

Consciousness and the Existence of God

A theistic argument

J.P. Moreland

Contents

<i>Preface</i>	x
1 The epistemic backdrop for locating consciousness in a naturalist ontology	1
2 The argument from consciousness	28
3 John Searle and contingent correlation	53
4 Timothy O'Connor and emergent necessitation	70
5 Colin McGinn and mysterian “naturalism”	95
6 David Skrbina and panpsychism	114
7 Philip Clayton and pluralistic emergentist monism	135
8 Science and strong physicalism	156
9 AC, dualism and the fear of God	175
<i>Notes</i>	195
<i>Bibliography</i>	221
<i>Indexes</i>	230

1 The epistemic backdrop for locating consciousness in a naturalist ontology

Leibniz's gauntlet

Consciousness is among the most mystifying features of the cosmos. During the emergence of the mechanical philosophy in the seventeenth century, Leibniz wrote the following as a challenge to mechanistic materialism:

It must be confessed, moreover, that *perception*, and that which depends on it *are inexplicable by mechanical causes*, that is, by figures and motions. And, supposing there were a machine so constructed as to think, feel and have perception, we could conceive of it as enlarged and yet preserving the same proportions, so that we might enter it as a mill. And this granted, we should only find on visiting it, pieces which push one against another, but never anything by which to explain a perception. This must be sought for, therefore, in the simple substance and not in the composite or in the machine.¹

And while different bells and whistles have been added to our conception of matter since Leibniz's time, scientific naturalist explanations for the emergence of consciousness are as inadequate today as they were when Leibniz threw down his gauntlet. Thus, Geoffrey Madell opines that "the emergence of consciousness, then is a mystery, and one to which materialism signally fails to provide an answer."² Colin McGinn claims that its arrival borders on sheer magic because there seems to be no naturalistic explanation for it: "How can mere matter originate consciousness? How did evolution convert the water of biological tissue into the wine of consciousness? Consciousness seems like a radical novelty in the universe, not prefigured by the after-effects of the Big Bang; so how did it contrive to spring into being from what preceded it?"³

Not only are adequate naturalistic explanations for irreducible consciousness hard to come by, there is a widespread suspicion, if not explicit acknowledgement that irreducible consciousness provides evidence for theism. Thus, Crispin Wright notes,

2 *The epistemic backdrop*

A central dilemma in contemporary metaphysics is to find a place for certain anthropocentric subject-matters—for instance, semantic, moral, and psychological—in a world as conceived by modern naturalism: a stance which inflates the concepts and categories deployed by (finished) physical science into a metaphysics of the kind of thing the real world essentially and exhaustively is. On one horn, if we embrace this naturalism, it seems we are committed either to reductionism: that is, to a construal of the reference of, for example, semantic, moral and psychological vocabulary as somehow being within the physical domain—or to disputing that the discourses in question involve reference to what is real at all. On the other horn, if we reject this naturalism, then we accept that there is more to the world than can be embraced within a physicalist ontology—and so take on a commitment, it can seem, to a kind of eerie supernaturalism.⁴

Similarly, William Lyons notes that

[physicalism] seem[s] to be in tune with the scientific materialism of the twentieth century because it [is] a harmonic of the general theme that all there is in the universe is matter and energy and motion and that humans are a product of the evolution of species just as much as buffaloes and beavers are. Evolution is a seamless garment with no holes wherein souls might be inserted from above.⁵

Wright's allusion to a commitment to "a kind of eerie supernaturalism," and Lyons' reference to souls being "inserted from above" appear to be veiled references to the explanatory power of theism for consciousness, viz., that if "souls" exist, they would have to be "inserted from above" since natural processes by themselves are "seamless." More generally, some argue that, while certain features of consciousness or other finite mental entities—construed as *sui generis* and non-physical—may be inexplicable on a naturalist worldview, they may be explained by theism, thereby furnishing evidence for God's existence.

It is clear that for the last two-thirds of the twentieth century, mental entities have been recalcitrant facts for naturalists. Indeed, for philosophers who take the issues and options in philosophy of mind to be significantly influenced by empirical considerations, the proliferation of a wild variety of physicalist specifications of a naturalist treatment of mental phenomena may fairly be taken as a sign that naturalism is in a period of Kuhnian paradigm crisis. The argument from consciousness for God's existence (hereafter, AC) provides a way of dethroning the naturalist hegemony. Moreover, by giving a more adequate analysis of and explanation for mental entities, it provides a way out of the crisis and, together with other lines of evidence, offers materials for a cumulative-case argument for theism, or so I shall argue in the pages to follow.

