analysis of mental terms using a purely physical vocabulary, a vocabulary borrowed from fundamental physics, for instance, or from biology or neuroscience. Reduction is not in the cards. But this important fact is consistent with the possibility that truthmakers for judgements framed in a mental vocabulary are invariably 'physically describable', by which I mean describable in terms at home in fundamental physics. The mental and the physical are names, not of families of substances and properties, but of ways we have of conceiving, describing, and explaining the universe.

You might be suspicious of my appealing to mental and physical categories without first making clear what distinguishes the one from the other. Although this is a problem for anyone who thinks that the mental–physical distinction is ontologically deep, it is not my problem. Characterize mental terms in whatever way you like. Ask the experts, consult psychologists and philosophers of mind. Make a list. My claim is that, however you elect to characterize them, whatever you regard as falling under the labels 'mental' and 'physical', judgements in which those terms figure, when true, are made true by ways the universe is that could be given a description that makes use of categories anchored in fundamental physics.

Although philosophers like to pretend that the mental forms a sharp-edged category of reality, it is notoriously difficult to find a characterization that encompasses all and only those things commonly classified as mental. Such a characterization would have to include conscious and unconscious reasons, thoughts, beliefs, desires, images, perceivings, passions, feelings, and sensations.

---

1 The denial of a real distinction between the mental and the physical is forcefully defended by Galen Strawson (see, for instance, Strawson 2006). I am in agreement with many of Strawson's conclusions, although not always with his formulation of them.

Most often philosophers are content to regard what counts as mental as intuitively obvious. We have our own personal, subjective experiences, and we have psychology, a science explicitly devoted to the study of mental states and processes. Psychological kinds have the feel of perfectly respectable natural kinds on all fours with biological or chemical kinds. Still, a moment's reflection on all that we ordinarily classify as mental and on all that is included in the subject matter of psychology supports another, very different, impression. What we classify as mental is a hodgepodge of states, conditions, and processes, a motley collection lacking inner unity. Compare your belief that the Earth orbits the Sun, an experience of vertigo, your remembering the words to *Casabianca*, your hearing chalk squeaking against a blackboard. Do these fall under an obvious kind? Maybe this impression is confused, however. Maybe apparent differences among items in the list are superficial. Maybe I could sell you the Brooklyn Bridge.

10.4 Property Dualism

One kind of response to the apparent lack of unity among properties and states commonly regarded as mental is to concede the point and divide the mental into a pair of rough, possibly open-ended categories. On the one hand, there are sensory states, E. C. Tolman's 'raw feels': conscious episodes, including perceptual experiences and bodily sensations (Tolman 1932, 426–7). On the other hand, there are the 'propositional attitudes': beliefs, desires, intentions, and the like. The latter might be thought to be susceptible to a broadly functionalist analysis. The former, however, resist this kind of reduction. Pains, for instance, seem to have, in addition to a functional profile (pains result from bodily disturbances and move sufferers to take ameliorative action), a distinctive *qualitative* nature, a nature that makes them especially apt for playing the 'pain role'. You might be tempted to ascribe beliefs and intentions to a cleverly programmed computing machine, but pains seem not to be the kinds of state that could be added solely by inspired programming.

David Chalmers has something like this division in mind when he recommends distinguishing *conscious* states from what he calls states of *awareness* (Chalmers 1996). Both kinds of state, Chalmers holds, are

kinds of functional state operative in sentient beings. A state of awareness is a state that could be explained in *exclusively* functional terms: states of awareness are a species of functional state, nothing more nor less, the kind of state studied by psychologists. Some functional states, however, in addition to exhibiting unproblematic functional characteristics, are *conscious*. Your experience, your fully conscious experience, of a brightly illuminated ripe tomato has *both* a particular functional profile *and* a distinctive qualitative character. This qualitative character is not 'functionally reducible'. It is something possessed by the state *in addition to* its manifestly functional features. Chalmers argues that it is a contingent fact about our universe that creatures with the right kinds of functional make-up are conscious.

Minds of sentient creatures, then, are functional systems, operating in accord with biologically tractable functional principles. But this is not the whole story. Some functional states give rise to consciousness. Conscious states are salient to those undergoing them, invisible to outsiders. Because the distinctive qualitative character of conscious experiences 'arises from' functional organization (Chalmers 1996, 41, 47), science can rest content with purely functional descriptions and explanations of the behaviour of sentient creatures.

According to Chalmers, the distinctive non-physical character of consciousness combined with the plausibility of functionalism as a psychological theory can be explained by positing new fundamental laws of nature. Although consciousness is rooted in the functional organization of physical systems, its occurrence is not explicable by reference to laws governing purely physical phenomena. To accommodate consciousness, science requires additional fundamental laws. The fundamental laws of nature, then, must include both laws governing the behaviour of the particles and fields and *psychophysical* laws governing the production of conscious qualities in systems with the right kinds of functional organization.

Chalmers's psychophysical laws are contingent and irreducible. Nothing in the fundamental physical laws, laws governing the behaviour of the fundamental physical things, could explain psychophysical laws. If you thought of fundamental natural laws on the model of an axiom system, Chalmers's psychophysical laws would be independent axioms, independent in the sense that they are not derivable as

theorems from axioms pertaining to the purely physical realm. If universes could differ with respect to the fundamental laws, then there could be universes exactly resembling ours with respect to all their fundamental *physical* laws and the distribution of particles in space–time, but ones in which consciousness is entirely absent. This is Chalmers's famous invocation of 'zombies', imagined creatures like us physically, 'molecule-for-molecule duplicates', but completely lacking in conscious experiences: 'all is dark inside' (Chalmers 1996, 96).

A word about zombies. Robert Kirk was the first to unleash philosophical zombies on a population of unsuspecting philosophers (Kirk 1974). Kirk notes that Descartes comes close to the zombie idea in arguing that the human body could be regarded as a machine that operates on purely physical, mechanical, principles (2005, 7; see § 9.9 above). Descartes accepts that God (or an army of pious scientists) might construct an automaton the behaviour of which would be mostly indistinguishable from the behaviour of an intelligent human being. Such a robot would fall short of zombiehood, however, because it would be unable to use language intelligently or behave in ways requiring rational deliberation. Such activities, Descartes thought, involve the deployment of a mind operating on non-mechanical—that is, non-physical—principles (see Descartes 1637, 1664). Unlike most philosophers today, Descartes did not regard the physical universe as 'causally closed'.

As noted already, Chalmers's guiding idea is that, when you consider any reasonable list of mental terms, you will find some that designate purely functional states of sentient creatures and some that designate states, functional states, that ground conscious qualities. Functional states are physically unproblematic. You can see, at least in principle, how God, or even a team of dedicated scientists, might go about creating a being with a particular functional profile using exclusively physical components. In fact, it is a functionalist article of faith that there are always going to be many different ways to do this, a point Chalmers emphasizes. Whereas functional states are physically transparent, conscious qualities—the *qualia*—are, in contrast, physically enigmatic, 'phenomenal' add-ons exhibiting markedly non-physical natures. Psychophysical laws prevailing in our universe

ensure that any creature with the right kind of functional organization will, in addition, be conscious.

On such a conception, our conviction that the qualitative character of our conscious experiences makes a difference in our behaviour—that the painfulness of a pain is what leads us to complain and seek relief, for instance—stems from the fact that qualitative character 'arises from' functional states that are themselves causally engaged. Our recognition that others resemble us functionally licenses our belief that they are like us in their conscious experiences. But, as the zombie possibility shows, this belief is defeasible.

10.5 Metaphysical Infelicities[2]

In enticing us down the rabbit hole, Chalmers paints a picture of a kind of naturalistic dualism: a dualism of states or properties meant to coexist peacefully with the precepts of fundamental physics. Conscious states are physically unremarkable functional states accompanied by non-physical 'phenomenal properties'. Qualities of conscious experiences 'arise from' these purely physical states and processes. In psychology, the science, functionalism rules: psychological states *are* functional states of sentient beings. In our universe such states happen to give rise to *qualia*, familiar qualities of conscious experiences. These qualities never get in the way of physical processes, however, and for that reason psychologists and neuroscientists, *qua* psychologists and neuroscientists, can safely ignore them. Still, because they have respectable functional anchors, there is no harm in referring to, even revelling in, conscious qualities in describing human experience.

A central tenet of Chalmers's thesis is that laws governing the production of *qualia* are fundamental in the sense that they are not derivable from, or entailed by, fundamental laws governing the fundamental physical constituents of the universe. A universe indiscernible from ours in every detailed *physical* respect except that it lacked these fundamental psychophysical laws would be a universe indiscernible from ours psychologically. Such a universe would differ from ours in one dramatic respect: its inhabitants would be utterly

2 A version of the argument developed in this section appeared originally in Heil 2003a, chap. 19.

lacking in consciousness. In such a universe there might be a zombie you, a being resembling you down to the last detail, a being functionally, hence behaviourally and psychologically, indiscernible from you, but utterly lacking in conscious experiences. (Come to think of it, what makes you so sure that *you* are not zombie you?)

For Chalmers, it is important that the fundamental laws responsible for the production of conscious qualities relate states *functionally* characterized to particular kinds of conscious quality. Imagine suffering a pain with a distinctive qualitative character. Your experiencing this pain stems from your being in a particular sort of functional state. Any being, or any being in *our* universe, in a functionally equivalent state would experience a pain with precisely *this* qualitative character. Psychophysical laws tie conscious qualities, the *qualia*, not to particular kinds of physical or biological state, but to particular kinds of functional state.[3] This pivotal feature of Chalmers's view has disquieting ontological implications.

You can see what I have in mind by starting with the idea that functional states are 'multiply realizable': creatures quite different physically could be in the very same functional state. On Chalmers's model, when you are in a particular conscious state, this is because your brain is in a state that plays the right sort of functional role. Suppose, for instance, you are savouring the smell of freshly baked bread, and suppose your savouring is a matter of your being in a particular kind of functional state that is now realized by a particular neurological state of your brain. If scientists could replace neurons in your brain responsible for your being in that state with transistors, or an array of gears and springs, or a legion of undergraduates equipped with mobile telephones linked to your brain wirelessly, you would be in the very same functional state.

A standard argument *against* functionalism begins with these kinds of wild possibility and concludes that functionalism cannot account for qualities of conscious experiences. It has seemed obvious to critics that beings functionally alike could fail to be alike consciously.

3 You could contrast Chalmers's view to John Searle's (see Searle 1983, esp. chap. 10; 1992, chaps. 4, 5; 1997; 2004, chaps. 4, 5). Searle traces consciousness to unspecified *biological* features of sentient creatures.

Chalmers agrees in part. He agrees that it is a merely contingent fact that functionally equivalent creatures are thereby alike with respect to their conscious experiences. In addition to a creature's functional make-up, you need contingent psychophysical laws of nature. These serve to guarantee that functional duplicates will be conscious duplicates. You can appreciate the force of these laws by recognizing that, if neuroscientists could cycle through the realizers of your savouring the smell of freshly baked bread (from neurons, to transistors, to gears and springs, to armies of undergraduates, and back again), you would notice no difference in your conscious experience (Chalmers 1996, chap. 7). You would notice no difference because there would *be* no difference.

Note that fundamental laws governing the occurrence of conscious experiences would have to relate particular conscious qualities to heterogeneous physical conditions, conditions that would realize particular functional states, hence conditions belonging to very different physical categories. Such laws would apparently resemble no other fundamental laws. Fundamental laws concern the nature and behaviour of the fundamental things. Functional states, conceived of as multiply realizable, are scarcely fundamental, however. So the laws in question would need to connect heterogeneous families of fundamental realizers of particular kinds of functional state to particular kinds of conscious quality.

You might not regard this as a serious problem. You might not regard it as a serious problem if, for instance, you thought of functional states or properties as 'higher-level' states or properties, states or properties dependent on, but distinct from, their 'lower-level' realizers. These higher-level states anchor the fundamental psychophysical laws.

I have tried to persuade you that hierarchical ontologies that posit higher-level properties and states are hopeless. Even if you have doubts about my arguments, however, it is worth noting that Chalmers's own official view concerning functional states is, in fact, apparently compatible with the one defended here. The issues are interesting in their own right, but they are interesting as well for the light they throw on ontological projects. In ontology, it is unwise to pick and choose theses because they provide the most direct route

to a desired conclusion. In ontology, one thing leads to another. Ontology is a package deal.

Chalmers frames his discussion in terms of supervenience. On the one hand, there are the fundamental physical facts, the '*A*-facts', facts concerning the fundamental particles, perhaps, their properties and their space–time trajectories. On the other hand, there are all the other physical facts, the '*B*-facts'. These would include facts about entities made up of the fundamental things: trees, planets, and the brains of sentient creatures, for instance. According to Chalmers, the *B*-facts 'logically supervene' on the *A*-facts. Think of logical supervenience as the 'nothing-over-and-above' relation: 'when logical supervenience holds, *all there is* to the *B*-facts being as they are is that the *A*-facts are as they are' (1996, 30). Judgements concerning *B*-facts are made true by the *A*-facts.

Logical supervenience is to be contrasted with what Chalmers calls 'natural supervenience'. Certain *B*-facts could 'naturally supervene' on the *A*-facts without logically supervening on those *A*-facts. This would be so when the *B*-facts depend in part on the holding of particular contingent laws of nature not included among the *A*-facts. One of Chalmers's examples is the relation between pressure of a gas and its temperature and volume. This relation depends on certain constants that, Chalmers is confident, could have been other than they are. Chalmers accepts that the fundamental laws are contingent; laws applicable in our universe might have been different. *Given* the laws governing the fundamental physical things, however, the *B*-facts logically supervene on the *A*-facts. You need *additional* laws, supplemental *psychophysical* laws, to accommodate qualities of conscious experiences.

Natural supervenience, then, unlike logical supervenience, is a kind of lawful correlation. The terminology in play is unfortunate to the extent that it encourages the unwary to regard natural and logical supervenience as species of a common genus. Lawful correlations are nothing at all like nothing-over-and-above relations, however.

Before discussing the bearing of all this on Chalmers's conception of the mental–physical relation, a reminder: if you are sympathetic to the idea that objects' powers are built into their properties, then the scope of logical supervenience is broadened considerably. Universes

alike with respect to kinds of fundamental substance will be governed by the same laws of nature. There is no holding the fundamental things and their properties constant and varying the laws.

Whatever your preferences, you should recognize that Chalmers helps himself to a substantive metaphysical thesis concerning natural necessities in making his case. Here, as elsewhere, the heavy lifting is done by background assumptions presumed to be wholly innocent, or, if not exactly innocent, at least uncontroversial.

Having introduced natural supervenience as a kind of lawful, but contingent, correlation, Chalmers, as noted, argues that consciousness, or qualities of conscious experiences, 'arise from' or naturally supervene on functional states of sentient creatures. Creatures indiscernible functionally will be indiscernible mentally. This is not because mental properties—the *qualia*—are functional properties. No, mental properties and states naturally supervene on functional properties and states. Your current conscious state 'arises from' your current functional state. This kind of functionalism supplemented by the postulation of fundamental psychophysical laws provides an answer to the standard complaint that functionalism leaves out what is important about consciousness.

A closer look at the metaphysical details, however, yields new, and in many ways, much more unsettling, worries. Suppose, with Chalmers, that your functional state *logically* supervenes on a particular collection of particles, those making up your nervous system, organized and interacting in a particular way. Functional facts are among the *B*-facts, and these logically supervene in the *A*-facts. Thus, your now being in this functional state, *is* this collection of particles interacting in just this way.[4] So your conscious state 'arises from' this dynamic, interacting collection. The law governing this arising-from relation is a fundamental law. It will have to be a law relating an unimaginably complex dynamic arrangement of fundamental particles to the production of familiar, apparently simple introspectable qualities of your conscious experience.

---

4 The particles interact with one another and with particles outside the collection in all sorts of important ways. Histories of the particles and relations they bear to hosts of other particles might be included in the 'supervenience base'. I omit these vital complications for the sake of simplicity.

But such laws would be utterly unlike other fundamental laws. J. J. C. Smart made the point more than fifty years ago.

It is sometimes asked, 'Why can't there be psycho-physical laws which are of a novel sort, just as the laws of electricity and magnetism were novelties from the standpoint of Newtonian mechanics?' Certainly we are pretty sure in the future to come across ultimate laws of a novel type, but I expect them to relate simple constituents: for example whatever ultimate particles are then in vogue. I cannot believe that the ultimate laws of nature could relate simple constituents to configurations consisting of billions of neurons (and goodness knows how many billions of billions of ultimate particles) all put together for all the world as though their main purpose was to be a negative feedback mechanism of a complicated sort. Such ultimate laws would be like nothing so far known in science. (Smart 1959, 142–3)

The situation envisioned by Chalmers is in fact *much worse* than this. Recall that a central tenet of Chalmers's theory of consciousness is that qualities of conscious experiences naturally supervene on *functional* states of sentient creatures. Two creatures very different in their physical make-up could nevertheless be functionally indiscernible. If you put this thesis together with the idea that functional states logically supervene on physical states—so that every particular functional state just *is* a particular physical arrangement—it follows that the fundamental laws relating physical conditions of sentient creatures to qualities of conscious experiences are going to be many–one laws. Every distinct kind of conscious state will be necessitated by many very different kinds of highly complex fundamental physical state. This reinforces Smart's point in spades.

   You can see the problem by working through an imaginary case. Suppose you, a robot, and an octopus are now experiencing a particular species of dull ache. Functionalism requires that you, the robot, and the octopus all be in the same functional state. This state is *multiply realized*, however. You, the robot, and the octopus are (we are supposing) very different physically, and, in particular, very different with respect to the physical condition in virtue of which each of you is in this functional state. This can be represented schematically by means of a simple diagram. Let C be a particular kind of conscious experience, an experience of a particular kind of dull ache, let F be the functional state responsible for C, the state on

which C naturally supervenes. Finally, let P_1, P_2, and P_3 be the physical states of you, the robot, and the octopus on which F logically supervenes, physical states that 'play the pain role'. We have something like the following picture:
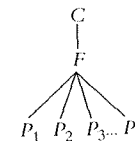
Figure 10.1

This makes it appear as though Chalmers's fundamental law governs the 'arising' of C from F. F, however, logically supervenes on P_1, P_2, and P_3 (and perhaps an open-ended collection of other actual and possible physical states, hence the '$\ldots P_n$'): F is really nothing over and above P_1, P_2, $P_3 \ldots P_n$. As Chalmers puts it in a passage quoted earlier, '*all there is* to the B-facts being as they are is that the A-facts are as they are'. Functional facts are paradigmatic B-facts; the A-facts are the fundamental physical facts. So a better, more candid, or at least less misleading, representation of the natural supervenience of C might be
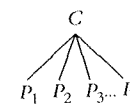
Figure 10.2

This makes it clear that the combination of

(1) conscious experiences or their qualities 'naturally supervene' on functional states of sentient beings;

(2) functional states and properties are multiply realizable; and

(3) functional states and properties logically supervene on heterogeneous fundamental physical states and properties;

implies that

(4) conscious experiences or their qualities 'naturally supervene' on heterogeneous fundamental physical states and properties.

The upshot is that Chalmers is apparently committed to fundamental laws that connect particular kinds of conscious state or property with complex *and heterogeneous* physical states in an apparently unprecedented way. The utter strangeness of this commitment is masked by the positing of a functional intermediary hovering between conscious states or properties—the *qualia*—and the physical realizers of these. Once this intermediary is recognized for what it is, a classificatory device, not an entity itself 'over and above' interactive arrangements of the fundamental things, it loses its ontological standing—and certainly its standing as a constituent of a fundamental physical relation falling under a fundamental physical law.

Stated in such a way as to make its ontological commitments fully explicit is enough to cast Chalmers's theory into doubt. The theory obliges us to embrace fundamental laws of a most seriously peculiar sort, laws that appear simultaneously *sui generis* and ad hoc: an ontological embarrassment.

10.6 Back to the Drawing Board

Problems bubbling to the surface here are symptomatic of a more general problem inherent in applications of metaphysics to immediately pressing philosophical concerns. It is all too easy, too tempting, to cherry pick, to bend metaphysics to the service of parochial philosophical interests. The upshot is, more often than not, a metaphysical patchwork with little or no internal cohesion or independent plausibility.

One option, an option motivating this volume, would be to recognize applications of metaphysics for what they are: *applied metaphysics*. Take the philosophy of mind, and, in particular, non-reductive physicalist approaches to the philosophy of mind. A point pressed at the outset of this chapter is that these, more often than not, lead off with a significant core of substantive metaphysical assumptions, a Big Picture. The assumptions collide with our experiences and with scientific practice in various ways. The problem of mental causation provides just one example of what I have in mind, the problem of conscious qualities, another. These problems arise in a

particular way owing to the character of the Big Picture embraced by non-reductive physicalists.

You could tinker with the concept of causation and various conceptions of property 'realization' so as to achieve the desired outcome—respectable mental–physical causal interaction among objects possessing utterly distinct kinds of property—*or* you could look more closely at the prevailing Big Picture in the belief—or hope—that once you get the metaphysics right, once you get the *ontology* right, problems besetting the philosophy of mind will take care of themselves.