In the last twenty years or so, versions of naturalism have multiplied like rabbits, so before we examine AC and some of its rivals, it is important to clarify two factors that constitute the appropriate dialectical background for the arguments to follow. First, I shall unpack the ideational structure of a version of naturalism that follows most plausibly from taking it as a worldview that claims explanatory, epistemic superiority to its rivals. Second, I shall lay out the central epistemic conditions relevant to assessing the force of AC *vis-à-vis* naturalism.

Exactly what are the central features of contemporary scientific naturalism?⁶ There will be different nuances given to naturalism by different thinkers, as one would expect with any widely accepted ideology. Nevertheless, it is both possible and desirable to give an accurate characterization of a specific form of philosophical naturalism (hereafter, simply naturalism or scientific naturalism) that is currently enjoying widespread acceptance. Moreover, by clarifying the relationship between a naturalist ontology on the one hand, and its epistemology and creation account on the other, a picture will emerge as to what *ought* to constitute that ontology. This picture will allow us to identify a substantial burden of proof for alternative naturalist ontologies that bloat naturalist metaphysical commitments beyond what is justifiable within the constraints that follow from the other two aspects of a naturalist worldview.

Fundamentally, and as a first attempt at characterization, naturalism is the view that the spatio-temporal universe of entities postulated by our best current (or ideal) theories in the physical sciences (or relevantly similar entities), particularly physics, is all there is. Scientific naturalism includes:

- 1 Different aspects of a naturalist epistemic attitude (e.g. acceptance of naturalized epistemology, a rejection of so-called first philosophy along with an acceptance of either weak or strong scientism);
- 2 An etiological account of how all entities whatsoever have come to be, constituted by an event-causal story (especially the atomic theory of matter and evolutionary biology) described in natural scientific terms; and
- 3 A general ontology in which the only entities allowed are ones that bear a relevant similarity to those thought to characterize a completed form of physics. Whether or not this ontology should be expanded to include *sui generis* emergent properties, e.g. secondary qualities, normative or mental properties, will occupy our attention shortly.

The ordering of these three ingredients is important. The naturalist epistemic attitude serves as justification for the naturalist etiology, which together justify the naturalist's ontological commitment. Moreover, naturalism seems to require coherence among what is postulated in these three different areas of the naturalistic turn. Thus, in setting up his naturalist project, David Papineau claims that we should set philosophy within

4 *The epistemic backdrop*

science in the sense that philosophical investigation should be conducted within the framework of our best empirical theories. It follows, says Papi-
neau, that “the task of the philosophers is to bring coherence and order to
the set of assumptions we use to explain the empirical world.”⁷ By way of
application, there should be coherence among third-person scientific ways
of knowing, a physical, evolutionary account of how our sensory and
cognitive processes came to be, and an ontological analysis of those pro-
cesses themselves. Any entities that are taken to exist should bear a rele-
vant similarity to entities that characterize our best (or ideal) physical
theories, their coming-to-be should be intelligible in light of the naturalist
causal story, and they should be knowable by scientific means.

The naturalist epistemic attitude

As with much of modern philosophy, naturalism is primarily an expression
of an epistemic posture, specifically, a posture called scientism. In the early
1960s, Wilfrid Sellars expressed this posture when he said, “in the dimen-
sion of describing and explaining the world, science is the measure of all
things, of what is that it is, and of what is not that it is not.”⁸ Steven
Wagner and Richard Warner claim that naturalism is “the view that only
natural science deserves full and unqualified credence.”⁹ Contemporary
naturalists embrace either weak or strong scientism. According to the
former, nonscientific fields are not worthless nor do they offer no intellec-
tual results, but they are notably inferior to science in their epistemic
standing and do not merit full credence. According to the latter, unquali-
fied cognitive value resides in science and in nothing else. Either way, natu-
ralists are extremely skeptical of any claims about reality that are not
justified by scientific methodology in the hard sciences. For example, that
methodology is a third-person one and the entities justified by it are capa-
ble of exhaustive description from a third-person perspective. Entities that
require the first-person perspective as the basic mode of epistemic access to
them are to be met with skepticism.

Naturalists believe that they are justified in this posture because of the
success of science *vis-à-vis* other fields of inquiry. In addition, some natu-
ralists justify this standpoint by appealing to the unity of science, though
this argument is employed less frequently today than it was a few decades
ago. For example, in the late 1970s, Roy Bhaskar asserted that “*naturalism*
may be defined as the thesis that there is (or can be) an essential unity of
method between the natural and the social sciences.”¹⁰ Moreover, as John
Searle notes, since for these naturalists science exhausts what we can know,
then belief in the unity of science turns out to be a belief in the unity of all
knowledge because it is scientific knowledge:

Every fact in the universe is in principle knowable and understandable
by human investigators. Because reality is physical, and because science

concerns the investigation of physical reality, and because there are no limits on what we can know of physical reality, it follows that all facts are knowable and understandable by us.¹¹

For such naturalists, the exhaustive or elevated nature of scientific knowledge entails that either the only explanations that count or the ones with superior, unqualified acceptance are those employed in the hard sciences.¹²

We have seen that scientism is the core epistemic posture of the contemporary naturalist. From this core commitment, at least three philosophical theses follow that in one way or another elaborate the epistemic and methodological constraints for philosophy that are part of taking the naturalistic turn. First, there is no such thing as first philosophy. According to David Papineau, there is a continuity between philosophy and natural science:

the task of the philosophers is to bring coherence and order to the total set of assumptions we use to explain the natural world.

The question at issue is whether *all* philosophical theorizing is of this kind. Naturalists will say that it is. Those with a more traditional attitude to philosophy will disagree. These traditionalists will allow, of course, that some philosophical problems, problems in *applied* philosophy, as it were, will fit the above account. But they will insist that when we turn to “first philosophy,” to the investigation of such fundamental categories as thought and knowledge, then philosophy must proceed independently of science. Naturalists will respond that there is no reason to place first philosophy outside of science.¹³

Second, the naturalist epistemology generates intellectual pressure to employ epistemology—or language which has become the surrogate epistemology for many naturalists—to deflate, eliminate or reduce ontological matters that are *prima facie* philosophical and not scientific to epistemic or linguistic ones. Thus, Paul Churchland replaces the first-person qualitative ontology of pain with a physicalist substitute because the latter is more epistemically acceptable from a naturalist perspective: the former is derived from a discredited first-person knowledge by acquaintance, whereas all one needs for explaining our notion of pain is linguistic know-how regarding the term “pain” discernible in terms of the third-person perspective.¹⁴ Keith Campbell reduces the ontological notion of “abstract” (not existing in space or time) in debates about abstract objects to an epistemic notion (attending to a property-instance by disregarding a number of features in its noetic environment; an abstract particular is one that is brought before the mind by certain acts of noticing and disregarding).¹⁵ David Papineau undercuts and, thus, places a substantial burden of proof on dualist ontological claims they take to be descriptive reports of events and properties with which they are introspectively acquainted. By contrast, Papineau employs a version of epistemic methodism to establish the terms of debate

6 *The epistemic backdrop*

and thereby dismiss such dualist claims: A claim about mental entities is justified if and only if it is required by the categories of an ideal physics.¹⁶

These are not isolated incidents. They are expressions of a proper understanding of the impact on ontology of the naturalist epistemology and its certified creation account by three of naturalism's brightest lights. It is widely recognized that the sorts of naturalistic *reduction*—Nagelian and linguistic – that flourished from the 1930s to the 1960s was intellectually motivated by epistemic, specifically positivistic, concerns. Moreover, while positivism has gone the way of the dodo, its epistemic (or linguistic) remains are what underwrite the sorts of ontological moves illustrated by the naturalist triumvirate above.

Third, scientific theories that are paradigm cases of epistemic/explanatory success, e.g. the atomic theory of matter and evolutionary biology, employ combinatorial modes of explanation. Thus, any process that constitutes the Grand Story and any entity in the naturalist ontology should exhibit an ontological structure analyzable in terms that are isomorphic with such modes of explanation. Perhaps more than anyone, Colin McGinn has defended this idea along with what he takes it to entail, viz., the inability of naturalism to explain genuinely unique emergent properties:

Can we gain any deeper insight into what makes the problem of consciousness run against the grain of our thinking? Are our modes of theorizing about the world of the wrong shape to extend to the nature of mind? I think we can discern a characteristic structure possessed by successful scientific theories, a structure that is unsuitable for explaining consciousness. . . . Is there a “grammar” to science that fits the physical world but becomes shaky with applied to the mental world?