The hallmark of a Big Picture is the seeming inevitability of its commitments. The mission of philosophy, however, is to put everything up for examination, not all at once, but all in good time. To the extent that a philosophical puzzle depends on what turn out to be optional commitments, you will want to be in a position to evaluate the appeal of those commitments, then let the chips fall where they might. The chapter that follows comprises an attempt to move ahead on one front, to move ahead on the question of conscious qualities, the *qualia*, without the kinds of commitment I have been calling into question.

CHAPTER 11

Consciousness

Is there any reason, and if so what, for supposing that physical events differ in quality from mental events?

(Russell 1948, 230)

Science has nothing to say as to the intrinsic nature of the atom. The physical atom is, like everything else in physics, a schedule of pointer readings. The schedule is, we agree, attached to some unknown background. Why not then attach it to something of a spiritual nature of which a prominent characteristic is *thought*. It seems rather silly to attach it to something of a so-called 'concrete' nature inconsistent with thought, and then to wonder where the thought comes from. We have dismissed all preconception as to the background of our pointer readings, and for the most part we can discover nothing of its nature. But in one case—namely for the pointer readings of my own brain—I have an insight which is not limited to the evidence of the pointer readings. That insight shows that they are attached to a background of consciousness. Although I may expect that the background of other pointer readings in physics is of a nature continuous with that revealed to me in this particular case, I do not suppose that it always has the more specialized attributes of consciousness. *There is nothing to prevent the assemblage of atoms constituting a brain from being of itself a thinking object in virtue of that nature which physics leaves undetermined and undeterminable.*

(Eddington 1928, 259–60)

Our having recourse to an *immaterial principle*, to account for perception and thought, is only saying in other words, that we do not know in what they consist; for no one will say that he has any conception how the principle of thought can have any more relation to immateriality than to materiality.

(Priestley 1775, xx)

11.1 Conscious Experience

David Chalmers, who figured prominently in the previous chapter, is not alone in regarding it as beyond question that qualities of conscious experiences, the *qualia*, must be utterly unlike physical qualities, unlike kinds of property that figure in explanations in the physical sciences. Materialists (or physicalists; I shall use the labels interchangeably) disagree. Every property, they contend, is a material property. The difficulty for materialism has been to reconcile the 'rich phenomenology', the 'phenomenal character' of conscious experience, with the idea that all there is to conscious agents are dreary states and properties of the kind studied in the natural sciences.

The difficulty accosts philosophers and non-philosophers alike. Here is physiologist Charles Sherrington writing early in the twentieth century:

I see the sun; the eyes trained in a certain direction entrap a tiny packet of solar radiation covering certain wave-lengths emitted from the sun rather less than 10 minutes earlier. This radiation is condensed to a circular patch on the retina and generates a photo-chemical reaction, which in turn excites nerve-threads which relay their excitation to certain parts of the brain, eventually to areas in the brain-cortex. From the retina to the brain the medium of propagation is wholly nervous; that is, the reaction can be subsumed as electrical. Some of this electrical reaction generated in the eye does not reach the brain-cortex but diverges by a side-path to nerve-threads which relay it to a small muscle, which by contracting prevents excess of light attaining the retina. The chain of events stretching from the sun's radiation entering the eye to, on the one hand, the contraction of the pupillary muscle, and on the other to the electrical disturbances in the brain-cortex are all straightforward steps in a sequence of physical 'causation', such as, thanks to science, are intelligible. But in the second serial chain there follows on or attends, the stage of brain-cortex reaction an event or set of events quite inexplicable to us, which both as to themselves and as to the causal tie between them and what preceded them science does not help us; a set of events seemingly incommensurable with any of the events leading up to it. The self 'sees' the sun; it senses a two-dimensional disk of brightness, and overhead shaped as a rather flattened dome, coping the self and a hundred other visual things as well. Of hint that this scene is in the head there is none. Vision is saturated with this strange property called

'projection', the unargued inference that what it sees is at a 'distance' from the seeing 'self'. Enough has been said to stress that in the sequence of events a step is reached where a physical situation in the brain leads to a psychical, which however contains no hint of the brain or any other bodily part. (Sherrington 1906, xx–xxi)

Sherrington concludes, 'the supposition has to be, it would seem, two continuous series of events, one physic-chemical, the other psychical, and at times interaction between them' (1906, xxi).

Sherrington speaks as a pioneering physiologist and neural anatomist. B. A. Farrell finds a similar chain of reasoning among the psychologists.

Like physiologists, the ordinary working psychologist would be quite pleased in a way to get rid of sensations, feelings, etc., as items of experience and deal solely with reactions, discriminations, behaviour-readinesses, and so on. But he cannot bring himself to do so. For he sees himself faced by the old unpalatable alternative. To get rid of 'experience' can be done only by denying that we have sensations, etc., or by refusing to bother with them. But to assert that we do not have sensations, or that no experiences occur, is to assert what is palpably false; and to refuse to bother with them is to leave out certain phenomena, or aspects of phenomena, that psychologists are supposed to investigate. (1950, 174)

Farrell continues:

Some psychologists, chiefly American I think, have paraded their embarrassments (being less inhibited than their British colleagues), and have tried to deal with them. For example, Tolman (of California), in his *Purposive Behaviour in Animals and Men*, distinguishes between discriminations, discriminatory readinesses and so forth, on the one hand, and what a psychology of discriminations leaves out on the other. What it leaves out he calls 'raw feels'. These raw feels, he says, are not capable of scientific treatment, and he admits at the end of the book that he does not know what on earth to do with them. He suggests three scrap heaps onto which they may be thrown. They may be ignored as scientific will of the wisps. They may be assumed to correlate consistently with our responses and response readinesses, so that in so far as X and Y behave alike, they have the same sort of experience. Thirdly, 'Raw feels may be the way physical realities are intrinsically' so that, e.g., experienced qualities may be 'the intrinsic nature of a nervous process'. (1950, 174; see Tolman 1932, 426–7)

While Sherrington flirts with mental–physical parallelism, Farrell depicts Tolman as foreshadowing Chalmers's distinction between states of awareness ('discriminations, discriminatory readinesses and so forth') and conscious qualities ('raw feels'). Tolman's second option closely resembles Chalmers's invocation of 'natural supervenience', and it should be noted that Chalmers, elsewhere in *The Conscious Mind*, indicates that he is not unsympathetic to something like Tolman's third option: 'raw feels', the qualia, are 'the way physical realities are intrinsically'. The first option is not a live one for either Tolman or Chalmers, nor for Galen Strawson, who speaks for many in noting that 'experience is itself the fundamental given natural fact...there is nothing more certain than the existence of experience' (2006, 4). The problem is that it is hard to see how experiences and experiential qualities *could* be reconciled with the physical universe as revealed by the physical and biological sciences.

11.2 Investigating Conscious Experiences

Pretend for a moment that you are a psychologist who accepts, however grudgingly, the importance of conscious experience. You set out to study the phenomenon of consciousness. How might you do this? 'Let us remind ourselves how psychologists do set to work to deal with "experience"—with our sensations, our feelings, and so on; and what sorts of discoveries they make about them.'

Consider sensation. What a psychologist does is to take Mr. X and use him as a subject in a laboratory. He might use Mr. X as a subject by, for example, subjecting him to the important classes of stimuli that are likely to affect X's sense organs—viz., mechanical, thermal, acoustic, chemical, and photic stimuli. In this way [the psychologist] will discover from X's responses and discriminations whether, for instance, X is colour blind or not, or has abnormally acute acoustic sensitivity, and so on. (Farrell 1950, 172)

As an observer, you might also put yourself in the experimental subject's place.

What is important to note is that by playing the role of observer–subject, [the psychologist] does not add anything to the discoveries of psychological science that he could not in principle obtain from observation of X alone;

and no new concepts are required to deal with what his own subject–observation reveals which are not also required by what was, or can be, revealed by his observation of *X*. (Farrell 1950, 173)

Could this be right? When you take the place of the experimental subject, you discover *what it is like* to be in the subject's shoes! How could you possibly work *that* out merely by observing the subject? In an influential paper published three decades ago, Frank Jackson argued that, in learning what it is like to undergo a particular kind of experience, you learn 'new facts', you learn something impossible to learn simply by studying experiences of others 'from the outside'.[1] The 'new facts' seem not to be among the 'physical facts'.

Even if this were true, however, it is orthogonal to Farrell's point. Farrell's idea is that describing your experiences is a matter of describing *what* you are experiencing. When experimental subjects describe their visual sensations, they are describing how things look to them, what they see or think they see. Their experiences might differ from yours—or those of Jackson's scientist, Mary—in that things look different to them. If a subject is red–green colour blind, for instance, things will not look the same to him as they would to a normal subject.

In describing how things look (or sound, or smell, or feel) to you, what are you describing? One thought is that you are describing an image located in your mind, Sherrington's 'scene in the head', a mental tableaux distinct from, although almost certainly intimately related to, an external 'stimulus' for which it purports to go proxy. This picture is encouraged by Sherrington and Tolman, but even without their help, it is one that might occur to you quite naturally.

Think of the familiar response to the question whether a tree falling in a deserted forest makes a sound. Ask any clever undergraduate and you will be assured that this just depends on whether you mean by 'sound' *sound waves* (in which case the tree does make a sound) or *heard sound*, sound *experience* (in which case the tree makes no sound). Your *hearing* a sound is a matter of your undergoing a particular kind of experience. That experience is what occurs *as a*

---

result of being stimulated in a particular way. In perception the stimulation is external and your experience, in important but hard to specify respects, *resembles* its cause; in dreams and hallucinations experiences are triggered internally. What seems puzzling is that characteristics of experiences, so familiar to us as subjects, are omitted when you study subjects in the psychological laboratory or when you probe their brains directly.

This appears not to be a matter of our lacking the right instrumentation. However well equipped, no neuroscientist expects to find imagistic entities, *eidola*, in the brain. Ask yourself, what would such entities be *like*? Visual images, perhaps, would be picture-like, but what of auditory images? Might you detect these with tiny, embedded microphones? And would olfactory experiences give off faint odours, pains be painful to the touch? No, experiences and their qualities evidently elude scientific scrutiny, detection from the 'outside'.

This is just what you would expect were conscious qualities non-physical, purely mental, items. But what are the alternatives? A scientist can afford to ignore what could not be studied anyway, but this does not let the philosopher off the hook. Nor, as Sherrington and Tolman note ruefully, does it really let scientists off the hook. We want our story about the fundamental nature of the universe to be inclusive and complete. It would be unseemly simply to accept that there are fundamental, highly salient, and to all appearances entirely natural, features of the universe fenced off from scientific investigation.

11.3 Representationalism

To see more clearly what is at issue, return to the picture painted by Sherrington. When you see the sun you undergo a visual experience caused by the sun. The sun emits light radiation that is intercepted by your retina and transduced into signals passed along the optic nerve and onward to various brain centres. Activation of one or more of these centres, or perhaps a holistic pattern of activation, 'corresponds' to your visually experiencing the sun.

The weasel word here is 'correspond'. What is it for a pattern of activation, a dynamic state of your brain, to *correspond* to your visual experience? Pretend that the state in question plays the role of your experience in your psychological economy. That is, the state is brought about by your visually encountering the sun, it disposes you to believe that you see the sun, and it equips you to give a description of your experience. This state, then, is in *some* way intimately connected with your experience.

One possibility is that your brain is in a state that *causes* your experience. Chalmers would put this by saying that your experience 'arises from' your brain state. Others would describe your brain state as 'grounding' or serving as the 'substrate' of your experience. You might think in one of these ways if you were already convinced that your experience must be non-physical. And you might think your experience must be non-physical if you thought that its characteristics, its nature, could not possibly belong to something physical, and in particular could not possibly belong to a state of your brain.

Although it can seem obvious that your experience is one thing, your brain state something else entirely, *is* it obvious? What exactly *are* the characteristics of your experience that make it crazy to imagine that the experience might turn out to be a brain state or process? Recall Sherrington's description. When you see the sun, you sense 'a two-dimensional disk of brightness, and overhead shaped as a rather flattened dome, coping the self and a hundred other visual things as well'. As he then observes, 'of hint that this scene is in the head there is none'.

Note, however, that Sherrington's description is a description, not of your experience per se, but of *what* you are experiencing. This is Farrell's point. Your experience is an experience *of* the sun, or *of* a tomato, or *of* freshly baked bread. It would be a mistake to confuse features of *what* you are experiencing with features of your *experience*. When you perceptually encounter a ripe tomato you have, let us suppose, a visual experience *of* something red and spherical, and an olfactory experience of something pungent, but it does not follow that your *experience* is red, spherical, and pungent. In fact, *whatever* your conception of the mind, it is hard to make sense of the thought that your *experience* might be red, spherical, or pungent.

Perhaps, then, the conviction that experiences must be non-physical accompaniments of physical goings-on in the brain, stems from an elementary mistake: the mistake of conflating characteristics of whatever your experience is an experience *of* with characteristics of the experience itself. U. T. Place dubs this kind of mistake the *phenomenological fallacy* (Place 1956, 48–50). The tomato you experience is red and the sound you experience is shrill, but your *experiences* of these things are neither red nor shrill.

Consider an apparently more difficult case, a case discussed by J. J. C. Smart in his defense of the contention that sensations are brain processes. As a result, perhaps, of fixating on a bright blue-green colour patch then looking away, you undergo the experience of a 'yellowy-orange after-image'. The after-image is saturated with colour in a way nothing physical in the vicinity is, and certainly nothing in your brain is. How *could* your experience of the after-image be a process in your brain? Smart's comment:

I am not arguing that the after-image is a brain-process, but that the experience of having an after-image is a brain process. It is the *experience* which is reported in the introspective report. Similarly, if it is objected that the after-image is yellowy-orange but that a surgeon looking into your brain would see nothing yellowy-orange, my reply is that it is the experience of seeing yellowy-orange that is being described, and this experience is not a yellowy-orange something. So to say that a brain-process cannot be yellowy-orange is not to say that a brain-process cannot in fact be the experience of having a yellowy-orange after-image. There is, in a sense, no such thing as an after-image or a sense-datum, though there is such a thing as the experience of having an image. . . . Trees or wallpaper can be green, but not the experience of seeing or imagining a tree or wallpaper. (Smart 1959, 150–1)

Notice that the claim here is not that experiences have *no* qualities. If experiences are, as Smart contends, goings-on in the brain, then experiences have whatever qualities goings-on in the brain have. These qualities, however, are, it would seem, unproblematically physical.

Could it be this easy? Could it be that the chief stumbling block to the identification of conscious experiences with goings-on in the brain is nothing more than a fallacious pattern of reasoning? Certainly, the

mind–body problem as posed by Sherrington and by Tolman has an air of utter hopelessness about it. Maybe this is precisely because it stems from an assumption that undermines *any* potential solution, or any solution that does not embrace dualism. If that assumption is confused, it ought to be dispensable, and if you dispense with it, the problem would, it seems, no longer arise. The solution of the riddle is that 'the riddle does not exist' (Wittgenstein 1921, § 6.5).

The idea that features of experiences are one thing, features of what is experienced something else altogether has come to be associated with what G. E. Moore described as the 'diaphanous' or 'transparent' character of experience (Moore 1903). And insistence on the 'transparent' nature of experience has come to be a hallmark of contemporary 'representationalist' (or 'intentionalist') accounts of conscious experience.[2] Moore puts it this way.

The moment we try to fix our attention upon consciousness and to see what, distinctly, it is, it seems to vanish: it seems as if we had before us a mere emptiness. When we try to inspect the sensation of blue, all we can see is the blue: the other element is as if it were diaphanous. (Moore 1903, 450)

In the hands of representationalists, Moore's observation has spawned an aggressive, fully fledged research program founded on the conviction that conscious mental states are 'intentional' states. Beliefs, thoughts, judgements seem uncontroversially intentional; that is, they are *of*, *for*, or *about* this or that; they have *objects*, 'intentional objects', that might or might not exist. Your thoughts of Vienna and of Middle Earth, are *of* Vienna and *of* Middle Earth. According to the representationalists, the key to unlocking the mystery of consciousness is that this is the case for *conscious* states as well.

Your visual experience of a ripe tomato is an experience *of* something red, spherical, pungent. The tomato, not your experience, is red, spherical, and pungent. The scene you are experiencing is saturated with colour, perhaps, but it does not follow, indeed it makes no sense to suppose, that your *experience* is saturated with colour: tomatoes are coloured, not experiences. When you describe your experience of the tomato, you describe the tomato, ways the

2 See, for instance, Harman (1990), Dretske (1995), Lycan (1996), Tye (1995), and Byrne (2001). For a succinct overview, see Tye (2009).

tomato is or appears to be. You are not describing features of your perceptual state. When you are aware of your experience, you are aware of it *indirectly*, aware of it as an experience *of* a red tomato. In this way, representationalists locate potentially embarrassing mental qualities in representational, intentional *objects*. If these qualities seem not to fit comfortably into accounts of what goes on inside the heads of sentient creatures, this is not because they are non-physical phantasmagoria, but because, in taking them to be qualities of our *experiences*, we mislocate them.

Note that the qualities themselves—redness, for instance, or sphericality—are, by and large, physically unremarkable. The way the tomato appears to you is not an entity, an appearance that occurs inside your head. The tomato's appearance is something that could be captured in a photograph and measured with various instruments. Conscious qualities take on a mysterious air only when you imagine that they are qualities of *experiences* rather than qualities of objects experienced. You have the tomato's redness *plus* experienced redness. How could this *experiential* or *phenomenal* redness possibly fit into the physical universe?

Pretend for a moment that you are a no-nonsense neuroscientist convinced that conscious experiences are states of the brain, and that you are looking at a ripe tomato in bright sunlight. Judicious experimental work has convinced you that a particular dynamic state of your brain *is* your visually experiencing the tomato. If members of your research team were to look for the redness and sphericality of what you are experiencing in your brain, they would assuredly be disappointed. This is not because the redness and sphericality you experience are qualities of a private, non-physical, mental something, however, a something to which only you have 'access', it is because these qualities are qualities of what you are experiencing, the tomato. Were you to hallucinate—or imagine, or dream of—a ripe tomato, you might go into a state intrinsically similar to the state you go into (the very state your colleagues might be observing) when you see a ripe tomato. The difference would be that, in such cases, the cause of your going into the state would not be something red and spherical, not a tomato. Think here of Smart's 'yellowy-orange after-image'.

What of bodily sensations? You might consider the sharp pain you feel in your foot as purely sensory, as a kind of 'raw feel'. But this

feeling, according to the representationalists, is constituted by your representing your foot as being a certain way—significantly, a way it might or might not actually be: recall the phenomenon of 'phantom pain'.

11.4 'Diaphanous' Experiences

So where are we? Representationalists tell us that consciousness appears mysterious because, when we introspect conscious experiences, we see 'through' those experiences to their objects, the properties of which, as we have seen, are easily mistaken for properties of the experiences themselves. When asked to describe your experiences, you describe what you are experiencing—what else! The philosophical mistake is to take this description as a description of a state, a state of your brain or a state 'grounded in' a state of your brain. When you do that, you immediately notice that the characteristics you mention in describing your experience are utterly unlike characteristics of the brains of experiencing agents. We seem compelled to accept these characteristics as either non-physical or non-existent. The latter option is not a live one, however, so we find ourselves committed to some form of dualism.

Central to representationalism is the idea that there is nothing in the least remarkable about the characteristics you dwell on in describing your experience. (At any rate this is so for ordinary visual experiences.) These characteristics are perfectly ordinary characteristics of perceived objects. They come to seem remarkable only when you try to locate them, as Sherrington does, inside the heads of perceivers. This is the phenomenological fallacy.

Representationalism can be spelled out in many different ways. My interest here is not in the details, but in the thought that, what you might regard as intrinsic features of experiences, their 'feel', what it is like to undergo them, are not characteristics of the experiences, but characteristics of what the experiences are experiences *of*, 'representational contents'. Most representationalists would grant that conscious experiences *have* intrinsic features (how could they *not?*), but insist that these are not what you are aware of when you undergo, introspect, or describe an experience. The primary motive for a view

of this kind would be the conviction that qualities we commonly associate with conscious experiences are unlikely to be physical qualities of experiencing creatures or their brains.

Although representationalism is in many ways an attractive option, it is easy to doubt that representational theories succeed in fobbing off *all* introspectable experiential qualities. A visual experience of a red tomato need not itself be red or spherical. This is a true and important point. It does not follow that the experience has no qualitative character evident to the agent undergoing the experience, however (Heil 1970). Indeed, Moore himself makes this very point immediately after remarking on the 'diaphanous' character of conscious experience. Here is the whole passage.

The moment we try to fix our attention upon consciousness and to see what, distinctly, it is, it seems to vanish: it seems as if we had before us a mere emptiness. When we try to inspect the sensation of blue, all we can see is the blue: the other element is as if it were diaphanous. Yet it *can* be distinguished if we look attentively enough, and if we know there is something to look for. (Moore 1903, 450)

Philosophers have sometimes expressed this thought by noting that 'there is something it is like' to undergo a visual experience of (or imagine, or dream, or hallucinate) a red tomato (Farrell 1950, 183; Nagel 1974). Representationalists contend that what an experience is like is exhausted by its representational object. When you reflect on what it is like to undergo a visual experience of a ripe tomato, the characteristics that occupy your thoughts are characteristics you represent the tomato as having.

Representationalists have offered ingenious responses to apparent counter-examples to the representationalist thesis. You represent the tomato as being uniformly red, for instance, but in one important sense the tomato does not *look* uniformly red any more than a round coin *looks* round when viewed obliquely. (One way to see this is to note that a 'realistic' painting of a uniformly red tomato would use a variety of pigments, a drawing of a round coin on a table would have an elliptical shape.) At one time it was common for theorists to account for such cases by positing 'sense data' as objects of 'immediate awareness'. The coin is round, but the object of your immediate awareness. The coin is round, but the object of your immediate

awareness is elliptical, so what you are immediately aware of must be something—a mental something—other than the coin.

Representationalists believe they can accommodate such cases. The appearance of a tomato or a coin, how such things look, *is* a perfectly objective feature of the universe. As noted earlier, the appearance of a tomato or a coin could be photographed. A photograph of a round coin would capture the coin's roundness by means of an oval shape, or a uniformly red surface by means of an array of distinct colours. Similarly, you might visually represent a round coin or a uniformly coloured surface by means of a representational state the content of which is how the coin and tomato look to you. In this case, experienced features that are not present in the coin *are* present in the 'optic array', structured ambient light responsible for visual perception.[3]

The representationalist insistence on distinguishing features of objects experienced and features of experiences is unassailable. What is less clear is whether experiences are thereby shown to be altogether lacking in qualitative conscious trappings. Suppose the representationalists are right. Suppose qualities you would naturally regard as qualities of conscious experiences are in fact qualities of what those experiences are experiences of. Even after subtracting these qualities from the mix, it might still be the case that, as Moore suggests, you could identify qualitative aspects of experiences themselves, aspects that have a bearing on what those experiences are like. These qualities would be intrinsic qualities of experiences, *conscious* qualities not accounted for representationally. The question then would be whether these qualities are irredeemably 'non-physical', whether they are incommensurate with qualities of perceivers' brains. My sense is that, although there is more to conscious experiences than representationalists allow, qualities of experiences themselves can indeed be accommodated without a leap to dualism.

You can see where this is going by imagining how things might be for an agent equipped with a visual apparatus very different from ours. Philosophers speculate as to how the universe might appear to bats or to Martians, but you can reflect on more accessible possibilities

without delving into science fiction. Consider the TVSS: the Tactile Visual Simulation System.[4] The original TVSS consisted of a television camera that transmitted an 'image', not to a television monitor, but to a 40 × 40 grid of metal pins in contact with a subject's skin, usually pressed against the subject's stomach or back. These pins vibrate in a pattern roughly corresponding to the pattern of illuminated pixels on a television monitor. Blind subjects equipped with a TVSS can learn to identify and characterize objects in their environment in ways analogous to those of sighted subjects. Researchers have described the apparatus as enabling subjects to 'see with the skin', a description I regard as entirely apt. 'Tactile vision' resembles ordinary vision functionally. Thus, as in the case of sight, the apparatus conveys information carried by structured ambient light radiation; it ceases to work if the scene is not illuminated; and objects can be 'visually occluded' by opaque (but not transparent) objects placed between them and the camera.

A TVSS cannot register colours, and the resolution of a TVSS does not approach the resolution of the human visual system. This does not affect the point I aim to make here, however. Imagine a blind subject, Mary, equipped with a high resolution TVSS and a colour blind, visually impaired subject, Harry, whose visual acuity roughly matched that of Mary's enhanced TVSS. Functionally, Mary and Harry are on a par. And to the extent that their visual experiences are representational, they are on a par representationally. Mary's descriptions of her experiences are indistinguishable from Harry's. Yet Mary's and Harry's visual experiences differ qualitatively. *What it is like* for Mary to experience a rotating coffee cup on the table in front of her differs dramatically from what it is like for Harry to undergo the corresponding experience. We can accept that, if their experiences are of a light-coloured, shiny, cylindrical object, it does not follow that their experiences are light-coloured, shiny, and cylindrical. Well and good. But neither does it follow that their experiences altogether lack 'experiential qualities'.

---

3 See Gibson 1966. See Tye 2002, 2009 for extensive discussion of, and responses to, apparent difficulties for representationalism.

4 For detailed discussion of the TVSS (with references), see Heil 1983, 13–18, 74–7; and Heil 2003a, 227–9.

11.5 Experiential Qualities

You might think of a normal (human) visual system and a TVSS as affording different modes of conscious visual perception. If perceiving is representing, it is representing in some particular way. This becomes clear when you ask how representationalists who hold that all there is to the qualitative character of conscious experiences is what is represented by those experiences could account for the distinction between *conscious* and *non-conscious* instances of representation.

By 'non-conscious representation', I have in mind the operation of 'subpersonal' representational mechanisms routinely posited by scientists studying the self-regulation of sentient creatures' bodies (see Martin 2008, chap. 9). The operation of mammalian thermoregulatory systems and digestive systems, for instance, would be difficult to explain unless components of those systems were taken to represent particular bodily states—and in considerable detail. The question is, what must be added to the deployment of representations by a biological system for it to constitute a conscious experiencing?

Might conscious representations be distinguished from non-conscious representations *functionally*? Thus, a conscious representational state might be a state available to the 'higher cognitive processes'. Representational states implicated in the operation of digestive or thermoregulatory systems do not report to a central cognitive command post, hence they are invisible to creatures in which they occur and might, on that account, safely be classified as non-conscious.

Here, as elsewhere, functionalism could be thought to get the order of explanation backwards. A system does not lack experiential qualities because of its role in a larger, more encompassing system. Rather its role in the system is partly determined by how it is qualitatively. One lesson of the TVSS is that conscious experiences include a qualitative 'aura' determined by the nature of the representational *medium*. This 'aura' is not simply an incidental accompaniment of various representational states. It is a vital part of what makes those states apt to play the roles they play in the psychological

economies of sentient creatures. More importantly for this discussion, it is what finally distinguishes conscious from non-conscious representational states.

The much-discussed phenomenon of 'blindsight' is suggestive in this regard (see Weiskrantz 1986; Heil 1983, 77–81). A patient with blindsight might be thought to harbour visual representations accessible to 'higher processes' that are, nevertheless non-conscious. Such representations might functionally resemble ordinary conscious visual experiences. They would, however, lack the 'aura' definitive of such experiences.

I mean these examples to be suggestive only. I am not prepared to offer anything approaching a definitive refutation of accounts of conscious experience that hope to dispense with qualities of such experiences. Definitive refutations are rare in philosophy and, in ontology, non-existent. My aim is only to make clear that a lingering worry about the accounts in question is perfectly legitimate and, in due time, to offer an alternative account, one that avoids the *prima facie* implausibilities of this one.

Let me be clear about what I take myself to have accomplished thus far. First, I accept an important aspect of representationalism. Conscious experiences, as experiences *of* this or that, are representational. Second, it is a mistake—an instance of Place's 'phenomenological fallacy'—to conflate qualities of what is represented in an experience with qualities of the experience itself. Just as a description of a red tomato can be 'of a red tomato' without being red, so your experience of a red tomato can be 'of the tomato' without being red. But, third, it is no less a mistake to imagine that experiences themselves, experiences *qua* representations, are qualitatively featureless, and, in particular, that they lack *conscious* qualities. The significance of representational states, I shall argue in chapter 12, stems from uses to which they are put by representing agents. This holds for conscious experiences just as it does for drawings, maps, and ordinary linguistic utterances or inscriptions. The important point, a point most representationalists could accept, is that qualities of experiences need not be at odds with qualities of the brains of agents undergoing those experiences.

11.6 Could Experiences Really Be 'Brain Processes'?

In this and the previous chapter, I have examined two approaches to the venerable mind–body problem. The first approach, a version of property dualism, *begins* with a commitment to the ineliminability of distinctive qualities of conscious experience, qualities taken to be wholly unlike anything comprehended in the physical sciences. The second sees states of mind as one and all representational and endeavours to accommodate qualities of conscious experiences by relegating them to representational objects, qualities of which can, but need not, correspond to qualities at home in the physical universe. (Qualities you represent a tomato as having reside in the tomato; qualities of an hallucinated yellowy-orange unicorn are nowhere.) On this second approach, representational states themselves, including those constituting perceptual experiences, need exhibit no non-physical properties, no properties the possession of which would give rise to puzzlement or mystery.

I have argued that property dualism, at least in the guise of non-reductive physicalism, is ontologically hopeless. Ontological embarrassments, however, have not deterred non-reductivists, who seem willing to sacrifice almost any metaphysical principle in order to preserve the faith. This is most clear in the current willingness to diddle with accounts of causation so as to reconcile the conviction that mental states and properties are distinct from, but dependent on, physical states and processes, with the evident reality of 'mental causation', the 'causal relevance' of the mental.

Dualists and representationalists evidently agree that *qualities* are the problematic entities. In earlier chapters I argued, on independent grounds, that genuine properties are powerful qualities. This forces qualities into the picture on the ground floor. Such a conception of properties fits naturally with the idea defended by Spinoza, Davidson, and more recently, Galen Strawson, that the mental–physical distinction is not a real distinction, but one of conception only. This will be true of the quality–power distinction as well. Every property can be considered as a quality or as a power. But this in no way implies a 'real distinction', in no way implies that qualities and powers are distinct kinds of property or distinct 'aspects' of properties.

I part company with Strawson in his invocation of panpsychism, however. I am sceptical of the thesis that *everything*—fundamental particles, as well as conscious agents—exhibits experiential qualities. This, Strawson contends, is implied by the ineliminability of conscious experiencing and the hopelessness of both non-reductivism and emergent dualism. Emergence, Strawson argues, is wholly mysterious if it is understood as the generation of an entirely new kind of quality solely by arranging material bodies that themselves altogether lack the quality. An emergent quality would have to be a quality of the arrangement, not a quality of one or more of its constituents. But unlike familiar qualities of arrangements—the redness and sphericality of a tomato, for instance—an emergent quality would not simply be what you get when you arrange the right things in the right ways, it would be something with its own independent ontological standing, an authentic 'addition of being'.

One way to see what is at stake here is to put the point in terms of truthmakers. You can say truly that the tomato is red, spherical, and has a particular mass. Truthmakers for such claims are the constituent substances, duly propertied, duly arranged, and duly interacting as they are. The tomato's redness, sphericality, and mass are not properties of the tomato, emergent or otherwise. The only properties on the scene are those belonging to the fundamental constituents of the tomato. In contrast, a genuinely emergent property of the sort that troubles Strawson, *would* be a property of a complex whole, something produced by—or 'arising from'—the whole's constituents organized as they are.

It is one thing to doubt that emergence of this kind is possible, another thing to give up on emergence altogether. Rightly understood, emergence is ontologically unobjectionable (see §§ 2.6–2.7). Properties emerge when substances emerge—in collisions among fundamental particles, for instance. Positrons are routinely produced by interactions among particles, none of which are positrons. Quantum entanglement might count as another kind of emergence (see § 3.8). In all such cases, what might properly be said to emerge are substances. These substances are various ways, and these ways are properties of emergent substances: emergent properties.

Qualities of conscious experience might possibly emerge as positrons do or as a consequence of quantum entanglement involving the

nervous systems of sentient creatures, but the actual story seems likely to be much less dramatic. Conscious systems might reasonably be thought to require particular sorts of highly complex arrangement of the fundamental things. Consciousness is one of many kinds of state into which such systems are capable of entering. Their occupying a conscious state, their being conscious, is a matter, not of emergence, but of their constituents being organized as they are, including all their various interrelations with one another and with their extra-systemic environments. These dynamic arrangements serve as truth-makers for claims about sentient creatures' conscious experiences.

Strawson is unhappy with such a picture. Conscious qualities, he argues, are not the *sorts* of quality that could be 'built up' from arrangements of constituents that themselves lacked such qualities. You can see how something could be spherical in virtue of being an arrangement of constituents, none of which is spherical. You can see how something could be red in virtue of being put together in the right way from constituents, none of which is red. But consciousness is another matter. The experiential is not the kind of thing that could be obtained by assembling non-experiential elements in the right way.

You can get liquidity from non-liquid molecules as easily as you can get a cricket team from eleven things that are not cricket teams. In God's physics, it would have to be just as plain how you get experiential phenomena from wholly non-experiential phenomena. But this is what boggles the human mind. (Strawson 2006, 15)

Indeed, Strawson insists, the idea that conscious experiential qualities could arise from organizations of constituents wholly lacking in such qualities is flatly unintelligible. The relation of the liquidity of water to water's non-liquid constituents is, he thinks, nothing at all like the relation between conscious experiences and sentient creatures' non-sentient constituents. According to Strawson, trying to get conscious wholes from non-conscious parts would be like trying to build up an object with spatial features, an object with a definite size and shape, for instance, using wholly non-spatial constituents. Both possibilities are unintelligible. Just as spatial parts are required for spatial wholes, so conscious experiential parts are required for conscious wholes. Spatial constituents are 'intrinsically suited to constituting spatial

phenomena' (Strawson 2006, 21). In the same way, constituents of sentient creatures must be 'intrinsically suited' to constituting creatures capable of conscious experiences.

So you need experiential qualities, sparks of consciousness, in on the ground floor. What about the possibility that the fundamental things, or some species of fundamental thing, are 'proto-conscious': although not conscious themselves, they contain *seeds* of consciousness that yield experiential qualities when they are appropriately organized? This would be analogous to the construction of objects with various shapes and sizes from constituents that, while not spatial themselves, contained *seeds* of spatiality, 'proto-spatial' entities. How, Strawson asks, is this an advance over the original case? To obtain an object with spatial qualities, you need constituents with spatial qualities. Just as the notion of 'proto-spatiality' is hopeless, so the notion of 'proto-experiential' lacks application. Experiential wholes require experiential parts.

That is what I believe: experiential phenomena cannot be emergent from wholly non-experiential phenomena. The intuition that drives people to dualism (and eliminativism, and all other crazy attempts at wholesale mental-to-non-mental reduction) is correct in holding that you can't get experiential phenomena from ... shape-size-mass-charge-etc. phenomena, or, more carefully now ... from non-experiential features of shape-size-mass-charge-etc. phenomena. So if experience like ours (or mouse experience, or sea snail experience) emerges from something that is not experience like ours (or mouse experience, or sea snail experience), then that something must already be experiential in some sense or other. It must already be somehow experiential in its essential and fundamental nature, however primitively or strangely or (to us) incomprehensibly; whether or not it is also non-experiential in its essential nature, as conventional physicalism supposes. (Strawson 2006, 24)

Is this right? I have suggested that the mental–physical distinction is one of conception only, not a real distinction. I understand this as Strawson's view as well. But if that is so, there is no question of trying to obtain mental qualities from non-mental physical qualities, qualities of the fundamental physical things duly arranged. This is not because there is an unbridgeable ontological gap between the physical and the mental, a gap analogous to the gap between spatial

and non-spatial qualities. It is because the distinction implicit in talk of mental qualities as distinct from physical qualities is ill-conceived from the start. Strawson's advocacy of 'micropsychism' and panpsychism is, in this regard, potentially misleading. Once you accept the idea that the mental–physical distinction is not a deep distinction, the question of whether the mental emerges or was there all along ceases to arise.

11.7 Manufacturing Qualities

Is this just terminological sleight of hand? How plausible is it to think that experiential qualities and qualities physicists ascribe to the fundamental things differ 'in conception only'? You will certainly not think it plausible if you start with vivid descriptions of your own conscious experiences and compare these with descriptions of electrons, quarks, and gravitational fields.

Such an approach would foreclose the possibility of progress in advance. You need not be a hard-core representationalist to accept the idea that many things you might want to say about our conscious experiences concern not qualities of the experiences themselves but qualities of what the experiences are experiences of, their 'objects'. And you need not be an emergentist to recognize that qualities found in familiar ordinary objects appear to differ in dramatic fashion from those physicists ascribe to electrons, quarks, and gravitational fields. Nevertheless, what makes it true that an ordinary object is as it is qualitatively will be a dynamic arrangement of the fundamental things. You have the particles as they are, interactively arranged as they are, making it true that this is a red, spherical tomato. You do not have the particles as they are, interactively arranged as they are, *plus* a tomato, *plus* the tomato's qualities.

Earlier, siding with Strawson, I noted that you can understand how you could obtain a liquid by putting together particles (none of which were liquid) in a particular way, and how you could get redness—this tomato's being red, for instance—by organizing particles (none of which were red) in the right way. But consider the tomato's redness more closely. Scientists could explain why, if you have colourless particles distributed as they are on the surface of a tomato, you have

something red. But I wonder how well we *do* understand this. It was precisely the difficulty of understanding how this is supposed to work for colours and other secondary qualities of complex objects that led Galileo and Berkeley to declare that colours were not 'out there', but existed only in the mind.

Nowadays, we are more comfortable with the idea that coloured surfaces are objective features of the universe. Yes, I know that colour science raises all manner of complications here. But the point I want to make does not depend on an aggressive form of 'colour realism'. My point concerns our attitudes toward colours and their 'physical basis', If you thought that colours are 'out there', and you thought that the tomato's redness *is* its having the right kind of distribution of (colourless) particles on its surface, then you are unlikely to find the 'gap' between particles making up coloured surfaces and those coloured surfaces mysterious or puzzling.

My suggestion is that the 'gap' many philosophers find between, on the one hand, goings-on in the nervous systems of conscious creatures and, on the other hand, qualities of conscious experiences, is no more mysterious than this. And, yes, no less mysterious either. It is overwhelmingly likely that in such cases what we need is not more theory, not more argument, but analogies and reminders of how we think of other phenomena. We, many of us anyway, are comfortable with the idea that observable features of complex objects result from arrangements of particles that themselves lack these features. The relations in such cases—the relations between objects' features and the arrangements that make them up—are neither more nor less mysterious than relations between qualities of conscious experiences and qualities of the particles that make up our nervous systems.

In reflecting on such matters, it is all too easy to make things harder on ourselves than necessary—by, for instance, implicitly conflating qualities of experiences with qualities of what is experienced. The idea is not to expunge conscious qualities, but to get clear on what we are talking about when we talk about conscious qualities. Once you are clear on that, you are in a position to ask yourself whether the 'gap' between those qualities and qualities of constituents of our nervous systems, is greater or more surprising than the 'gap' between qualities of ordinary objects and qualities of *their* constituents.

Appreciation of these points defuses much of the force of worries as to how something with mental qualities *could* turn out to be an arrangement of electrons and quarks or a fluctuation in a field. Still, I suspect that there are residual concerns. I have suggested that experiences themselves have qualities discernible by sentient creatures undergoing those experiences: experiences are not wholly 'transparent'. And how do we get *these* qualities into the picture? Must we locate them in the electrons and quarks, or assume that they are emergent?

Again, I think this question betrays an implicit commitment to a deep ontological divide between the mental and the physical. My suggestion is that we are in a position to make progress only once we are prepared to renounce this divide—and not by 'reducing' the mental to the physical. So the question here should be, what exactly *prevents* us from supposing that qualities of conscious experiences are to the nervous systems of sentient creatures as the qualities of a tomato are to the dynamic arrangement of particles that make it up? Russell and Eddington, in passages quoted at the outset of this chapter, echo Priestly, who sums up the idea elegantly:

> It is said we can have no conception of how sensation or thought can arise from matter, they being so very different from it, and bearing no sort of resemblance to *figure* or *motion*; which is all that can result from any modification of matter, or any operation upon it.
>
> But this is an argument which derives all its force from our *ignorance*. *Different* as are the properties of sensation and thought, from such as are usually ascribed to matter, they may, nevertheless, inhere in the same substance, unless we can shew them to be absolutely *incompatible* with one another. There is no apparent resemblance between the ideas of *sight*, and those of *hearing*, or *smelling*, &c. and yet they all exist in the same mind, which is possessed of the very different senses and faculties appropriated to each of them. Besides, this argument, from our not being able to *conceive* how a thing can be, equally affects the immaterial system: for we have no more conception how the powers of sensation and thought can inhere in an *immaterial*, than in a *material* substance. (Priestley 1782, 281)

What of the thought, a thought that would not have occurred to Priestly, that qualities of experiences are 'subjective', and qualities of tomatoes and nervous systems are 'objective'? The idea that the subjective–objective distinction is an ontological distinction, however, the idea that some things—substances? properties?—are

subjective, some objective, appears hopeless (Searle 1992, 93–100 disagrees). In the present context, the distinction seems to reflect a difference between

(a) *A*'s *undergoing* a particular conscious experience, *C*; and
(b) *B*'s *observing A*'s undergoing *C*, perhaps by observing *A*'s brain.

(I leave open the possibility that *A is B*.)

Imagine that you are self-consciously visually perceiving a ripe tomato in bright sunlight. You are undergoing a particular experience. Suppose that your undergoing this experience *is* your brain's being in a particular dynamic state, and that an observer is perceiving your experience by perceiving this state. To keep matters simple, imagine that the observer's perception is visual as well. These experiences, yours and the observer's, are altogether different. But this is not because one is subjective, the other objective, whatever that might mean. *As* experiences, they are dynamic states of the individuals undergoing the experiences, and in that regard 'subjective'. What they are experiences of, in contrast, are perfectly 'objective', entirely open to the public.