Perhaps the most basic aspect of thought is the operation of *combination*. This is the way in which we think of complex entities as resulting from the arrangement of simpler parts. There are three aspects to this basic idea: the atoms we start with, the laws we use to combine them, and the resulting complexes . . . I think it is clear that this mode of understanding is central to what we think of as scientific theory; our scientific faculty involves representing the world in this combinatorial style.¹⁷

We have looked at a number of philosophers who express different aspects of the naturalist epistemic attitude. Let us now turn to an overview of the naturalist's view of how things came to be.

The naturalist Grand Story

The naturalist has an account of how all things whatever came to be. Let us call this account the Grand Story. The details of the Grand Story need not concern us here. Some version of the Big Bang is the most reasonable

view currently available. On this view, all of reality—space, time, and matter—came from the original “creation” event and the various galaxies, stars, and other heavenly bodies eventually developed as the expanding universe went through various stages. On at least one of those heavenly bodies—earth—some sort of pre-biotic soup scenario explains how living things came into being from non-living chemicals. Moreover, the processes of evolution, understood in either neo-Darwinian or punctuated equilibrium terms, gave rise to all the life forms we see including human beings. Thus, all organisms and their parts exist and are what they are because they contributed to (or at least did not hinder) the struggle for reproductive advantage, more specifically, because they contributed to the tasks of feeding, fighting, fleeing, and reproducing.

There are four important things to note about the Grand Story. First, at the core of the Grand Story are two theories that result from combinatorial modes of explanation: the atomic theory of matter and evolutionary theory. If we take John Searle to be representative of naturalists here, this means that causal explanations, specifically, bottom-up but not top-down causal explanations, are central to the (alleged) explanatory superiority of the Grand Story.¹⁸

Second, it is an expression of a scientific version of philosophical monism. According to this view, everything that exists or happens in the world is susceptible to explanations by natural scientific methods. Whatever exists or happens in the world is natural in this sense. *Prima facie*, the most consistent way to understand naturalism in this regard is to see it as entailing some version of strong or strict physicalism: everything that exists is fundamentally matter, most likely, elementary “particles” (whether taken as points of potentiality, centers of mass/energy, units of spatially extended stuff/waves or reduced to/eliminated in favor of fields), organized in various ways according to the laws of nature.¹⁹ By keeping track of these particles and their physical traits, we are keeping track of everything that exists. No non-physical entities exist, including emergent ones. This constitutes a strong sense of physicalism. When naturalists venture away from strong physicalism, however, they still argue that additions to a strong physicalist ontology must be depicted as rooted in, emergent from, dependent upon the physical states and events of the Grand Story. Whether or not emergent properties should be allowed in a naturalist ontology will occupy our attention throughout this book.

Third, the Grand Story is constituted by event causality. It eschews both irreducible teleology and agent causation in which the first relatum of the causal relation is in the category of substance and not event. Moreover, the Grand Story is deterministic in two senses: diachronically such that the state of the universe at any time t coupled with the laws of nature determine or fix the chances for the state of the universe at subsequent times; synchronically such that the features of and changes regarding macro-wholes are dependent on and determined by micro-phenomena.

Finally, while some naturalists eschew questions about the nature of existence itself, others have formulated an analysis of existence—sometimes called the Eleatic principle of existence—based on a naturalist epistemology and consistent with the Grand Story. Thus, Bruce Aune defines *a exists* as “a belongs to the space-time-causal system that is our world. Our world is, again, that system of (roughly) causally related objects.”²⁰ Along similar lines, D. M. Armstrong says that for any entities, the following question settles the issue of whether or not those entities can be said to exist: “Are these entities, or are they not, capable of action upon the spatio-temporal system? Do these entities, or do they not, act in nature?”²¹ Daniel Dennett claims that when we are trying to find out whether or not some entity like the self exists, what we must do is locate the entity within the causal fabric.²² Keith Campbell applies the same reasoning to the question of the existence of emergent entities like social characteristics by claiming that the test of their existence turns on their ability to exhibit independent causality because “power has been recognized as the mark of being.”²³ Finally, Jaegwon Kim says, “Causal powers and reality go hand in hand. To render mental events causally impotent is as good as banishing them from our ontology.”²⁴ The sort of causal power characteristic of the entities most consistent with the processes, properties, relations and particulars that constitute the Grand Story is passive liability and not active power.

The naturalist ontology

Weak vs. strong naturalism

In order to characterize a naturalist ontology, we must distinguish global vs. local naturalism and weak vs. strong naturalism. Roughly, global naturalism is the view that the spatio-temporal universe of natural entities studied by science is all there is. Global naturalists (e.g. Wilfrid Sellars) reject abstract objects of any kind, including traditional realist properties. Local naturalists (e.g. Jeffrey Poland) either accept or are indifferent towards abstract objects but they insist that the spatio-temporal universe consists only of entities studied by the natural sciences. Local naturalists reject Cartesian souls, Aristotelian entelechies, and so forth.