But surely there is something about your experience, its experiential quality, that is *not* public, not observable by anyone other than you! I have made things hard on myself here by supposing that experiences have qualities of which experiencers can be aware. If you are a representationalist, you would want to deny this. But in either case, the first important point is that what it is like for a creature to undergo a conscious experience, and what it is like to undergo an experience of that experience, need be nothing at all alike. In the case under consideration, you are undergoing a visual experience of a tomato; the observer is undergoing a visual experience of your experience of a tomato. Why should anyone imagine that the objects of these experiences or the experiences themselves must be alike qualitatively?

If you continue to think that there must be a chasm between the mental and the physical, I invite you to reflect on your reasons for this thought. Ask yourself how many of those reasons are founded on questionable ontological assumptions (that qualities are inevitably and irreducibly mental, for instance), or on a tacit acceptance of the idea that an experience of redness or sphericality must itself be red and spherical, or on the thought that an experience of an experience

(a visual experience of a visual experience of a tomato, for instance) ought somehow to resemble the original experience. Otherwise, the burden would be on you to identify qualities of experiences that could not possibly be either qualities of paradigmatic physical substances or, as is more likely, what you get when you put the right kinds of physical substances together in the right ways.

11.8 Qualities *of* Experiences

What of the aforementioned options? I have accepted the idea that experiences themselves are qualitatively imbued. My suggestion was that two 'visual' perceivers, one with attenuated 'normal' vision, the other equipped with a TVSS, would have qualitatively different visual experiences. You do not have to accept this case to accept its point. But what of these qualitative differences? If they are genuine, then we are, or could be, aware of them. But how does this kind of 'introspective' awareness fit the picture I have been promoting?

When you visually apprehend a tomato, you are aware, not of your experience of the tomato, but of the tomato. Nevertheless, your being aware of the tomato is a matter of your being in a particular qualitative state, and this is something—get ready for this—something of which you are implicitly aware.

Imagine dropping your keys in back of the refrigerator.[5] The refrigerator is too large and heavy to move, so you track down a yardstick, and use it to prod behind the refrigerator, feeling for the keys. Eventually you succeed, you feel the keys at the end of the stick. The stick, and the system of which it is a part, a system that includes your arm and hand, is, in Moore's sense, 'transparent'. You feel 'through' the yardstick to the keys. But the system is not *entirely* transparent. Your making use of the yardstick as you do depends on your feeling *through* the stick. Matters would be very different if your hand and arm were anaesthetized and the system comprising your hand, your arm, and the stick were being manoeuvered by means of ropes and pulleys, guided by a 'yardstick cam' fastened to the end

5 A version of this example appears in Heil 1983, 32, where a somewhat different moral is drawn.

of the stick. The feeling you have of the stick in your hand is the vehicle of your feeling your car keys at the end of the stick.

This is how it is with vision and our other sensory modalities. In fact, something like a qualitative 'aura' is a distinguishing feature of conscious experiences, perceptual and otherwise. Our deployment of mental imagery (discussed in chapter 12) relies on this feature of experience and would be impossible without it.

None of this poses any sort of threat to the idea that the universe owes its nature to the natures of the fundamental things. Dynamic arrangements of the fundamental things serve as truthmakers for *all* the truths that have truthmakers, including all the truths concerning conscious experiences. If you organize *these* fundamental things in a particular way, the result will be an arrangement of which it is true that this is a tomato, this is red, this is spherical. If you take the very same fundamental things and organize them differently, you will produce an arrangement of which it is true that this is a sentient creature undergoing a particular kind of experience. Just as in the case of the tomato you do not have the arrangement *plus* the tomato (or *plus* its redness or *plus* its sphericality), in the case of a sentient creature, you do not have the arrangement *plus* sentience.

Bear in mind that a tomato is not just a 'tomato-shaped collection of particles'. A ceramic tomato is a tomato-shaped collection, but not a tomato. If a tomato is a collection of particles, it is a collection of particles possessing various powerful qualities that enable them to interact with one another and with particles elsewhere in a way that amounts to a tomato. There is no question of our being in a position to derive truths about tomatoes from truths about the particles (or the fields, or the One . . .), from truths expressed in the vocabulary of fundamental physics. This is not because tomatoes and their colours and shapes 'emerge', but because application conditions for 'this is a tomato' are orthogonal to application conditions for terms at home in fundamental physics.

Need I repeat that none of this implies that tomatoes are linguistic artifices, none of it implies anti-realism about tomatoes? Tomatoes are one kind of arrangement of the fundamental things, one kind of arrangement among many. The fact that we have labels for some of these kinds of arrangement and not others detracts not at all from the reality of the arrangements.

This is how it is with conscious experiences as well. If you put the fundamental things together in the right ways, you get a sentient being capable of conscious experiencing, a being for which there is something it is like to be that being. Priestley, Eddington, and Russell are right: *nothing* physics tells us about the fundamental things casts doubt on this possibility. If you reject it, you will need an argument.

This is the picture. Many will spurn it. If you start with vivid descriptions of conscious experiences, then try to derive these from descriptions of the atoms and the void, you are likely to be disappointed. If, in contrast, you start with the atoms and the void, the fundamental things, and recognize that these provide resources for all the diversity we encounter in the sciences and in everyday reflection, you have at least a fighting chance of unravelling the basis of consciousness. At some level we *know* this is how it must be. If we cannot find a way to get conscious experiences into our conception of the universe, it is time to question that conception. Almost certainly it will be ontologically impoverished. One mistake is to deprive the fundamental things of qualities, then wonder at the prominence of the qualitative in experience.

This approach reverses the burden of proof as it is usually presented. My thought is that reasons operative in producing the belief that the mental and the physical differ fundamentally are all unsatisfactory, all rely on tacit presumptions that do not survive scrutiny. The mystery of consciousness is that it is so widely regarded as mysterious. We can do better. The place to start is with serious ontology.

CHAPTER 12

Conscious Thought

Man, though he have great variety of Thoughts, and such, from which others, as well as himself, might receive Profit and Delight; yet they are all within his own Breast, invisible and hidden from others, nor can of themselves be made appear. The Comfort and Advantage of Society, not being to be had without Communication of Thoughts, it was necessary, that Man should find out some external sensible Signs, whereby those invisible *Ideas*, which his thoughts are made up of, might be made known to others. For this purpose, nothing was so fit, either for Plenty or Quickness, as those articulate Sounds, which with so much Ease and Variety, he found himself able to make. Thus we may conceive how *Words*, which were by Nature so well adapted to that purpose, came to be made use of by Men, as *the Signs of* their *Ideas*; not by any natural connexion, that there is between particular articulate Sounds and certain *Ideas*, for then there would be but one Language amongst all Men; but by a voluntary Imposition, whereby such a Word is made arbitrarily the Mark of such an *Idea*. The use, then, of Words, is to be sensible marks of *Ideas*; and the *Ideas* they stand for are their proper and immediate Signification.

(Locke 1690, II, ii, 1)

A dog believes his master is at the door. But can he also believe his master will come the day after tomorrow?

(Wittgenstein 1953: § 174; see also § 650)

A creature cannot have thoughts unless it is the interpreter of the speech of another.

(Davidson 1984, 157)

12.1 Conscious Thinking

The previous chapter addressed the problem of conscious experience, in particular the problem of reconciling qualities of conscious experiences with characteristics of sentient creatures taken to be responsible for those experiences. My contention was that the problem is largely of our own making. It arises from a pair of confusions.

(1) a confusion over what qualities of conscious experiences might be, coupled with a tendency to mis-identify those qualities;

(2) a confusion over the qualitative and dispositional nature of properties.

If you mistake qualities of what is experienced for qualities of experiences, you will find no room for such qualities in the nervous systems of sentient creatures. You will be pushed towards mental–physical dualism. If you take the fundamental properties to be 'pure powers', the very fact that conscious experiences have a qualitative nature becomes utterly mysterious. Again, dualism beckons.

In contrast, if you distinguish experiential qualities from qualities of things experienced and if you agree that properties quite generally are powerful qualities, the way is opened to a rejection of dualism. This is not 'physicalism', it is not a matter of 'reducing' the mental to the physical. It is a matter of giving qualities their due.

In this chapter I extend these points to the cognitive dimension of the mind. I begin with a discussion of the idea that thinking is a kind of inner speech. My target is a conception of language as a privileged vehicle for determinate thoughts, or determinate thoughts of particular kinds. My suggestion is that cognition—conscious thought—is qualitatively imbued, irreducibly imagistic. Language is an indispensable tool, perhaps, but the familiar contrast between linguistic and imagistic thought is ill-founded and misleading.

12.2 Thinking and Speaking

For human beings thought and language go hand in hand. But *must* they? This is most often put as a question concerning the relation of

language to thought. Could thought (or 'genuine thought') occur in the absence of language (or 'genuine language')? Could a creature who lacked a language (a 'genuine language') entertain thoughts ('genuine thoughts')? Even if thought of a rudimentary sort were possible without language, might there be kinds of thought thinkable *only* linguistically, thinkable *only* in language?

Before venturing into deeper waters, a word on terminology. I take the question whether thought requires language to boil down to the question whether conscious thinking, or some species of conscious thinking, is inevitably linguistic. And I take 'linguistic' thinking to involve thoughts 'in' some natural language: English, Portuguese, Urdu. I do not mean to exclude the possibility that some thinking is unconscious or 'subpersonal'. This is a topic to which I shall return in § 12.9. Meanwhile, the focus will be on conscious thought and reflection.

What of the language of honeybees, computer languages, Ockham's *sermo interior*, Fodor's (1975) Language of Thought? These raise interesting questions, but I shall set them aside here. The plan is to concentrate on the clear case of natural language and see where this leads. It might then be possible to move ahead on honeybees, computing machines, and the Language of Thought.

My broader aim in this chapter is to explore a pattern of reasoning about language and thought that seems to me virtually to guarantee a distortion of what precisely constitutes thinking. My guiding idea is a simple one, an idea officially embraced by many philosophers, but an idea the implications of which are too easily neglected. I have in mind the thesis that something functions as a representation only when it is given a *use* by a representing agent. Let me hasten to add that I am thinking of agency here in a broad sense that encompasses, not merely intelligent human beings, but intelligent systems generally. In this sense, a digestive system—the 'brain in the gut'—could constitute an intelligent system, hence count as a candidate agent.

This thesis—that representation requires a representing agent—applies in equal measure to linguistic and to non-linguistic, 'pictorial' representations. What I would like to call attention to is a tendency among philosophers to conflate *thinking* and materials *used* in thinking. We sometimes think 'in language', soliloquizing privately. On other occasions, we reflect non-linguistically. Some philosophers

distinguish these, describing the first as 'propositional' or 'sentential', the second as 'imagistic'. But *both* kinds of thinking are imagistic: we deploy visual, auditory, olfactory, tactile, and kinesthetic *images*. Some of this imagery is linguistic, verbal. Verbal imagery can be auditory (as when you 'hear' utterances in your head), kinesthetic (you 'feel' yourself uttering sentences), or a combination of these. Some cognizers can visualize inscriptions, mentally 'sign', and 'feel' embossed letters or Braille sequences. The point to appreciate is that verbal imagery is no less 'imagistic' than imagery of other sorts.[1]

It is easy to lose sight of the imagistic character of conscious thought when we start philosophizing about language and thought. We suppose that, although there is no problem at all as to how interior *utterances* could mean what they do, there *is* a special problem as to how pictorial images could have determinate meanings. Philosophers on all sides agree that meanings are not built into representations, sentences, for instance, or maps. Representations acquire significance by being put to use by representing agents. The apparent fact that pictorial representations lack 'intrinsic meaning' is simply a reflection of a wholly unsurprising fact about representations generally. Interior utterances—verbal images—are no less bereft of intrinsic meaning than images of sugarplums.

Regarded in this light, discussions of language and thought take on a new colouring. Thinking—conscious thinking—is not merely the having or entertaining of images, verbal or otherwise. Thinking is a matter of an agent's *using* such images, putting them to work. And whatever it is to put images to work, it is not solely a matter of entertaining further images. Nor is conscious thinking something occurring behind the scenes when you deploy representations: it *is* the deployment of those representations—in your head or otherwise.

Your consciously thinking that the cat is on the mat might be a matter of your saying to yourself, 'the cat is on the mat', or it might involve your conjuring a visual image of a cat on a mat. In either case, however, you do not do two things: (1) entertain an image; (2) think. Thinking occurs *in* your deploying the image. This point about use was made forcefully by Wittgenstein, and it has been endorsed by theorists as different as Locke and Andy Clark. Although philosophers

---

[1] The points emphasized in this paragraph were impressed on me in the 1980s by C. B. Martin.

routinely pay lip service to the idea, its consequences have not always been appreciated.

12.3 Davidson on Higher-Order Thought

One interesting question is whether some kinds of thought might be possible only for creatures possessing a natural language. Thoughts, or some thoughts about particular natural languages require proficiency in the pertinent language, but this is not what philosophers have in mind when they ask the question. Rather the question concerns whether certain kinds of thought about goings-on in the universe—thoughts of cats on mats, for instance, or thoughts about the thoughts of other thinkers—might be possible only for creatures who have mastered a natural language.

Consider thoughts about thoughts. Some philosophers have held that thoughts about thoughts—'higher-order thoughts'—would be thinkable only by linguistically endowed creatures. Donald Davidson goes further. Davidson contends that a capacity for thought per se requires a capacity for higher-order thought (see Davidson 1975, 1982; Heil 1992: chap. 6). This thesis, coupled with the contention that higher-order thought requires a capacity for language, implies that only creatures possessing a language harbour genuine thoughts. Davidson puts it this way: only a creature capable of interpreting (and so ascribing thoughts to) another creature is capable of any thoughts at all. Interpretation—which includes the mapping of sentences onto sentences—is impossible in the absence of language, however, so a creature lacking a language—a 'mute creature'—is perforce incapable of thought.

This is not an argument but the statement of a theory for which Davidson supplies elaborate support. I think it best to read Davidson, not as denying that mute creatures are capable of representing their surroundings mentally and acting on those representations, but as contending that there is a stark discontinuity between representations that might be deployed by such creatures and the kinds of representation found in agents equipped with a capacity for higher-order representation (see Bennett 1964). This is why, according to

Davidson, ascriptions of beliefs, desires, and intentions to mute creatures are at best figurative and at worst flatly misleading.

You can see what Davidson is driving at by noting that ascriptions of 'propositional attitudes' to mute creatures are 'semantically transparent'. Spot barks at an intruder lurking beneath the window. Which of the following is true:

(a) Spot believes there is an intruder lurking beneath the window.
(b) Spot believes there is someone beneath the window.
(c) Spot believes there is a noise beneath the window.
(d) Spot believes there is a noise *there*.

Such questions seem puzzling, according to Davidson, not because we lack ready access to Spot's inner life, but because, so long as Spot remains incapable of expressing his inner state linguistically, so long as Spot fails to be a candidate for *interpretation*, there *is* no fact of the matter as to which of these characterizations is correct. Genuine belief ascription, in contrast, is 'semantically opaque'. Substitution of co-referring terms in such ascriptions can alter their truth value. The belief that there is an intruder beneath the window differs from the belief that there is someone beneath the window. Thoughts about particular objects or states of affairs are thoughts about objects or states of affairs *under some description*. Attributions of representational states of mind to Spot, in contrast, identify those states of mind wholly by reference to objects on which they are directed. (In most cases, these will be the objects that elicited the states.) Because Spot lacks a language, there is no question of Spot's representing a given object, the intruder, say, under one description rather than another.[2]

Davidson encapsulates all this in a slogan: only a creature possessing the concept of belief could have beliefs. But that seems crazy! A creature could be furry without possessing the concept of being furry; a creature could fly without possessing the concept of flight. Why should the mere possession of beliefs require possession of the belief *concept*?

Davidson is thinking about the question this way. Start with the concept of belief, and consider what this concept includes. Beliefs can

be true or false. Possession of the concept of belief includes possession of a conception of something—a representation—that purports to fit the facts but could fail to do so. This is the concept of a truth *bearer*. Such a concept could be deployed only by a creature in a position to *ascribe* beliefs. Why? Consider your place in the universe. As you move through your surroundings, you come to represent those surroundings in various ways. If you are careful and attentive, changes in your surroundings will be mirrored by changes in your representations of those surroundings. You represent the water in a pond as deep. You subsequently step into the pond and discover it is shallow. You now represent the water as shallow. None of this requires you to distinguish what you *believe* to be the case from what *is* the case, none of this requires that you *deploy* a notion of truth.

A slightly more tendentious way to put this point is to note that a concept of truth comes into play only when you are in a position to put to work a distinction between truth *bearers*—representational states with definite, truth evaluable contents—and truth*makers*.[3] This amounts to the idea that your possession of a 'truth concept' requires possession of a capacity to represent representations and their contents, requires, that is, that you be capable of higher-order representation. What is striking about Davidson's argument is his contention that your deploying representations with definite contents *at all* requires your possessing a notion of truth, hence (if the reasoning above is sound) your being in a position to comprehend higher-order representations. This is something you do when you ascribe beliefs to others.

You would invoke higher-order representations in reflecting on your own beliefs, but Davidson's idea seems to be that in the absence of an opportunity to ascribe representational states of mind to others, agents are in no position—or perhaps have no reason—to ascribe such states of mind to themselves.[4] Given just the universe and your experiences of it, you have no use for a concept of representation, no grounds for a distinction between the universe as it is and the universe as you represent it as being.

---

2 This is not to say that there might not be more or less perspicuous ways of describing Spot's states of mind, only that pertinent factors fall well short of those constraining belief ascription to creatures endowed with a language.

3 As noted in chapter 8, the notion of a truthmaker and the notion of a truth bearer are correlative concepts: grasp of one requires grasp of the other. The notion of truth (as a relation between a truth bearer and a truthmaker) incorporates both.

4 A related argument can be found in Wittgenstein 1921: §§ 5.632–5.6331.

These are deep waters, but pretend for the moment that something like this is right: the notion of representation finds purchase only when you are in a position to ascribe representational states to others. Once the notion takes hold, you are in a position to extend it to your own case. This way of putting it is infelicitous, however. First- and third-person recognition go hand in hand. It would be misleading, then, to imagine that you reason from the third- to the first-person case: the relation, rather, is meant to be *reciprocal*.

Wait! This argument shows at most what is required in order for an agent to possess the *concept* of belief. What has any of it to do with our having *beliefs*, our deploying ordinary first-order representations of states of affairs?

Here the argument becomes, if possible, even denser. Imagine a solitary creature who represents the universe in particular ways. The creature's representations are dynamic: they 'update', reflecting changes in what the creature encounters. What are we to say about the contents of such a creature's representational states of mind? Return to Spot's thought about an intruder outside the window. As it happens, there are many ways to characterize the content of this thought. Davidson would say that there is no obvious reason to prefer any of these characterizations to any other: there is no fact of the matter as to how Spot represents his universe. You could grant that the creature is in a state that changes in ways that reflect changes in the universe, and you could grant that the creature's behaviour is guided by this state, without supposing that the state has any definite content. Its 'content', such as it is, can be characterized indifferently so long as it includes the pertinent object. Because the creature has no use for distinctions of the sort appearing in (a)–(d), the distinctions play no role in the creature's mental life.

Matters look different when higher-order representations appear on the scene. In *representing* a representation, you represent it, of necessity, as a representation possessing a definite content. What you have is another package deal: a creature capable of representing representations is a creature capable of ascribing representations—to other creatures and to itself—with definite contents, contents that could vary without the objects on which they are directed varying. And just as self-reflection is born of reflection on others, so contents—definite contents—of first-order representational

states of mind are born of higher-order representations of those states of mind.

Pretend this is right, or at least that it is not obviously wrong. The final move is to tie a capacity for language to the capacity for higher-order representation. Let us grant that an agent's entertaining higher-order representations requires that the agent be in a position to ascribe representations with definite contents to other agents who are *themselves* ascribers. How is this to be brought off, Davidson asks, without the kind of definiteness you find in language? Sentences (a)–(d) differ, and so are apt for the ascription of distinct representational states, states possessing determinate contents. The idea is that only a *linguistic* representation could be a representation of a representational state with a determinate content.

There you have it. The possession of representational states with determinate contents requires a capacity for higher-order representation. Higher-order representation requires possession of a language. So the possession of representational states with definite contents requires possession of a language.

You are likely to remain unmoved by all this, but my aim is not to convince you of the truth of Davidson's view, only to convince you that the view is not entirely mad. Many philosophers have been attracted to the idea that a capacity for thought, or a capacity to entertain particular *kinds* of thought, requires a capacity for language. The hope is that we might make progress in assessing this thesis by reflecting on one important class of thoughts, the class of thoughts about thoughts, higher-order thoughts, or, more generally, thoughts about representations *qua* truth bearers. Indeed, when we try to think of examples, it is easy to convince ourselves that a capacity for higher-order representation might well require a capacity for language. Philosophers have been attracted to such views for reasons that, when they are articulated at all, appear to have nothing to do with Davidson's arguments. What is important here is that Davidson does make explicit one line of reasoning to a popular, but unevenly defended, conclusion.

I shall argue that all this is off the mark, but before going on the offensive, it will be useful to look at another, independent line of argument that appeals to higher-order thoughts en route to the conclusion that certain prominent species of thought must be linguistic in character.

12.4 Bermúdez's Argument

The line of argument I have in mind is advanced by José Bermúdez in the course of an illuminating examination of 'thinking without words' (Bermúdez 2003). Bermúdez, unlike Davidson, accepts the idea that mute creatures are capable of entertaining a variety of interesting and sophisticated thoughts about actual and non-actual, merely possible, states of affairs. He draws the line, however, at second- (or, more generally, higher-) order thoughts. Bermúdez's contention is that thoughts on which higher-order thoughts are directed must themselves be linguistic. If, as seems obvious, only creatures who deploy a language are capable of linguistic thoughts, it would follow that only linguistically endowed creatures are capable of higher-order thought.

To see how Bermúdez's argument works, consider our capacity for what Andy Clark calls 'second-order cognitive dynamics':

Perhaps it is public language which is responsible for a complex of rather distinctive features of human thought viz, our ability to display *second order cognitive dynamics*. By second order cognitive dynamics I mean a cluster of powerful capacities involving self-evaluation, self-criticism and finely honed remedial responses. Examples would include: recognizing a flaw in our own plan or argument, and dedicating further cognitive efforts to fixing it; reflecting on the unreliability of our own initial judgements in certain types of situations and proceeding with special caution as a result; coming to see why we reached a particular conclusion by appreciating the logical transitions in our own thought; thinking about the conditions under which we think best and trying to bring them about. The list could be continued, but the pattern should be clear. In all these cases, we are effectively thinking about our own cognitive profiles or about specific thoughts. (Clark 1998: 177)

The characteristically human capacity for complex thoughts about thoughts could thus be tied to our capacity for language. This, according to Clark, is due, not to the distinctiveness of the *contents* of linguistic thoughts, but to the way they are embodied: as mental utterances or inscriptions of *sentences*. Imagine a particular state of affairs, the cat's being on the mat. This state of affairs might be represented pictorially or sententially. Your thought about the cat's

being on the mat could take the form of an image of a cat on a mat or an interior utterance of a sentence: 'The cat is on the mat.' In either case, I have suggested, you make use of imagery: the 'vehicle' of your thought might be a visual image of a cat on a mat or a verbal image, what Bermúdez calls an 'imaged sentence' (160).

Against this background, you can see Bermúdez's argument proceeding in two stages. First, he notes that thoughts on which second-order thoughts are directed ('target thoughts') 'must be at the personal level' (159), they must be kinds of thought concerning which you are or could be conscious. This apparently uncontroversial proviso is meant to set to one side Fodor's Language of Thought, sentences of which are not consciously available, hence evidently unsuitable as 'targets' for higher-order thoughts.

So, the Language of Thought aside, might there be kinds of thought possible only for creatures possessing a natural language? Trivially, thoughts about linguistic representations (that include a grasp of the significance of those representations) are possible only for creatures possessing a language.[5] The interesting question is whether there might be other kinds of thought that would require linguistic competence. Bermúdez's contention is that there are: Clark's 'second-order cognitive dynamics' would be possible only for creatures linguistically endowed.

This is stronger than the thesis defended by Clark. Clark argues that it is at best a contingent fact about human beings that much of our reasoning involves consideration of thoughts expressed sententially. Bermúdez, in contrast, holds that certain kinds of thought are flatly impossible in the absence of language.

There are certain types of problem that we solve by manipulating mental images and exercising the visual imagination. And we are, of course, conscious of bodily sensations, emotional feelings, and other such qualitative states (although these are not properly described as types of thinking at all). But we are not, I think, ever conscious of *propositional* thoughts that do not have linguistic vehicles. (160)

---

5 The parenthetical rider is important. A cat might entertain thoughts about your utterances without having the slightest idea what they meant or even that they *have* meaning. Your thoughts about sentences in an unfamiliar tongue might be no better in this regard than the cat's.

Next, Bermúdez advances an 'argument by elimination' to the conclusion that, in the case of such thoughts, 'the only available vehicles at the personal level are public language sentences' (159). Representational 'vehicles' could be of two sorts. On the one hand, 'representation might be secured symbolically through the complex symbols of a natural language' (160). In this case, 'a thought would be represented . . . through its linguistic expression and would appear as an object of thought qua linguistic entity' (160). On the other hand, 'representation might be secured in an analogue manner, through some kind of pictorial model' in which 'the vehicle of thought is a pictorial representation of the state of affairs being thought about' (160).

Although first-order thoughts can be 'vehicled' pictorially, such thoughts, Bermúdez contends, cannot serve as objects of second-order thoughts. This is a surprising and unexpected claim, given that this is exactly what anyone does who reads Bermúdez's colourful accounts of non-linguistic thoughts.[6] Consider your own case. In the right circumstances you can entertain a 'pictorial' thought of the cat's being on the mat and you can reflect—as you are now reflecting—on that thought. Bermúdez is after bigger game, however. Representation, he contends, evidently requires the holding of a 'structural isomorphism' between representation and represented state of affairs. There are, however, at least two senses in which a representation could be said to share a structure with what it represents.

In the weak sense there is structure whenever a structural isomorphism can be identified between the vehicle and what it represents. In the strong sense, . . . structure requires the existence of basic representational units combined according to independently identifiable combinatorial rules. Natural language sentences . . . are clearly structured in the strong sense, whereas mental maps/models possess structure only in the weak sense. (161)

Non-sentential representations—'mental maps and models'—represent what they represent by virtue of the holding of 'some combination of isomorphic resemblance and exemplification holding between the model/map as a whole and the represented state of affairs as a whole' (161). Bermúdez notes that it is easy to doubt that isomorphism could suffice for representation, but the real problem, he thinks, is in seeing how 'analogue' representations could serve as targets for particular kinds of higher-order thought.

Second-order cognitive dynamics involves a sensitivity to inferential relations between thoughts, and we do not yet have an understanding of how images can be inferentially connected to each other. The problem . . . derives from the intimate relation between inference and structure. . . . There is a sense in which mental models and maps are structured, since they contain elements that can feature in further mental models/maps. Nevertheless, they do not seem to be structured in the right sort of way to permit the reflexive type of second-order cognitive dynamics under discussion. (161)

Consider cognitive 'maps' of the sort discussed by David Braddon-Mitchell and Frank Jackson (1996), in particular consider a 'canonical example of second-order cognitive dynamics': recognizing 'the evidential basis for a particular belief and then evaluating the inferential transition made on that basis' (161).

It is perfectly easy to see how there could be some very basic forms of inferential transition between maps. Such transitions might be modeled on broadly associationist lines, and it is the possibility of such transitions that enables maps to serve as guides to action. What is not possible, however, is for such transitions to be understood and evaluated in terms of either deductive validity or probabilistic support. (162)

Such evaluations require the interpretation of the maps 'in broadly propositional terms'. Thus, 'we must interpret one map as expressing one proposition and the second as representing a further proposition, and then evaluate the inferential relations . . . between these two propositions'. The problem is that 'our only understanding of how to do this rests on the two propositions being linguistically formulated' (162).

Bermúdez's idea is that purely 'pictorial' imagery in general, and cognitive 'maps' in particular, could function as cognitive vehicles only when given an interpretation by a cognizer. The interpretation is fully linguistic. And in that case, it is the interpretation, not the item interpreted, that is cognitively efficacious.

---

6 Perhaps what Bermúdez has in mind is that, in reading about them, we do not represent non-linguistic representations but merely linguistic *descriptions* of non-linguistic representations. However, Bermúdez's discussion of non-linguistic representations purports to be a discussion of the representations themselves, and not merely descriptions of representations. In discussing such representations, Bermúdez's own thoughts must be directed on them.

Here we are meant to understand the vehicle of inference as a *syntactic* mechanism. P's implying Q is a semantic relation: P implies Q just in case P cannot be true if Q is false. This relation is *implemented* in an agent syntactically, a lesson you might have learned in mastering derivations in a first-order logic. Semantic relationships are mirrored by the syntax. From, 'If it's raining, the streets are wet', together with 'It's raining', it follows that 'The streets are wet'. The inference is captured by a sequence that would be familiar to anyone with a passing acquaintance with logic:

If p, then q

$$\frac{p}{q}$$

Symbolically,

$$\frac{p \supset q}{\frac{p}{q}}$$

An inferential move from premises to conclusion is captured syntactically. A logic student or a cognitive mechanism need only react to the syntax, to the shapes of the symbols, to the *structure* to move from premises to conclusion. In subtracting the syntax, you subtract the mechanism.

What of 'mental models' (see, e.g. Craik 1967; Johnson-Laird 1983, 1999; Rips 1994)? Mental models have been proposed as alternatives to syntactic models in accounts of reasoning. The idea is that, given an argument from premises to conclusion, agents 'construct' models of the premises and judge that the conclusion is implied when the conclusion is true in each of these models.

What goes for 'pictorial' imagery and mental maps, however, goes as well for mental models, according to Bermúdez. Mental models, he argues, are constructed on the basis of 'sententially encoded propositions' (163). The theory

construes . . . transactions between sententially encoded propositions as taking place in virtue of relations between analogue representations of the states of affairs portrayed in those sententially encoded propositions (as opposed to formal relations holding between the syntactic structures of the relevant sentences). (163)

In this regard,

mental models theory offers a particular way of developing the sentential conception, not of supplanting it. Mental models are not the vehicles of inference but rather, as their name suggests, models of those inferences. (163)

Bermúdez's contention is not that agents do not use mental maps and models, nor that agents cannot think about such things.

The point is that we cannot use mental maps or mental models for thinking about thoughts in the manner demanded by second-order cognitive dynamics. Natural language sentences are the only proxies that will permit thoughts to function as the objects of thought in this manner. (164)

This passage makes it clear that, despite suggestions to the contrary, Bermúdez's argument applies only to certain *kinds* of higher-order thought: thoughts pertaining to evidential and inferential relations among representations. Thus construed, the thesis is that we can make sense of evidential and inferential reasoning only when the relata are sentential in form.

In summarizing his conclusion, however, Bermúdez endorses a stronger thesis: 'there can be no intentional ascent without semantic ascent. We think about thoughts through thinking about the sentences through which those thoughts might be expressed' (164). This makes it sound as though higher-order thought, quite generally, requires a capacity for linguistic thought. In what follows I shall offer a challenge to those sympathetic to either a weak or a strong reading of Bermúdez's contention. The message is simple: where cognition is concerned, there is nothing special about language.[7]

12.5 Thoughts and their Expression

Return to Spot who, upon hearing the key in the lock, might naturally be said to believe his master is at the door. Why does Wittgenstein's rhetorical question—could Spot believe his master

---

7 I do not deny that, as a matter of empirical fact, human beings lean heavily on language. My interest is in the question whether what we accomplish by means of language could *only* be accomplished through the use of language.

will arrive the day after tomorrow—give us pause? If Spot can have thoughts about the here and now, why should we be reluctant to ascribe to Spot 'tensed' thoughts: thoughts about temporally distal occurrences? Or, to turn the question around, if we are reluctant to ascribe to Spot beliefs about occurrences the day after tomorrow, might our confidence that Spot could harbour beliefs about the here and now be misplaced? It is hard not to imagine that an answer to this question presumes an answer to the question whether thought is inevitably linguistic.

Here are two reasons you might think that thought requires language.

First, persuaded by Fodor, you might imagine that thoughts must themselves exhibit a definite syntax and lexical structure. To think is to engage in something like inner speech: the *medium* of thought is linguistic. Non-verbal imagery lacks the kind of determinate syntax required for genuine thought.

Second, you might want to connect the having or entertaining of thoughts with actual or possible *manifestations* of those thoughts. You might have verificationist motives here: 'an "inner process" stands in need of outward criteria' (Wittgenstein 1953, §580). Or you might simply take it to be part of the nature of a thought that it manifest itself in particular ways in particular circumstances. An unmanifestable thought would be one the possession of which made no difference to what its possessor does or might do. This, coupled with the idea that the only unambiguous manifestation of thoughts—or some thoughts—is linguistic, implies that only linguistically endowed creatures could have thoughts—or thoughts of certain kinds.

Consider this second reason more closely as it applies to Spot. The idea is that thoughts require manifestations (actual or possible), and only linguistic utterances could count as unambiguous manifestations of tensed thoughts. Spot could, the day after tomorrow, look expectantly at the door. This, however, would be most naturally taken to be a manifestation of Spot's belief that his master *is* at the door, not a manifestation of a tensed belief.

Perhaps, however, the problem is not that Spot lacks a capacity for the linguistic manifestation of tensed beliefs, but that spot lacks a *use* for tensed representations. One sort of use of such representations

would be the conveying of tensed information to others. This is the kind of use manifested in utterances of tensed sentences.

The suggestion on the table is that the idea that Spot is linguistically challenged hence unable to entertain thoughts about temporally remote states of affairs is a red herring. To be sure, Spot lacks a linguistic vehicle of expression. But it is not this that makes it unlikely that Spot could entertain tensed thoughts. Spot lacks a *use* for tensed representations, hence a capacity for tensed thinking.[8] One prominent sort of use might be communicative, a fact that points to language. What is doing the work here, however, is not language, but *use*.

You might be sceptical that there *could* be anything like a tensed representation in the absence of language. If you thought that, you might accept what I have said about use, but note that, in many cases, including Spot's, only linguistic use is available. Maybe. But maybe not.

12.6 Thought

Suppose you take seriously the venerable idea that language is a tool. Intelligent agents employ language as you might employ a screwdriver to achieve various ends. If language is a tool, it is a multipurpose tool, a Swiss Army knife. (To be sure, even a 'single-purpose' tool—a screwdriver, for instance—can have multiple uses.) If this is right, it would be a mistake to regard thoughts as being *in* language. Thinking is linguistic when you think *with* language. If a linguistic episode—an inner utterance, for instance—is analogous to a screwdriver, thinking is analogous to driving a screw with a screwdriver. Regarded in this light, the question whether thought, or some thought, requires language, reduces to the question whether there might be tasks that could *only* be performed by using language. And just as it seems unlikely that *any* tool is irreplaceable, so it seems unlikely that language is, for any particular task, irreplaceable.

What exactly is it to deploy language as a tool? And what *other* intellectual tools might be available to intelligent agents? Just as you

8 This, to be sure, is an empirical claim, one that could turn out to be false.

make use of utterances—written or spoken—to articulate ideas to others, you can use *inner* utterances in the articulation of ideas to yourself. You can talk through a problem, recall the details of an earlier conversation, or plan a course of action by listing steps to its completion in your head. In these cases, inner utterances are not manifestations or copies of thoughts; you are thinking with language just as you might open a can with a can opener. There is no mystery here any more than there is a mystery in your spontaneously engaging in intelligent conversation. You need not first think—rehearse—then speak. This is so whether speech is overt or covert.

Inner utterances (I say, siding with Bermúdez) are a species of mental imagery, where the images are images of what their audible, visual, or tactile counterparts sound, look, or feel like. There is no logical or conceptual gulf between linguistic ('propositional') imagery and imagery of other sorts, 'pictorial' imagery. Conscious thought quite generally is imagistic.[9]

Not all thoughts incorporate *linguistic* imagery, however. Much of our thought involves non-linguistic visual, auditory, tactile, or olfactory imagery. Indeed, your thought about a particular person might include verbal imagery (an inner utterance of a name, for instance) accompanied by a visual image of the person and perhaps other imagery as well.

The association of imagery with thought is not a matter of identifying thoughts with images. Thinking is a matter of *using* imagery. Bare images, what Martin (1987) calls the *materials* of thought, are devoid of intrinsic significance. You can grant this, grant that images themselves lack 'intrinsic intentionality', without thereby giving up on the idea that conscious thinking is a matter of deploying images. Interior utterances are, in this regard, on a par with 'pictorial' imagery. Intentionality, significant thought, arises when images (or signs, or representations generally) are put to work by intelligent agents. Without use, images or signs are empty; severed from use, representations fail to represent.

---

9 This, it would seem, is what lies behind the venerable debate concerning the possibility of 'imageless thought'. See, for instance, Angell (1897), Woodworth (1906), Danzinger (1980), Thomas (1999, 2001).

It is easy to lose sight of this simple idea so long as we persist in the illusion that interior utterances are not themselves images altogether on a par with images of sunsets and romantic Greek isles. Debates over whether imagery is 'pictorial' or 'propositional' are beside the point (see Kosslyn and Pomerantz 1977; Pylyshyn 1981, 1984; Rollins 1989; Tye 1991). Linguistic thoughts are as 'pictorial' as any others.

At various points in previous chapters I have made disparaging remarks about propositions. What do I have against propositions? Philosophers posit propositions as 'abstract entities' in the service of various philosophical projects. Some find the abstract nature, the ontology, of propositions objectionable. Even if you were sanguine about such things, however, you might find propositions dodgy for a completely different reason.

Most philosophers are happy to accept the idea addressed in this section, an idea defended with equal verve by philosophers as different as Locke and Wittgenstein, the idea that meaningful signs owe their significance, not to intrinsic features of those signs, but to the use to which they are put by intelligent agents. But what of propositions? A proposition is an entity that is *itself* essentially meaningful, an entity with *built-in* meaning. Propositions wear their meanings on their sleeves, a proposition could not fail to mean what it means. This makes propositions like no temporal entity. Their dodgyness lies not in their being *abstracta*, but in their being intrinsically meaningful *abstracta*.

It is easy to doubt that there are or could be such entities. The wonder is that philosophers whose theories would make these kinds of entity impossible are happy to appeal to them to advance their projects.

12.7 Images, Sentential and Otherwise

Suppose I am right: conscious thinking is inevitably imagistic; to entertain a thought consciously is to deploy images of one sort or another. Imagery can be 'pictorial' or 'sentential'. You can imagine how something looks (did look, will look, or might look), feels (did feel, will feel, or might feel), tastes (did, will, or might taste), sounds

(did, will, or might sound), or smells (did, will, or might smell). One species of such imagining is verbal: you imaginatively utter, or hear, or feel yourself uttering, words.

In considering such cases, we philosophers are apt to be misled in at least two ways. First, we are likely to regard verbal imagery as privileged. On the one hand, there is ordinary 'pictorial' imagery, on the other hand, there is the entertaining of 'sentential' thoughts. This is the kind of distinction Peter Carruthers (1996) appears to have in mind in distinguishing 'imagism' and 'sententialism'. An 'imagist' regards thoughts as invariably pictorial. 'Imagists' hold that 'thoughts inherit their representational properties from the representative powers of the images that constitute them' (1996, 32).

Carruthers mentions Locke, Hume, and Russell as proponents of imagism. 'Imagists', he argues, hold that thinking consists of a succession of picture-like ideas parading through the mind. 'Imagism' could be contrasted with 'sententialism', the view that thoughts, or some thoughts, are linguistic in character. 'Tokens' of such thoughts are mental sentences, which are to be distinguished sharply from 'pictorial' images. But what are mental sentences if not images: verbal images? The contrast here, if there is one, is not between images and non-images, but between two kinds of image: one non-sentential, one sentential. There is no sense in which verbal images—mental 'sentence tokens', interior utterances or inscriptions—are less imagistic than nightmares or the remembered taste of a *madeleine*.

A second way in which we philosophers are likely to be misled in thinking about verbal and non-verbal imagery is more pernicious. Consider a passage from Carruthers.

An image will always carry with it *excess content* beyond mere entailment relations. The sentence, 'A cat is on a mat' carries no more, and no less, content than the proposition [that a cat is on a mat]. So it will contain the content, [that a mammal is on a mat], since this is entailed, but it will not contain the content, [that a tabby cat is on a mat], nor the content, [that a cat is sitting on a mat]. Not so with an image. Although an image can be in various respects indeterminate or vague, it must always be *more* determinate than a proposition. My image of a cat on a mat must always be an image of some particular kind of cat (fluffy or short-haired, tabby or black) in some particular position (sitting or standing, facing or turned away) on some particular sort of mat. (1996: 38; brackets are in the original)

Carruthers's comment echoes Berkeley's criticism of Locke's doctrine of 'abstract general ideas' (see Locke 1690: III, iii, 6; Berkeley 1710: *Introduction*, §9). Locke's doctrine of abstract ideas is not a thesis about kinds of idea, however. Abstraction is, for Locke, 'partial consideration' (see, for instance, Locke 1690: II, viii, 13). An abstract idea is not a generic idea, whatever that might be, but an idea *used* generically. In just the same way, you can use a particular, determinately hued rectangular red paint chip to represent redness-in-general by considering it 'partially'—just in so far as it is red.

Berkeley's own account of general ideas begins with the observation that images inevitably exhibit a kind of definiteness (Carruthers's 'excess content') lacked by linguistic representations. A pictorial representation of a cat on a mat will be a representation of a cat sitting or lying on a mat, for instance. This, Carruthers thinks, makes it ill suited as a vehicle for the thought that the cat is on the mat. But Berkeley—and Locke—note what Carruthers does not: generality is a product of our *deployment* of images. An image of a reclining tabby could be used to represent a cat's being on a mat and nothing more, to represent a cat's being on a mat, for instance, rather than on the sofa. Think of road signs depicting an automobile skidding on a wet road. The image is of an automobile of a particular shape, but sign makers count on our grasping the image as representing motor vehicles generally.

The point applies across the board: the key to understanding the nature of thought is the recognition that thinking is something agents do *with* imagery. Thinking is not the having or entertaining of images, the mere *occurrence* of imagery, sentential or otherwise. Thinking is something *done* imagistically.

12.8   Proto-Language

Philosophizing is a largely verbal enterprise. It is scarcely surprising, then, that when we engage in philosophical reflection—when we think about philosophical topics—we do so in a manifestly verbal idiom. There is nothing special about words, however, nothing magical about utterances and inscriptions. Words take on significance when they are put to use by intelligent agents or intelligent systems.