A distinction also exists between strong and weak naturalists. Strong naturalists (e.g. David Papineau) accept a strong version of physicalism (all individuals, events, states of affairs, properties, relations and laws are entirely physical) for the natural world, while weak naturalists (e.g. John Searle) embrace various emergent entities. Elsewhere, I have argued against global naturalism because naturalists should eschew universals (construed as abstract objects) and other abstract objects yet such entities do, in fact, exist.²⁵ The debate about global naturalism will not be of concern in this book, except in chapter five in connection with evaluating McGinn’s

rejection of certain dualist rivals to his attempt to reconcile consciousness with naturalism.

The location problem

For our purposes, it is important to say a bit more about criteria for naturalist ontological commitments. A good place to start is with what Frank Jackson calls the “location problem.”²⁶ According to Jackson, on the basis of the superiority of scientific ways of knowing exemplified by the hard sciences, naturalists are committed to a fairly widely-accepted physical story about how things came to be (the Grand Story) and what they are. The location problem is the task of locating or finding a place for some entity (for example, semantic contents, mind and agency) in that story.

For Jackson, the naturalist must either locate a problematic entity in the basic story or eliminate the entity. Roughly, an entity is located in the basic story just in case it is entailed by that story. Otherwise, the entity must be eliminated. At this point, it is worth recalling that Kim and others have complained that one does not *explain* a phenomenon by labeling it supervenient. Likewise, one might think that someone hasn’t really “located” a puzzling phenomena if all one has done is point out that it necessarily covaries with this or that sort of physical phenomenon. In any case, Jackson provides three examples of location. First, just as density is a different property from mass and volume, it is not an additional feature of reality over and above mass and volume in at least this sense: an account of things in terms of mass and volume implicitly contains, i.e. entails the account in terms of density. Second, Jones being taller than Smith is not an additional feature of reality over and above Jones’ and Smith’s heights because the relational fact is entailed, and in this sense located by the latter.

More importantly, Jackson focuses on the location of macro-solidity. He acknowledges that prior to modern science there was a widely accepted commonsense notion of macro-solidity, viz., being everywhere dense. However, due to modern science, this notion has been replaced with being impenetrable. So understood, macro-solidity may be located in the basic micro-story: given a description of two macro-objects in terms of their atomic parts, lattice structures, and sub-atomic forces of repulsion, this description entails that one macro-object is impenetrable with respect to the other.

Jackson believes there are four important sorts of troublesome entities that the naturalist must locate: mental properties/events, facts associated with the first-person indexical, secondary qualities, and moral properties. Focusing on mental properties/events, Jackson claims that the naturalist must argue that they globally supervene on the physical. He unpacks this claim with two clarifications. First, he defines a minimal physical duplicate of our world as “a world that (a) is exactly like our world in every

physical respect (instantiated property for instantiated property, law for law, relation for relation), and (b) contains nothing else in the sense of nothing more by way of kinds or particulars than it *must* to satisfy (a).²⁷ Second, he advocates B*: “Any world which is a *minimal* physical duplicate of our world is a psychological duplicate of our world.”²⁸

Jackson concludes in this way:

Let ϕ be the story as told in purely physical terms, which is true at the actual world and all the minimal physical duplicates of the actual world, and false elsewhere; ϕ is a hugely complex, purely physical account of our world. Let ψ be any true sentence which is about the psychological nature of our world in the sense that it can only come false by things being different psychologically from the way they actually are: every world at which ψ is false differs in some psychological way from our world. Intuitively, the idea is that ψ counts as being about the psychological nature of our world because making it false requires supposing a change in the distribution of psychological properties and relations. . . . [E]very world at which ϕ is true is a world at which ψ is true—that is, ϕ entails ψ .²⁹

The logic of the mereological hierarchy

Jackson grasps the connection between accepting the epistemic superiority of naturalism and deciding between weak and strong naturalism. For Jackson, if naturalism is to have superior explanatory power, this entails strong naturalism. Jackson correctly understands that there are at least three constraints for developing a naturalist ontology and locating entities within it:

- a Entities should conform to the naturalist epistemology.
- b Entities should conform to the naturalist Grand Story.
- c Entities should bear a relevant similarity to those found in chemistry and physics or merely be capable of one-to-one or one-to-many correlation with entities in chemistry or physics or be shown to depend necessarily on entities in chemistry and physics.