If it is doubtful that Spot could entertain the thought that his master will arrive home the day after tomorrow, this is not because Spot lacks words to express this thought. It is because it is hard to see how Spot could have a use for a representation with this content.

Here it might be useful to consider what Martin (1987) calls 'proto-language':

Nonlinguistic activity at its more sophisticated and structured levels has a remarkable pattern of parallels to that of linguistic activity. It is a matter of degree, but when an agent shows enough of this pattern of parallels, this structured network of procedures can be called 'proto-language'. (1987: 277)

'Proto-language' is not 'a kind of language as sign language *is* a kind of language' (278). Rather it is 'a structured rule-governed network of semantic, procedural activity prior to and basic to linguistic activity, having an almost totally unnoticed and surprising pattern of parallels to language itself' (278). Martin's idea is that semantic features of natural languages can have non-linguistic, 'procedural' analogues in assorted non-linguistic activities—overt and covert—of intelligent agents. It would follow that reference, quantification, modality, tense, and the like are language-independent. The challenge is to find roles for such things in the activities of creatures lacking a means of expressing them linguistically.

You can see what Martin is driving at by considering an example. Counterfactual and subjunctive conditionals might be taken *obviously* to depend on a distinctively linguistic medium. The difficulty is to identify a non-linguistic procedure aimed at discovering some counterfactual or subjunctive truth. Imagine an early hominid who has learned that what fish eat can be found in their stomachs.

He has also noticed that they eat different things at different times. When he catches a fish, he opens its stomach to see what it has been feeding on, so that he can use it as bait.

On one unsuccessful day's fishing he notices an approaching storm that looks like spoiling the fishing for a long time. Frustrated, he intends not to return to the fishing hole until the weather changes. He picks up his fishing gear and starts for the cave. He happens to frighten a mink eating a fish. His curiosity overtaking him, he opens the stomach to see what the fish had eaten and takes out some grasshoppers. This is a procedural action whose

projected outcome is information about the past. It *also* has the point of finding out what *would* have helped him to catch fish if he had used it as bait. (287)

In the case described, you have a procedure aimed at establishing a particular outcome, an outcome you could express linguistically by means of a tensed counterfactual conditional. If it is hard to imagine Spot entertaining thoughts with tensed or subjunctive conditional contents, this is not because Spot lacks a capacity to express such thoughts linguistically, but because it is hard to credit Spot with a capacity to engage in *any* sort of procedure the aim of which is to establish a tensed or subjunctive conditional truth. Spot's deficit is cognitive not linguistic.

This is not to deny that cognitive and linguistic capacities might be importantly linked in terrestrial species (see Clark 1998). The point, rather, is that a capacity for the production of utterances or inscriptions cannot, by itself, amount to a capacity for sophisticated thought. We knew this already: mynah birds and parrots can be taught to produce utterances.[10] What they (apparently) lack is a capacity to put those utterances to intelligent use.

12.9 Non-Conscious Thought

Suppose, then, that ordinary conscious thought is best understood as the manipulation of images *for* various purposes. Not all thought is conscious thought, however. Thinking apparently occurs behind the scenes outside our awareness. Can the imagistic conception I am suggesting be extended to cases of non-conscious or 'subpersonal' thought?

The idea that non-conscious thought might be imagistic is not something I intend to pursue here. I do not intend to pursue it because I believe there is a better way of understanding the nature of at least *some* non-conscious thought. A simple example will set the stage.

10 Perhaps such creatures produce, as well, silent imagistic counterparts of these utterances!

If you are like me, you can sometimes resolve nagging intellectual or practical problems by sleeping on them. You go to bed puzzled, wake up enlightened. Something similar can happen when you are trying to remember a forgotten name: you turn your thoughts elsewhere and discover that the name subsequently 'comes to you'. How should we understand such cases?

One possibility is that something is going on backstage altogether analogous to what goes on when we consciously reflect. This is the shoemaker-and-elf model: when the shoemaker retires for the night, elves materialize, performing just as the shoemaker would, but secretly.

Another possibility is that our unconscious mental life differs structurally and qualitatively from its conscious counterpart. Think of the brain as a complex, dynamic dispositional system. The system manifests itself in myriad ways: in bodily behaviour, in speech, and, often enough, in conscious thought. The system manifests itself as well in the regulation of various bodily functions, blood flow, digestion, thermal regulation, the maintenance of equilibrium, and the like. The finely-tuned, focused dispositions constituting the nervous system are not static but dynamic: they evolve continuously over time. The system that, a few minutes ago, was unable to produce the name of an acquaintance, has resettled itself and can now, in concert with other, reciprocal, dispositional systems, yield the name.

In many cases—in cases of 'unconscious inference', for instance—this could be what is going on. Were that so, non-conscious thinking would not be a matter of something like conscious thought going on behind the scenes, thinking minus its imagistic trappings. Rather, a different kind of process is occurring, a process that produces an alteration of the dispositional *basis* of thought.

12.10 Language in Its Place

What is the upshot? First, thinking—ordinary conscious thinking—is imagistic. This is so for 'linguistic', or 'sentential' thoughts as well as for patently non-linguistic thoughts. Second, thinking does not consist merely in having or entertaining images, but in deploying those images, putting them to work in the service of various ends. Third,

verbal and non-verbal imagery are representationally on a par. The significance of *any* sort of representation lies in the use to which it is put by intelligent representers. Finally, what we regard as non-conscious thought need not resemble conscious thought occurring out of sight. Non-conscious thought could turn out to be the dynamic unfolding of a dispositional base, one typical manifestation of which is conscious thinking.

Creatures lacking a capacity for language could well be unable to entertain thoughts of various kinds. But this might be so, if it is so, not because thoughts or some kinds of thought require linguistic vehicles. An inability to use language and an inability to think certain thoughts could, and most likely does, have a common source. A creature could lack the kind of make-up required for the pursuit of projects the satisfaction of which would afford opportunities for the deployment of representations of particular sorts. This puts language in its place.

CHAPTER 13
The Ontological Turn

For us in America today the contrast between the high-hearted metaphysics of naturalism and all the fine evasions of obscurantism and agnosticism may be literally of epochal importance. The culture of America, by reason of its unique provenance, may choose either to be old or to be young, to be Alexandrian or to be Milesian. Whether we are thus at the end of a career or the beginning of one, will in large part depend upon whether our citizens in this century learn their lessons from mystic evangels who would purge us of scientific understanding, from resigned sophisticates who set up languages and toy with thoughts of future possible sensations, or from philosophers who explore the nature of things.

(Williams 1944, 442–3/238)

The real constitution of things is accustomed to hide itself.

(Heraclitus, frag. 123, in Kirk and Raven 1964, 193)

13.1 Ontological Convergence

Previous chapters include numerous reminders of what I consider to be pitfalls inherent in linguisticized metaphysics. Now, I hold out a peace offering meant to alleviate potential worries that the ontological picture sketched here eviscerates the special sciences and undermines beliefs about cherished everyday objects. The goal has been to *account* for such things, to say what they *are*, not to cast doubt on their authenticity. In this context, it is worth bearing in mind that differences among anti-reductionist philosophers—and I count myself as anti-reductionist—are easily exaggerated.

Suppose I am right in thinking that arrangements of the fundamental things of the kind revealed by fundamental physics serve as truthmakers for all the truths that have truthmakers. This, I insist, would not cast doubt on tables, trees, and planets. Rather it points to the deep story about the nature of such things. One alternative would be to imagine that tables, trees, and planets occupy 'higher levels'. But the fact that talk of tables, trees, and planets cannot be reduced to or analysed in terms of categories of fundamental physics is best explained, not by positing intervening levels of entities, but by recognizing that categories we deploy to cope with the universe exhibit a diversity that resists regimentation. Psychological taxonomies and categories are orthogonal to biological taxonomies and categories, which in turn are orthogonal to physical and chemical taxonomies and categories.

Reduction is best understood as a relation among categories, taxonomies, theories, or terms, not a relation among entities or kinds of entity. Showing that one category is reducible to or replaceable by another is to show that the category is expendable. To say that there are tables, trees, and planets, is to say that the terms we use to identify tables, trees, and planets can serve to express truths. This is perfectly consistent with the idea that the deep story concerning the nature of such things is to be found in fundamental physics.

It might appear that I am recommending replacement of biology, or psychology, or anthropology by fundamental physics, but this is not the aim. Take biology as an example of a special, 'higher-level' science. Biologists operate with biological categories and taxonomies, illuminating features of the universe that fall under those categories and taxonomies. The legitimacy of biology in no way depends on a derivation of biological truths from truths expressed in a vocabulary of fundamental physics. In this regard, biology is an autonomous science.

This is what anti-reductionist arguments show, and this is *all* they show. Whatever their proponents might think, anti-reductionist arguments do not establish that, *in addition to* the substances and properties postulated by fundamental physics and arrangements of these, there are biological (or psychological, or anthropological) substances and properties. When you put the fundamental things together in the right ways, you assemble truthmakers for the biological, psychological, and anthropological truths.

Pretend for a moment that this tomato is a particular arrangement of particles. By this I mean only that, were you to look closely at the tomato, were you to subject the tomato to minute scientific scrutiny, this is what you would discover. The particles are interrelated in important ways, causally and otherwise. The particles have definite causal histories, and stand in definite relations to particles outside the tomato. These particles being in this configuration, *is* the tomato's being red, *is* the tomato's weighing six ounces, *is* the tomato's being disposed to roll. Do not imagine that, if these truths hold for the tomato, they hold of something distinct from, something in addition to, these particles in these relations interacting in these ways.

I am happy to say that the particles compose or make up the tomato, that the 'is' in 'the tomato is a collection of interrelated particles' is the 'is' of composition. Truths about the tomato are made true by this collection. The tomato's capacities and qualities are what you get when the tomato is made up of these particles interrelated as they are (with their various histories and relations to particles elsewhere). The view I am opposing is the view that the tomato's qualities and powers are somehow dependent on, but distinct from, the powers and qualities of its constituents with their assorted relations. I oppose this kind of view for three reasons.

First, it is gratuitous. You do not *need* to invoke levels of properties and property bearing entities to account for the success of the special sciences and the character of the manifest image. The irreducibility of categories and taxonomies can be explained without resorting to a hierarchy of levels of being.

Second, the levels view creates more problems than it solves. Suppose, for instance, you thought that the tomato and its properties were higher-level entities. And suppose you thought that entities' scientific legitimacy depended on their capacity for influencing other entities causally. In that case, you would want the tomato and its properties to have distinctive causal profiles. So now you would have powers of the tomato and its properties *and* powers of the interrelated collection of particles making up the tomato: the tomato's 'realizers'. This would seem to portend, at best, an unnecessary duplication of powers and, at worst, a pre-emption of the tomato's powers by those of the collection: the tomato and its properties would be epiphenomenal.

Assuming that observation includes a causal component, the thought that tomatoes and their properties are distinct, higher-level entities, and that these entities are epiphenomenal would have the consequence that tomatoes are unobservable. The manifest image would not be manifest! You might try to rescue higher-level items by introducing levels of causation corresponding to levels of entities, but higher-level causation looks deeply problematic. I hope that I have said enough in earlier chapters to convince you that a bad way to deal with this problem would be to move the goalposts by massaging the concept of causation.

Third, and most importantly, the view I am recommending is independently attractive. The view supports anti-reductive convictions without introducing entities that then threaten to undermine the very enterprise they were meant to preserve. Truthmakers for claims about the tomato are the particles duly arranged and interrelated to one another and to other things. You can accept these truthmakers without embracing reduction, without imagining that you could derive truths about tomatoes from truths about the particles.

I recognize that there remains room for misunderstanding here, but I have the sense that the view I am endeavouring to articulate is consistent with positions that might at first glance appear sharply at odds with it. These assurances will not do much for you if you are already firmly convinced that tomatoes and their properties are distinct from, but dependent on—maybe 'realized by'—collections of particles that make them up, if you think that tomatoes and their properties exhibit causal powers not possessed by the particles arranged as they are.

My hope is that the ontology developed here provides a kind of vaccine against layered views of this kind. The ontology does so, *if* it does so, not by yielding point-by-point criticisms of those views, but by offering an alternative intended to stand on its own.

13.2 Bottom–Up Ontology

Now, at the risk of losing readers who have figured out what I have to say and how to respond to it, I shall use the remainder of this

chapter to summarize a few central tenets of the proposed ontology. Readers for whom this exercise is superfluous are encouraged to skip to § 13.3.

The aim of the book has been to advance an ontology meant to accommodate the most general features of the universe as we find it and to illuminate the relation between the manifest image we absorb from everyday experience, from our culture, and from the special sciences, and the scientific image as we have it in physics. Physics, fundamental physics, is in the business of providing an account of the truthmakers for all the truths made true by the universe, including truths issuing from the manifest image.

Fundamental physics is inevitably a work in progress. I make no assumptions as to how the universe ultimately is, what form physics will take in the future. A couple of broad possibilities spring to mind. On the one hand, the universe might be quantized, discrete, corpuscular. In that case, ordinary things—common-sense objects, items discussed in the special sciences—would be dynamic, interrelated arrangements of corpuscles. On the other hand, the universe might be continuous. The universe might be a field, or a small number of fields, pervading space–time; or the universe might be space–time itself. Were that so, objects of common sense and the special sciences would be properties of a continuous substance or the intersection of several continuous, overlapping substances. These possibilities need not be thought to be either exhaustive or exclusive. The universe might be neither quantized nor continuous. The universe might be continuous at some times or in some places, quantized at other times or in other places.

Two features of this picture stand out. First, whatever the nature of the fundamental truthmaker or truthmakers, the ontology will be one of substance and property. The only question is whether the substances are continuous, or corpuscular, or something else. Second, although fundamental physics is pre-eminent whenever our goal is to uncover the deep story concerning the fundamental truthmakers, when it comes to the inventory of worldly truths, physics enjoys no privileged status. Ordinary, everyday assessments of the universe and assessments originating in the special sciences are or can be true, and no less true than those stemming from work in fundamental physics. The project is in no sense reductive.

I can hear a small voice in the background: 'the ontology *is* reductive; the real existents are those posited by physics, the rest are mere appearances'.

This is a misleading way of describing the ontology, however. In the first place, reduction is most naturally regarded as a relation among theories, taxonomies, categories, not a relation among entities or kinds of entity. How would you reduce one *thing* to another? As I see it, categories of common sense, and those deployed in the special sciences, are largely orthogonal to categories at home in fundamental physics. All carve, or purport to carve, nature at the joints. All figure, or could figure, in true judgements about the universe. The idea that we might dispense with biological or psychological categories and descriptions, replacing them with categories and descriptions borrowed from physics is hopeless. Categories we use to describe and explain our universe do not line up in the neat way reductive schemes require.

I admit that I have no unassailable argument for this assertion, but I suspect it would be accepted even by my most ardent critics. The failure of categories to align reflects our experience and honours the appearances. It is just what you would expect given a complex universe with endless joints and real divisions, a universe that admits of innumerable and incommensurable was of counting. These are all points inspired by one of the heroes of this enterprise, Locke.

The deep story about the universe's ontology is to be found in fundamental physics, but this does not augur the replacement of ontology by physics. Ontology serves at least two indispensable functions. First, it provides an accounting of the basic ontological categories. On the view defended here, these categories include substance and property. Truthmakers for all the worldly truths will be propertied substances. It is more natural, perhaps, to think of the truthmakers as being particular interrelated *arrangements* of propertied substances, but I believe (perhaps overly optimistically) that truthmakers for relational truths will turn out to be non-relational features of the universe.

A second indispensable role for ontology is that of providing an account of relations among the various sciences, a role no science, including fundamental physics, is equipped to play. To be sure, philosophers are not the only philosophers. Nothing prevents

scientists from engaging in metaphysical speculation when they are so moved. Still, it is philosophers, primarily, who have self-consciously addressed questions about 'reducibility' in the sciences and the status of the manifest image.

This makes ontology sound cut and dried. There are propertied substances and perhaps arrangements of these, and there are different ways of marking off the endless divisions among the substances. Is that *it*?

That might be it, were there widespread agreement over the nature of substances and properties. Many philosophers would join me in endorsing the primacy of substance and property, but would want to embrace markedly different conceptions of these categories. For this reason, four of the book's chapters have been devoted to discussions of the nature of substances and properties.

Substance and property are correlative, reciprocal categories. Substances are property bearers, properties are particular ways substances are. To be a property, to be a way, is to be a way some substance is; to be a substance is to be some way or other. Substances and properties alike are abstractions, not in the sense that they reside somehow independently of space and time, but in the traditional sense: they are by nature partial, separable only in thought. You can consider a substance as a property bearer, and you can consider ways the substance is. In so doing, you are engaging in Locke's 'partial consideration'.

I have introduced substances as property bearers. This leaves open various interesting questions about properties. Prominent among these is the question whether properties are *universals*. I follow Locke (and many, many others) in regarding properties as *modes*: particular ways particular substances are. If you thought of the redness of *this* tomato as a property of a particular tomato, and the redness of *that* tomato as a property of a distinct tomato, these would be distinct properties, distinct modes. If you thought of properties as universals, the two tomatoes, or perhaps their two rednesses, would be distinct *instances* of a single universal.

Some defenders of universals, Armstrong, for instance, regard universals as *immanent*. A universal is said to be wholly present 'in' each of its many instances. The two rednesses are 'numerically identical', the selfsame entity. This nicely accommodates the thought that the two tomatoes are *the same* colour, that they *share* a colour. If you

are like me, however, you will have trouble understanding how anything *could* be multiply locatable, how anything *could* be wholly present in distinct places at the same time. I have not dwelled on this point because, whatever its merits, I am convinced that purported advantages of immanent universals are illusory. That, coupled with the thought that an ontology that includes general entities, alongside particulars, is profoundly mysterious, leads me to prefer modes.

A willingness to describe distinct objects as *sharing* a property or as having *the same* property provides only the flimsiest of reasons to accept universals, immanent or otherwise. Your eating the same breakfast every morning is a matter of your eating similar breakfasts, your sharing a dislike for Brussels sprouts with your dining companion is a matter of you and your companion's having similar culinary preferences. 'The same' can mean 'selfsame', 'one and the same', but it can also mean 'similar' or 'exactly similar'.

Linguistic evidence aside, proponents of immanent universals argue that taking properties of objects to be universals enables us to explain laws of nature. Laws connect universals. This is why one law can cover many distinct cases. If properties are particulars you lose this advantage. 'Same cause, same effect', becomes miraculous, coincidental. If this tomato rolls because of its sphericality, what would entitle you to expect another tomato to roll because if *its*—'numerically distinct'—sphericality?

Would the particularity of objects and their properties defeat the kind of generality sought in the sciences? Why should it? If two tomatoes are similar with respect to their shapes—this one is spherical and that one is spherical—why *shouldn't* you expect the one to roll if the other does? What advantages does identity enjoy over similarity? This thought, coupled with the difficulty of understanding how an entity could manage to be wholly present in many distinct places at once, tells in favour of the view that properties are modes.

What if universals were 'Platonic'? What if universals resided in a 'Platonic heaven' outside space–time. Their instances *are* distinct, so the puzzle of multiple locatability does not arise. Would such a conception of universals afford the advantages of immanent universals without the embarrassments?

What advantages? Given that our interest is in ontology, universals of any stripe enjoy no advantages over modes. More particularly,

when you take universals to be 'Platonic', you face additional problems. If universals subsist 'outside' space–time, how could they make a difference *within* space–time? They would, it appears, need local representatives. If the tomato rolls because it is spherical, something here-and-now about the tomato, the tomato's here-and-now sphericality, is responsible for its rolling. Plato had the forms, but he had, as well, the 'moving forms', proxies for the forms in the mundane universe. So you have universals and modes, or mode-equivalents. If the modes are doing the work, why credit the universals?

For the record, I do not believe that Plato was a Platonist, at least not a Platonist in the sense in which this label is often deployed. 'Plato's heaven' was never supposed to be a realm of objects residing independently of the spatio-temporal universe. Plato's forms are not *abstracta* in the sense in which philosophers today use this term. Plato's forms are concrete and, I reckon, meant to be causally involved in day-to-day operations of the universe.

This is but a brief reminder of just how much recent philosophical themes have strayed from tradition. Throwing off tradition can be exhilarating. The thought that we should simply footnote positions embraced by our philosophical ancestors is unappealing and ultimately self-defeating. But in our haste to move ahead, it is easy to lose our bearings. Pity the contemporary philosopher who thinks that Locke's or Spinoza's ontology is scientifically naïve or conceptually laughable, while straight-facedly embracing alternative worlds and realms of sets, numbers, universals, and propositions 'present' non-spatio-temporally in every world.

Armstrong and a handful of others aside, why are so many philosophers today attracted to the idea that properties are, or must be, universals? I can see no good reason. As far as I can tell, today's widespread acceptance of the idea that properties are universals stems from intellectual inertia: for whatever reason, the 'standard view', the default view, the view that needs no defence, is that properties are universals. In some circles 'property' *means* universal. If you think otherwise, the burden is on you to establish that properties might *not* be universals, an apparent contradiction.

If I am right, all the standard view has going for it is that it *is* the standard view, and even this is doubtful when you take up a broader perspective. If you look at the history of philosophers' discussion of

properties, you will find that, far from being standard, the view that properties are universals is in fact exceptional. We are prevented from seeing this in part because we tend to read the history of philosophy through our own lenses, we interpret our past, as we are bound to, in terms of our present.

I am aware that these assertions are controversial. Nevertheless I stand by them, convinced that many of the things earlier generations of philosophers have said about properties are best understood as attempts to say that truthmakers for judgements about properties are fully particular ways the universe is. This is consistent with the idea that you can regard any property, any way a substance is, as a way something else could be and thereby regard it as a universal. D. C. Williams puts it this way:

> Our description of physical realism patently assumes the reality of universals or essences, at least in their instances, just as any description does, but it need not be interpreted as hypostatizing them any more than does a child's remark that the ball is round and that the orange is round too. (1944, 428/223)

These universals, far from belonging to a special class of non-particular, general entities, are particular ways particular things are, modes.

Properties, then, are modes. I prefer 'mode' to 'trope' because 'mode' is the traditional designation for what I have in mind, and because philosophers who approve of tropes are most often 'bundle theorists'. According to the bundle theory, substances are assemblages of 'compresent' tropes; substances are constructed from tropes; tropes are substances' component parts. This yields an austere one-category ontology of the kind favoured by Williams, Keith Campbell, and David Robb.

When you look at the history of trope bundle theories in the twentieth century, you see that they have two principal sources. The first resides in the idea that the chief function of substance is to introduce particularity into a universe that would lack it otherwise. Universals need to be associated with substances to yield 'concrete' objects and states of affairs. Once you give up universals, once you accept properties as themselves particulars, *this* function of substance goes by the board.

A second source of the rejection of substance is the conviction that substances are obscure posits, unobserved and unobservable entities concerning which we could have no clear conception.

The lack of a need for substances as particularizers coupled with the opinion that substances are inherently mysterious we-know-not-whats, naked particulars that, once you have dispensed with universals, are ontologically superfluous, begets the bundle theory. Bundle theories lack obvious historical precedent. Hume is sometimes described as a bundle theorist, but it is always difficult to be sure when Hume is pursuing an epistemological point, and when to interpret his claims as metaphysical pronouncements. Berkeley treats material objects as bundles of ideas, but ideas are modes of immaterial substances, spirits. Hume and Berkeley aside, it is not easy to find bundle theorists among earlier philosophers who regarded properties as particulars. On the contrary, philosophers who took properties as modes— Descartes and Locke, for instance—considered the thought that modes could exist without being modes of some substance as unintelligible on the face of it.

Scholastic philosophers endeavouring to cope with the metaphysics of transubstantiation come close to allowing that objects might, under extraordinary circumstances, amount to bundles of properties. Some, if not all, properties of Eucharistic bread and wine could survive the miraculous subtraction of the substance to which they originally belonged. These properties, the so-called *real accidents*, were potentially embarrassing oddities and, outside ecclesiastical circles, decidedly unpopular. Descartes, for instance, expended considerable energy attempting to work out a metaphysically respectable story about transubstantiation that dispensed with real accidents.

Transubstantiation and the Eucharist aside, opponents of substance promote the idea that accepting substances poses a dilemma: either

(1) substances lack properties, substances are bare somethings-we-know-not-what, bizarre, unobservable entities with no shape, no size, no mass; or

(2) substances themselves *do* have properties, in which case positing substances to explain property bearing leads to a regress of bearers.

No friend of substance ought to be moved by this dilemma. Substances are not hidden beneath, or masked by, their properties. To encounter a substance is to encounter something that *is* various ways. The idea that substances are ontologically extravagant rests on a caricature. An electron is a candidate substance. An electron has a definite mass, charge, spin; these are ways the electron is. The electron is not an assemblage of these properties, they do not constitute the electron, nor is the electron an entity separable from them. An electron could cease to be some way—the electron could cease to have a particular spin—by becoming some other way.

If you are willing to think of properties as modes, ways particular substances are, the idea that substances are mysterious unobservable posits loses traction. A substance is not a faceless entity that combines with properties to form a concrete object. Properties are not applied to a substance as coats of paint are applied to a chair. A substance is something that is various ways; the ways are ways the substance is. Knowing the nature of a substance is a matter of knowing ways it is.

Properties are modes, ways substances are. But what distinguishes one property from another? A property is a quality; substances are as they are qualitatively, owing to their properties. Properties that differ qualitatively are distinct. Because properties are particulars, however, distinct properties can be qualitatively indiscernible. Properties indiscernible in this way are distinguished by distinguishing their bearers, the substances to which they belong. Socrates' paleness differs from Simmias' paleness in being *Socrates'* paleness.

Properties are not *mere* qualities. Properties are *powerful* qualities. Powers and qualities are not distinct kinds of property, not distinguishable aspects or components of properties, they are the properties themselves differently considered. The cubicity of a die is a quality of the die, a way the die is qualitatively. But in virtue of being cubical, the die would, when tossed, tumble in a distinctive way and eventually come to rest on one of its faces. So cubicity is a power. How this power manifests itself depends on its reciprocal partners. In virtue of being cubical, the die, when tossed, would tumble across a tabletop, would make a square impression if pressed against your forearm, would look and feel cubical: different kinds of reciprocal partner, different kinds of manifestation.

How does the idea that properties are powerful qualities extend to *abstracta* such as numbers or sets? The short answer is that it does not. Numbers and sets are not substances, not property bearers. Two is an even prime, but being even and being prime are not properties of an entity, the number two, that resides in a realm especially suited to house such things. Mathematical truths, like logical truths, make no demands on the universe and require no truthmakers. Whatever ways the universe is or could be are consistent with the mathematical and logical truths. This is not because numbers are eternal, incorruptible, present everywhere and nowhere. It is because talk of numbers enables us to work out relationships among worldly truths that would be difficult or impossible to discover otherwise and to work out all that these truths encompass.

One initially surprising corollary of the idea that substance and property are correlative categories, that substances are property bearers, properties ways substances are, is that substances must be simple. A substance cannot have parts that are themselves substances. A 'complex substance' would not be a substance, but a complex *of* substances, hence not an apt property bearer. Complex objects are not substances, apparent properties of complex objects are not properties. Or, if this sounds excessive, complex objects are substances by courtesy, *quasi* substances, their characteristics are properties by courtesy, *quasi* properties.

Think of a complex object that is, at a particular time, made up of particular simple substances in a particular arrangement. The object, at this time, *is* the arrangement, particular substances arranged in a particular way. Why not allow that this complex object has properties? You might reason as follows. Suppose properties are modes, and modes are ways substances are. And suppose the complex object *is* these substances arranged in this way, then the complex object is a particular way. The object's being this way *is* its being duly propertied. Properties are borne by substances, so complexes can be substances.

This conclusion could be turned around: the object's being duly propertied *is* its being this way, *is* its constituents being arranged as they are. Complex objects are various ways. These ways are ways objects' propertied constituents are arranged at a particular time. If you think of these ways as properties of the complexes, *all there is* to a

complex object's being propertied is its propertied constituents being arranged as they are. The object here and now *is* these constituents so arranged. You do not have an arrangement of constituent substances plus an object with *its* properties. The arrangement serves as truthmaker for claims about the object and its characteristics.

Thinking of complex objects, objects that have substances as parts, as *themselves* having properties encourages the kind of double counting that plagues strains of non-reductive physicalism. You do not have arrangements of propertied substances *and* a propertied complex object. To have the arrangement is just what it is to have the object as it is. One salutary feature of the thesis that substances *qua* property bearers must be simple is that it inhibits attempts to generate ontological levels. You have the substances and their properties variously arranged. These serve as truthmakers for all the truths about the complex objects—the tables, the trees, the planets—that populate the universe.

Talk of truthmaking brings us back to earth. Truthmaking, I contend, is an indispensable philosophical notion. When a philosopher makes a claim as to how things stand, it is almost always worth asking what the philosopher takes to be the truthmaker for the claim. The *nature* of truthmakers for many everyday and scientific truths could be largely beyond our ken. What makes it true that this tomato is red or that snow is white? For that matter, what makes it the case that this is a tomato or this is snow, what makes it the case that this satisfies application conditions for 'is a tomato' or 'is snow'? Answers to these questions take us eventually to fundamental physics.

The mistake, a mistake encouraged by linguisticized metaphysics, is to assume that truthmakers must be directly inferable from sentences used to express truths. A particularly egregious form of this mistake is to suppose that a property corresponds to every predicate (or every significant predicate, or every predicate that figures in respectable scientific discourse), and that distinct predicates must designate distinct properties.

We use predicates to express truths, truthmakers for which are rarely properties of substances. Only when you get to fundamental physics, do predicates begin to line up with properties. Once you see that none of this threatens the significance of truths that outstrip

truths of fundamental physics, that none of this threatens ordinary beliefs about the universe, you will have made the big move.

13.3 Serious Ontology

One measure of success in ontology is the extent to which a given account is internally coherent, the extent to which it fits together in sensible ways to yield a plausible picture of the universe. A second measure of an ontology's success is the degree to which it reconciles the manifest image of the universe, the universe as we find it in the special sciences and in ordinary life, with the conception of the universe that emerges in fundamental physics, where fundamental physics is understood as being in the business of providing the deep story concerning the truthmakers for every truth that has a truth-maker. A third measure of an ontology is its success in resolving pressing philosophical puzzles. I believe that the ontological picture recommended in these pages stacks up well on all these counts.

Some readers will have worked out why they do not need to agree, how the doctrines that come under criticism in these pages can be preserved, and when it is safe to ignore apparent threats to the status quo. This is the way of philosophy. We sometimes engage, sometimes talk past one another, sometimes operate in distinct worlds. Our motives are partly philosophical, partly sociological, partly personal. Philosophy as an academic profession makes many non-philosophical demands on its practitioners. This is a fact of life anyone would be foolish to ignore.

As philosophers, we labour under a multitude of assumptions inculcated in the course of learning the trade. These assumptions are like a well-worn pair of comfortable shoes. We can get around in them easily. But comfortable shoes cannot always take us where we need to be. It can be liberating to try on something different, to move beyond familiar well-worn forms of thought.

The decision to follow a new course is not one likely to be compelled by philosophical argument. It is the result of a choice based partly on purely philosophical considerations and partly on instinct, gut feeling. Hume advises, in the *Inquiry*, 'Be a philosopher; but, amidst all your philosophy, be still a man.' This passage is most

often read as supporting the idea that rational reflection is one thing, passion, instinct, gut feeling, another. But Hume says, '*amidst* all your philosophy'. No good philosopher imagines that philosophy (or science, or mathematics) is a purely intellectual endeavour, no good philosopher bifurcates the intellect and the gut, and no one lives philosophy more passionately than the philosophers to whom this book is dedicated. Let them be your guides.

Bibliography

Albritton, R. (1957) 'Substance and Form in Aristotle'. *Journal of Philosophy* 54: 699–708.

Angell, J. R. (1897) 'Thought and Imagery'. *Philosophical Review* 6: 646–51.

Armstrong, D. M. (1961) *Perception and the Physical World*. London: Routledge & Kegan Paul.

—— (1968) 'The Headless Woman Illusion and the Defense of Materialism'. *Analysis* 29: 48–9.

—— (1975) 'Towards a Theory of Properties: Work in Progress on the Problem of Universals'. *Philosophy* 50: 145–55.

—— (1978) *Universals and Scientific Realism*, 2 vols. Cambridge: Cambridge University Press.

—— (1983) *What is a Law of Nature?* Cambridge: Cambridge University Press.

—— (1989) *Universals: An Opinionated Introduction*. Boulder: Westview Press.

—— (1997) *A World of States of Affairs*. Cambridge: Cambridge University Press.

—— (1999) 'The Causal Theory of Properties: Properties According to Ellis, Shoemaker, and Others'. *Philosophical Topics* 26: 25–37.

—— (2004) *Truth and Truthmakers*. Cambridge: Cambridge University Press.

Bacon, J. (1995) *Universals and Property Instances: The Alphabet of Being*. Oxford: Blackwell.

—— (2008) 'Tropes'. *Stanford Encyclopedia of Philosophy*. Ed. E. N. Zalta. <http://plato.stanford.edu/archives/spr2008/entries/tropes/>

—— K. Campbell, and L. Reinhardt, eds. (1993) *Ontology, Causality, and Mind: Essays in Honour of D. M. Armstrong*. Cambridge: Cambridge University Press.

Beckermann, A., H. Flohr, and J. Kim, eds. (1992) *Emergence or Reduction?* Berlin: De Gruyter.

Below, C. (1987) *The Vedgymight History of Australia*, 2nd edn. Hawthorne, Vic.: N. H. Hudson Publishing Services.

Bennett, J. F. (1964) *Rationality: An Essay Towards an Analysis*. London: Routledge & Kegan Paul.

—— (2001) *Learning from Six Philosophers*, vol. 2. Oxford: Clarendon Press.

Bennett, K. (2007) 'Mental Causation'. *Philosophy Compass* 2: 316–37.

Bergmann, G. (1967) *Realism*. Madison: University of Wisconsin Press.

Berkeley, G. (1710/1998) *A Treatise Concerning the Principles of Human Knowledge*. Ed. J. Dancy. Oxford: Oxford University Press.

Bermúdez, J. L. (2003) *Thinking Without Words*. Oxford: Oxford University Press.

Bigelow, J. (1988) *The Reality of Numbers*. Oxford: Clarendon Press.

Bird, A. (2007) *Nature's Metaphysics: Laws and Properties*. Oxford: Clarendon Press.

Blackburn, S. (1984) *Spreading the Word: Groundings in the Philosophy of Language*. Oxford: Clarendon Press.

—— (1990) 'Filling in Space'. *Analysis* 50: 62–5.

Block, N., ed. (1980) *Readings in Philosophy of Psychology*, vol. 1. Cambridge, MA: Harvard University Press.

Boyd, R. (1980) 'Materialism Without Reductionism: What Physicalism Does not Entail'. In Block (1980): 67–106.

—— (1999) 'Homeostasis, Species, and Higher Taxa'. In R. A. Wilson, ed., *Species: New Interdisciplinary Essays*. Cambridge, MA: MIT Press: 141–86. The abridged version cited here appears in Lange (2007): 406–26.

Braddon-Mitchell, D., and F. C. Jackson (1996) *Philosophy of Mind and Cognition*. Oxford: Blackwell.

Brand, M., and D. Walton, eds. (1976) *Action Theory*. Dordrecht: Kluwer Academic Publishers.

Brower, J. (2009) 'Medieval Theories of Relations'. *Stanford Encyclopedia of Philosophy*. Ed. E. N. Zalta. <http://plato.stanford.edu/archives/win2010/entries/relations-medieval/>

Bueno, O., and O. Linnebo, eds. (2010) *New Waves in Philosophy of Mathematics*. Basingstoke: Palgrave Macmillan.

Byrne, A. (2001) 'Intentionalism Defended'. *Philosophical Review* 110: 199–240.

Cameron, R. P. (2008) 'Turtles All the Way Down: Regress, Priority, and Fundamentality'. *Philosophical Quarterly* 58: 1–14.

—— (2010) 'Necessity and Triviality'. *Australasian Journal of Philosophy* 88: 401–15.

Campbell, K. (1976) *Metaphysics: An Introduction*. Encino: Dickenson Publishing Co.

Campbell, K. (1981) 'The Metaphysics of Abstract Particulars'. *Midwest Studies in Philosophy* 6: 477–88.

——(1990) *Abstract Particulars*. Oxford: Basil Blackwell.

Capitan, W. H., and D. D. Merrill, eds. (1967) *Art, Mind, and Religion*. Pittsburgh: University of Pittsburgh Press.

Carruthers, P. (1996) *Language, Thought, and Consciousness*. Cambridge: Cambridge University Press.

——and J. Boucher, eds. (1998) *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press.

Cartwright, N. (1989) *Nature's Capacities and Their Measurement*. Oxford: Clarendon Press.

——(1999) *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.

Chakravartty, A. (2005) 'Causal Realism: Events and Processes'. *Erkenntnis* 63: 7–31.

——(2007) *A Metaphysics for Scientific Realism: Knowing the Unobservable*. Cambridge: Cambridge University Press.

Chalmers, D. J. (1996) *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.

Clark, A. (1998) 'Magic Words: How Language Augments Human Computation'. In Carruthers and Boucher (1998): 162–83.

Collingwood, R. G. (1940) *An Essay on Metaphysics*. Oxford: Clarendon Press.

Collins, J., E. J. Hall, and L. A. Paul, eds. (2004) *Causation and Counterfactuals*. Cambridge, MA: MIT Press.

Cottingham, J., R. Stoothoff, and D. Murdoch, trans. (1984) *The Philosophical Writings of Descartes*, vol. 2. Cambridge: Cambridge University Press.

——R. Stoothoff, and D. Murdoch, trans. (1985) *The Philosophical Writings of Descartes*, vol. 1. Cambridge: Cambridge University Press.

Craik, K. (1967) *The Nature of Explanation*. Cambridge: Cambridge University Press.

Crane, T., ed. (1996) *A Debate on Dispositions*. By D. M. Armstrong, C. B. Martin, and U. T. Place. London: Routledge.

Danzinger, K. (1980) 'The History of Introspectionism Reconsidered'. *Journal of the History of the Behavioral Sciences* 16: 241–62.

Davidson, D. (1970) 'Mental Events'. In Foster and Swanson (1970): 79–101. Reprinted in Davidson (1980/2001): 207–225 (citations are to the latter).

——(1973) 'Freedom to Act'. In Honderich (1973): 137–56. Reprinted in Davidson (1980): 63–81 (citations are to the latter).

——(1975) 'Thought and Talk'. In Guttenplan (1975): 7–23. Reprinted in Davidson (1984): 155–70.

——(1980) *Essays on Actions and Events*. Oxford: Clarendon Press.

——(1982) 'Rational Animals'. *Dialectica* 36: 317–27. Reprinted in LePore and McLaughlin (1985): 473–80.

——(1984) *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press.

Demos, R. (1948) 'Note on Plato's Theory of Ideas'. *Philosophy and Phenomenological Research* 8: 456–60.

Descartes, R. (1637) *Discourse on Method*. Trans. J. Cottingham. In Cottingham et al. (1985): 111–51.

——(1641/1984) 'Reply to the Points Which May Cause Difficulty to Theologians'. Trans. J. Cottingham in Cottingham et al. (1984): 172–8.

——(1641/1986) *Meditations on First Philosophy*. Trans. J. Cottingham. Cambridge: Cambridge University Press.

——(1644/1985) *Principles of Philosophy*. Trans. J. Cottingham. In Cottingham et al. (1985): 179–291.

——(1664) *Treatise on Man*. Trans. R. Stoothoff. In Cottingham et al. (1985): 99–108.

Dipert, R. R. (1997) 'The Mathematical Structure of the World: The World as Graph'. *Journal of Philosophy* 94: 329–58.

Dretske, F. I. (1988) *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.

——(1995) *Naturalizing the Mind*. Cambridge, MA: MIT Press.

Dupré, J. (1993) *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.

Dyke, H. (2008) *Metaphysics and the Representational Fallacy*. New York: Routledge.

Eddington, A. S. (1928) *The Nature of the Physical World*. London: Macmillan.

Ellis, B. D. (2001) *Scientific Essentialism*. Cambridge: Cambridge University Press.

——(2002) *Philosophy of Nature: A Guide to the New Essentialism*. Chesham: Acumen.

Farrell, B. A. (1950) 'Experience'. *Mind* 59: 170–98.

Feigl, H. (1958) 'The "Mental" and the "Physical"'. In Feigl, Scriven, and Maxwell (1958): 370–497. Reprinted in Feigl (1967).

——(1967) *The 'Mental' and the 'Physical': The Essay and a Postscript*. Minneapolis: University of Minnesota Press.

Feigl, H., M. Scriven, and G. Maxwell, eds. (1958) *Concepts, Theories, and the Mind–Body Problem*. Minnesota Studies in the Philosophy of Science, vol. 2. Minneapolis: University of Minnesota Press.

Feynman, R. P. (1965) *The Character of Physical Law*. London: British Broadcasting Corporation.

Fine, K. (1991) 'The Study of Ontology'. *Noûs* 25: 263–94.

——(1994) 'Essence and Modality'. *Philosophical Perspectives* 8: 1–16.

Fodor, J. A. (1974) 'Special Sciences (Or: The Disunity of Science as a Working Hypothesis)'. *Synthese* 28: 97–115.

——(1975) *The Language of Thought*. New York: T. Y. Crowell.

Foster, J. (1982) *The Case for Idealism*. London: Routledge & Kegan Paul.

Foster, L., and J. Swanson, eds. (1970) *Experience and Theory*. Amherst: University of Massachusetts Press.

Fox, J. (1987) 'Truthmaker'. *Australasian Journal of Philosophy* 65: 188–207.

Freeman, A., ed. (2006) *Consciousness and Its Place in Nature: Does Physicalism Entail Panpsychism?* Exeter: Imprint Academic.

French, S., and D. Krause (2006) *Identity in Physics: A Historical, Philosophical, and Formal Analysis*. Oxford: Clarendon Press.

Garcia, R. K. (2015) 'Is Trope Theory a Divided House?', in G. Galluzzo and M. J. Loux, eds., *The Problem of Universals in Contemporary Philosophy* (Cambridge: Cambridge University Press, 2015).

Gibbons, J. (2006) 'Mental Causation Without Downward Causation'. *Philosophical Review* 115: 79–103.

Gibson, J. J. (1966) *The Senses Considered as Perceptual Systems*. Boston: Houghton Mifflin.

Gillett, C., and B. Loewer, eds. (2001) *Physicalism and Its Discontents*. Cambridge: Cambridge University Press.

Goodman, N. (1966) *The Structure of Appearance*, 2nd edn. Indianapolis: Bobbs-Merrill.

Guttenplan, S., ed. (1975) *Mind and Language*. Oxford: Clarendon Press.

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

Harman, G. (1990) 'The Intrinsic Quality of Experience'. *Philosophical Perspectives* 4: 31–52.

Harré, R., and E. H. Madden (1975) *Causal Powers: A Theory of Natural Necessity*. Oxford: Basil Blackwell.

Healey, R. A. (1991) 'Holism and Nonseparability'. *Journal of Philosophy* 88: 393–421.

Heil, J. (1970) 'Sensations, Experiences and Brain-Processes'. *Philosophy* 45: 221–6.

——(1983) *Perception and Cognition*. Berkeley: University of California Press.

——(1992) *The Nature of True Minds*. Cambridge: Cambridge University Press.

——(1998) 'Supervenience Deconstructed'. *European Journal of Philosophy* 6: 146–55.

——(2003a) *From an Ontological Point of View*. Oxford: Clarendon Press.

——(2003b) 'Multiply Realized Properties'. In Walter and Heckmann (2003): 11–30.

——(2005a) 'Dispositions', *Synthese* 144: 343–56.

——(2005b) 'Real Tables'. *The Monist* 88: 493–509.

——(2008a) 'Critical Notice of Peter Unger's *All the Power in the World*'. *Noûs* 42: 336–48.

——(2008b) 'Anomalous Monism'. In H. Dyke, ed., *From Truth to Reality: New Essays in Metaphysics*. London: Routledge, 2008: 85–98.

——(2009) 'Relations'. In R. Le Poidevin and R. Cameron, eds., *Routledge Companion to Metaphysics*. London: Routledge: 310–21.

——(2011) 'Powers and the Realization Relation'. *The Monist* 94: 34–53.

——and Robb, D. M. (2003) 'Mental Properties'. *American Philosophical Quarterly* 40: 175–96.

Holton, R. (1999) 'Dispositions All the Way Round'. *Analysis* 59: 9–14.

Honderich, T., ed. (1973) *Essays on Freedom of Action*. London: Routledge & Kegan Paul.

Huemer, M., and B. Kovitz. (2003) 'Causation as Simultaneous and Continuous'. *Philosophical Quarterly* 53: 556–65.

Hume, D. (1739/1978) *A Treatise of Human Nature*. Ed. L. A. Selby-Bigge 2nd edn. rev. P. H. Nidditch. Oxford: Clarendon Press.

Ingthorsson, R. D. (2002) 'Causal Production as Interaction'. *Metaphysica* 3: 87–119.

Jackson, F. C. (1982) 'Epiphenomenal Qualia'. *Philosophical Quarterly* 32: 127–36.

——(1998) *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. Oxford: Clarendon Press.

Johnson-Laird, P. N. (1983) *Mental Models*. Cambridge: Cambridge University Press.

——(1999) 'Mental Models'. In Wilson and Keil (1999): 525–7.

Kant, I. (1787/1929) *Critique of Pure Reason*. Trans. N. K. Smith. London: Macmillan & Co.

Kim, J. (1976) 'Events as Property Exemplifications'. In Brand and Walton (1976): 159–77.

——(1984) 'Concepts of Supervenience'. *Philosophy and Phenomenological Research* 45: 153–76.

——(1989) 'The Myth of Non-reductive Materialism'. *Proceedings of the American Philosophical Association* 63: 31–47. Reprinted in Kim (1993): 265–84 (citations are to the latter).

——(1990) 'Supervenience as a Philosophical Concept'. *Metaphilosophy* 12: 1–27.

——(1993) *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press.

——(1999) 'Making Sense of Emergence'. *Philosophical Studies* 95: 3–36.

——(2005) *Physicalism, or Something Near Enough*. Princeton: Princeton University Press.

Kirk, G. S., and J. E. Raven (1964) *The Presocratic Philosophers: A Critical History with a Selection of Texts*. Cambridge: Cambridge University Press.

Kirk, R. (1974) 'Zombies versus Materialists'. *Aristotelian Society Proceedings* (suppl.) 48: 135–52.

——(2005) *Zombies and Consciousness*. Oxford: Clarendon Press.

Kosslyn, S. M., and J. R. Pomerantz (1977) 'Imagery, Propositions, and the Form of Internal Representations'. *Cognitive Psychology* 9: 52–76.

Küng, G. (1967) *Ontology and the Logistic Analysis of Language*. Dordrecht: Reidel.

Ladyman, J. (2007) 'Structural Realism'. In *The Stanford Encyclopedia of Philosophy*. Ed. E. N. Zalta. <http://plato.stanford.edu/archives/win2007/entries/structural-realism/>

——and D. Ross, with D. Spurrett and J. Collier (2007) *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.

Lange, M., ed. (2007) *Philosophy of Science: An Anthology*. Oxford: Blackwell Publishing.

Le Poidevan, R., P. Simons, A. McGonigal, and R. P. Cameron, eds. (2009) *The Routledge Companion to Metaphysics*. London: Routledge.

Leibniz, G. W. (1715) 'Leibniz–Clarke Correspondence'. Trans. M. Morris and G. H. R. Parkinson. In *Leibniz: Philosophical Writings*. Ed. G. H. R. Parkinson. London: J. M. Dent & Sons, 1973.

LePore, E., and B. P. McLaughlin, eds. (1985) *Actions and Events*. Oxford: Basil Blackwell.

Levinson, J. (1978) 'Properties and Related Entities'. *Philosophy and Phenomenological Research* 39: 1–22.

Levinson, J. (1980) 'The Particularization of Attributes'. *Australasian Journal of Philosophy* 58: 102–15.

——(2006) 'Why There Are No Tropes'. *Philosophy* 81: 563–79.

Lewis, D. K. (1986a) *Philosophical Papers*, vol. 2. New York: Oxford University Press.

Lewis, D. K. (1986b) *On the Plurality of Worlds*. Oxford: Basil Blackwell.

——(1991) *Parts of Classes*. Oxford: Basil Blackwell.

——(2000) 'Causation as Influence'. *Journal of Philosophy* 97: 182–97.

——(2001) 'Truthmaking and Difference-Making'. *Noûs* 35: 602–15.

——(2004) 'Void and Object'. In Collins et al. (2004): 277–90.

Locke, J. (1690/1978) *An Essay Concerning Human Understanding*. Ed. P. H. Nidditch. Oxford: Clarendon Press.

——(1706) 'An Examination of P. Malebranche's Opinion'. In *Posthumous Works of Mr. John Locke*. London: Printed by W. B. for A. and J. Churchill). Reprinted in *Works of John Locke* (10 vols.). London: Thomas Tegg (1823), vol. IX: 211–255.

Lowe, E. J. (1994) 'Vague Identity and Quantum Indeterminacy'. *Analysis* 54: 110–14.

——(1996) *Subjects of Experience*. Cambridge: Cambridge University Press.

——(1998) *The Possibility of Metaphysics: Substance, Identity, and Time*. Oxford: Clarendon Press.

——(2002) 'Metaphysical Nihilism and the Subtraction Argument'. *Analysis* 62: 62–73.

——(2006) *The Four-Category Ontology: A Metaphysical Foundation for Natural Science*. Oxford: Clarendon Press.

——(2008) *Personal Agency: The Metaphysics of Mind and Action*. Oxford: Oxford University Press.

Lycan, W. (1996) *Consciousness and Experience*. Cambridge, MA: MIT Press.

McCall, S., and E. J. Lowe (2006) 'The 3D/4D Controversy: A Storm in a Teacup'. *Noûs* 40: 570–8.

Macdonald C., and G. Macdonald (2006) 'The Metaphysics of Mental Causation'. *Journal of Philosophy* 103: 539–76.

——and G. Macdonald, eds. (2010) *Emergence in Mind*. New York: Oxford University Press.

McGinn, C. (1989) 'Can we Solve the Mind–Body Problem?' *Mind* 98: 349–66.

McLaughlin, B. P. (1992) 'The Rise and Fall of British Emergentism'. In Beckermann et al. (1992): 49–93.

——A. Beckermann, and S. Walter, eds. (2009) *The Oxford Handbook of Philosophy of Mind*. Oxford: Clarendon Press.

Martin, C. B. (1980) 'Substance Substantiated'. *Australasian Journal of Philosophy* 58: 3–10.

—— (1987) 'Proto-Language'. *Australasian Journal of Philosophy* 65: 277–89.

—— (1993) 'Power for Realists'. In Bacon et al. (1993): 175–86.

—— (1994) 'Dispositions and Conditionals'. *Philosophical Quarterly* 44: 1–8.

—— (1996) 'How It Is: Entities, Absences, and Voids'. *Australasian Journal of Philosophy* 74: 57–65.

—— (1997) 'On the Need for Properties: The Road to Pythagoreanism and Back'. *Synthese* 112: 193–231.

—— (2008) *The Mind in Nature*. Oxford: Clarendon Press.

—— and J. Heil (1999) 'The Ontological Turn'. *Midwest Studies in Philosophy* 23: 34–60.

Matthews, G. B., and S. M. Cohen. (1968) 'The One and the Many'. *Review of Metaphysics* 21: 630–55.

Maurin, A.-S. (2002) *If Tropes*. Dordrecht: Kluwer Academic Publishers.

Mertz, D. W. (1996) *Moderate Realism and its Logic*. New Haven: Yale University Press.

Molnar, G. (2003) *Powers: A Study in Metaphysics*. Ed. S. Mumford. Oxford: Oxford University Press.

Moore, A. W. (2001) *The Infinite*, 2nd edn. London: Routledge.

Moore, G. E. (1903) 'The Refutation of Idealism'. *Mind* 12: 433–53. Reprinted in Moore (1922): 1–30.

—— (1919) 'External and Internal Relations'. *Proceedings of the Aristotelian Society* 20: 40–62. Reprinted in Moore (1922): 276–309.

—— (1922) *Philosophical Studies*. London: K. Paul, Trench, Trubner & Co.; New York: Harcourt, Brace & Co.

Morrison, J. S. (1977) 'Two Unresolved Difficulties in the Line and the Cave'. *Phronesis* 22: 212–31.

Mulligan, K., P. Simons, and B. Smith (1984) 'Truth-Makers'. *Philosophy and Phenomenological Research* 44: 287–321.

Mumford, S. (1998) *Dispositions*. Oxford: Clarendon Press.

Nagel, T. (1974) 'What is it Like to be a Bat?' *Philosophical Review* 83: 435–50.

O'Connor, T. (1994) 'Emergent Properties'. *American Philosophical Quarterly* 31: 91–104.

—— and J. Jacobs (2003) 'Emergent Individuals'. *Philosophical Quarterly* 53: 540–55.

—— and H. Y. Wong (2005) 'The Metaphysics of Emergence'. *Noûs* 39: 658–78.

Parsons, J. (1999) 'There is no "Truthmaker" Argument against Nominalism'. *Australasian Journal of Philosophy* 77: 325–34.

—— (2009) 'Are there Irreducibly Relational Facts?' In E. J. Lowe and E. Rami, eds., *Truth and Truth-Making*. Stocksfield: Acumen Press: 217–26.

Pereboom, D. (2002) 'Robust Non-reductive Materialism'. *Journal of Philosophy* 99: 499–531.

Place, U. T. (1956) 'Is Consciousness a Brain Process?' *British Journal of Psychology* 47: 44–50.

Priestley, J. (1775/1973) *Hartley's Theory of the Human Mind, on the Principle of the Association of Ideas; with Essays Relating to the Subject of It*. Ed. T. Verhave. New York: AMS Press.

—— (1782/1972) *Disquisitions of Matter and Spirit*, 2nd edn. In *The Theological and Miscellaneous Works of Joseph Priestley*, Vol. 3. New York: Kraus Reprint Co.

Prior, E. W., R. Pargetter, and F. C. Jackson (1982) 'Three Theses about Dispositions'. *American Philosophical Quarterly* 19: 251–7.

Putnam, H. (1967) 'Psychological Predicates'. In Capitan and Merrill (1967): 37–48. Reprinted as 'The Nature of Mental States' in Putnam (1975): 429–40.

—— (1975) *Philosophical Papers*, vol. 2. Cambridge: Cambridge University Press.

Pylyshyn, Z. W. (1981) 'The Imagery Debate: Analogue Media Versus Tacit Knowledge'. *Psychological Review* 88: 16–45.

—— (1984). *Computation and Cognition: Toward a Foundation for Cognitive Science*. Cambridge, MA: MIT Press.

Quine, W. V. O. (1948) 'On What There Is'. *Review of Metaphysics* 2: 21–38. Reprinted in Quine (1961a): 1–19.

—— (1961a) *From a Logical Point of View*, 2nd edn. Cambridge, MA: Harvard University Press. Reprinted in 1963 by Harper & Row.

—— (1961b) 'Logic and the Reification of Universals'. In Quine (1961a): 102–29.

Rayo, A. (2010) 'Towards a Trivialist Account of Mathematics'. In Bueno and Linnebo (2010): 239–62.

Reichenbach, H. (1956) *The Direction of Time*. Berkeley: University of California Press.

Resnik, M. (1997) *Mathematics as a Science of Patterns*. Oxford: Clarendon Press.

Rips, L. (1994) *The Psychology of Proof*. Cambridge, MA: MIT Press.

Robb, D. M. (1997) 'The Properties of Mental Causation'. *Philosophical Quarterly* 47: 178–94.

—— (2005) 'Qualitative Unity and the Bundle Theory'. *Monist* 88: 466–92.

—— (2009) 'Substance'. In Le Poidevin et al. (2009): 256–64.

Rollins, M. (1989) *Mental Imagery: On the Limits of Cognitive Science*. New Haven: Yale University Press.

Russell, B. (1903/1997) 'Asymmetrical Relations'. In *The Principles of Mathematics*. London: Routledge, 1903: 218–26; reprinted 1997.

—— (1912) 'On the Notion of a Cause'. *Proceedings of the Aristotelian Society* (New Series) 13: 1–26. Reprinted in Russell (1918): 180–208.

—— (1918) *Mysticism and Logic*. London: Longmans, Green and Co.

—— (1948) *Human Knowledge: Its Scope and Limits* 3rd edn. New York: Simon & Schuster.

Ryle, G. (1949) *The Concept of Mind*. London: Hutchinson and Co.

Santayana, G. (1930) *The Realm of Matter*. New York: Charles Scribner's Sons.

Saunders, A. (2009) 'Getting it Across'. *Australasian Journal of Philosophy* 87: 1–3.

Schaffer, J. (2001) 'The Individuation of Tropes'. *Australasian Journal of Philosophy* 79: 247–57.

—— (2004) 'Causes Need Not Be Physically Connected to their Effects: The Case for Negative Causation'. In C. Hitchcock, ed., *Contemporary Debates in Philosophy of Science*. Oxford: Blackwell Publishers: 197–216.

—— (2010) 'Monism: The Priority of the Whole'. *Philosophical Review* 119: 31–76.

Schrödinger, E. (1951) *Science and Humanism: Physics in Our Time*. Cambridge: Cambridge University Press.

Seargent, D. A. J. (1985) *Plurality and Continuity*. Dordrecht: Martinus Nijhoff Publishers.

Searle, J. R. (1983) *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.

—— (1992) *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.

—— (1997) *The Mystery of Consciousness*. New York: A New York Review Book.

—— (2004) *Mind: A Brief Introduction*. Oxford: Oxford University Press.

Sellars, W. I. (1957) 'Substance and Form in Aristotle'. *Journal of Philosophy* 54: 688–99.

Shapiro, S. (1997) *Philosophy of Mathematics: Structure and Ontology*. Oxford: Oxford University Press.

Sherrington, C. S. (1906/1947) *The Integrative Action of the Nervous System*. Yale: Yale University Press. Republished with a new foreword in 1947 by Cambridge University Press. Citations in this volume are to the 1947 edition.

Shoemaker, S. (1980) 'Causality and Properties'. In P. van Inwagen, ed., *Time and Cause*. Dordrecht: Reidel Publishing Co.: 109–35. Reprinted in Shoemaker (1984): 206–233.

—— (1984) *Identity, Cause, and Mind: Philosophical Essays*. Cambridge: Cambridge University Press.

—— (1998) 'Causal and Metaphysical Necessity'. *Pacific Philosophical Quarterly* 79: 59–77.

—— (2001) 'Realization and Mental Causation'. In Gillett and Loewer (2001): 74–98.

—— (2007) *Physical Realization*. Oxford: Oxford University Press.

—— (2011) 'Realization, Powers, and Property Identity'. *Monist* 94: 3–18.

Simons, P. (1994) 'Particulars in Particular Clothing: Three Trope Theories of Substance'. *Philosophy and Phenomenological Research* 54: 553–75.

Smart, J. J. C. (1959) 'Sensations and Brain Processes'. *Philosophical Review* 68: 141–56.

—— (1963) *Philosophy and Scientific Realism*. London: Routledge & Kegan Paul.

Stout, G. F. (1921) 'The Nature of Universals and Propositions'. *Proceedings of the British Academy* 10 (1921–23): 157–72. Oxford University Press. Reprinted in Stout (1930): 384–403.

—— (1930) *Studies in Philosophy and Psychology*. London: Macmillan and Co.

—— (1936) 'Universals Again'. *Proceedings of the Aristotelian Society*, Supplementary Volume 15: 1–15.

Strawson, G. (2006) 'Realistic Monism: Why Physicalism Entails Panpsychism'. In Freeman (2006): 3–31.

—— (2008a) 'The Identity of the Categorical and the Dispositional'. *Analysis* 68: 271–82.

—— (2008b) *Real Materialism and Other Essays*. Oxford: Clarendon Press.

Stump, E. (2005) *Aquinas*, 2nd edn. Abingdon: Routledge.

Teller, P. (1986) 'Relational Holism and Quantum Mechanics'. *British Journal for the Philosophy of Science* 37: 71–81.

—— (1995) *An Interpretive Introduction to Quantum Field Theory*. Princeton: Princeton University Press.

Thomas, N. J. T. (1999) 'Are Theories of Imagery Theories of Imagination?' *Cognitive Science* 23: 207–45.

Thomas, N. J. T. (2001) 'Mental Imagery'. *Stanford Encyclopedia of Philosophy*. Ed. E. N. Zalta. <http://plato.stanford.edu/archives/win2001/entries/mental-imagery/>

Tolman, E. C. (1932) *Purposive Behavior in Rats and Men*. New York: Century Co.

Tye, M. (1991) *The Imagery Debate*. Cambridge, MA: MIT Press.

——— (1995) *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, MA: MIT Press.

——— (2002) 'Representationalism and the Transparency of Experience'. *Noûs* 36: 137–51.

——— (2009) 'Representationalist Theories of Consciousness'. In McLaughlin et al. (2009): 253–67.

Unger, P. (2006) *All the Power in the World*. New York: Oxford University Press.

Van Inwagen, P. (2007) 'A Set of Accidents'. *Times Literary Supplement*, 21 and 28 December: 36.

Walter, S., and H.-D. Heckmann, eds. (2003) *Physicalism and Mental Causation: The Metaphysics of Mind and Action*. Exeter: Imprint Academic.

Weiskrantz, L. (1986) *Blindsight: A Case Study and Implications*. Oxford: Oxford University Press.

Williams, D. C. (1934) 'Scientific Method and the Existence of Consciousness'. *Psychological Review* 41: 461–79. Reprinted in Williams (1966): 23–40.

——— (1944) 'Naturalism and the Nature of Things'. *Philosophical Review* 53: 417–43. Reprinted in Williams (1966): 212–38.

——— (1953) 'On the Elements of Being'. *Review of Metaphysics* 7: 3–18, 171–92. Reprinted as 'The Elements of Being' in Williams (1966): 74–109.

——— (1954) 'Of Essence and Existence in Santayana'. *Journal of Philosophy* 51: 31–42.

——— (1959/1986) 'Universals and Existents'. A paper delivered to the Yale Philosophy Club and published posthumously in *The Australasian Journal of Philosophy* 64 (1986): 1–14.

——— (1966) *Principles of Empirical Realism*. Springfield: Charles C. Thomas.

Wilson, J. (2011) 'Non-reductive Realization and the Powers Sub-set Strategy'. *Monist* 94: 212–54.

Wilson, R. A., and F. C. Keil, eds. (1999) *Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press.

Wittgenstein, N. (1921/1961) *Tractatus Logico-Philosophicus*. Trans. D. F. Pears and B. F. McGuinness. London: Routledge & Kegan Paul.

——— (1953/1968) *Philosophical Investigations*. Trans. G. E. M. Anscombe. Oxford: Basil Blackwell.

Wolterstorff, N. (1970) *On Universals*. Chicago: University of Chicago Press.

Woodworth, R. S. (1906) 'Imageless Thought'. *Journal of Philosophy, Psychology, and Scientific Methods* 3: 701–8.

Index