Regarding the naturalist epistemology, all entities should be subject to combinatorial modes of explanation and be entirely publicly accessible and, thus, should be knowable entirely by third-person scientific means. Regarding the Grand Story, one should be able to show how any entity *had* to appear in light of the naturalist event-causal story according to which the history of the cosmos amounts to a series of events governed by natural law in which micro-parts come together to form various aggregates with increasingly complex physical structures.

As we shall see in subsequent chapters, these constraints seem to rule out the existence of genuinely emergent properties. When construed in terms of emergent properties, the second disjunct of (c) “solves” the so-called hard problem of consciousness, the explanatory gap, by simply naming the problem and dismissing the need for a naturalist to do any further explanatory work. For many philosophers, including many naturalists, this strategy is inadequate. The second disjunct also suffers from the difficulty of justifying the existence of *sui generis* emergent entities in light of criteria (a) and (b). The third disjunct of (c) suffers from this latter problem and from difficulties with justifying the claim that emergent entities are “necessitated” by their subvenient physical bases. Clarifying and defending these claims are central to the desiderata of this book. But it may be useful at this stage of reflection to show how (a) and (b) justify the standard layered mereological hierarchy as the proper naturalist ontology.

Let us construe this hierarchy in terms of individual entities and properties rather than in terms of concepts or linguistic descriptions. So understood, the standard mereological hierarchy consists in an ascending level of entities in the category of individual such that for each level above the ground level of elementary micro-physics (at which entities have no further physically significant separable parts), wholes at that level are composed of the separable parts at lower levels. Thus, from bottom to top we get micro-physical entities (strings, waves, particles, fields?), sub-atomic parts, atoms, molecules, cells, living organisms, and so on. The relationship between individuals at level n and $n+1$ is the part/whole relation. But there are two kinds of parts relevant to the hierarchy—separable and inseparable.

p is a separable part of some whole $W = \text{def. } p$ is a particular and p can exist if it is not a part of W .

p is an inseparable part of some whole $W = \text{def. } p$ is a particular and p cannot exist if it is not a part of W .

In contemporary philosophy, inseparable parts were most fruitfully analyzed in the writings of Brentano, Husserl and their followers.³⁰ The paradigm case of an inseparable part in this tradition is a (monadic) property-instance or relation-instance. The mereological hierarchy explicitly employs and only employs separable parts.

For present purposes, there is an important point to make about the hierarchy in the category of individual (and as we will see below, the category of property): *The “hierarchy” is not really a hierarchy.* There is no ascending, no going up anything. Rather, the levels form spatio-temporally wider and wider wholes. So we should think of the “hierarchy” as going out, not up. To see this, consider the relationship between a water molecule and its constituent atoms. There are two ways to analyze the water molecule.

First, we may adopt the eliminativist line of Peter van Inwagen and Trenton Merricks and eliminate water molecules in favor of certain collections

12 *The epistemic backdrop*

of proper separable parts arranged waterwise.³¹ Let “the p’s” stand for all and only the atomic simples that constitute such an arrangement. The difference between the p’s and the p’s being arranged waterwise is that the latter includes and the former does not include a relational structure. Elsewhere, I have defended a realist constituent ontology for properties/relations and their instances.³² Nothing of importance here turns on this assumption, so for ease of exposition, let us grant it. It follows that the structure itself is nothing but a combination of relation-instances that stand between and among and only between and among the p’s. On this view, it becomes clear that the molecule is not riding on top of anything because there is no such thing as a water molecule. Moreover, the-p’s-arranged-waterwise is not something such that its relationship to its separable parts (the p’s) is top/down. Rather, we have a relational structure that is spatio-temporally wider than any of its constituent proper separable parts. So if there is any hierarchy at all, that hierarchy does not move up and there is no top or down; it moves out: there are wider and wider relational arrangements in focus.

Second, we may resist the eliminativist line and provide an analysis of the water molecule that grounds its unity such that it is not exhaustively decomposable into a relevant collection of parts and an individuated structure. There are two main ways of doing this. The first is to attribute to the water molecule a particular sort of inseparable part called a boundary or surface. Arguably, entities at the microphysical level have boundaries/surfaces, so one may employ such an entity at “higher” levels of analysis without bloating one’s basic ontological categories. By providing a metaphysical limit such that entities within the boundary or surface are constituents of the whole and those outside the boundary are not, it unifies and reifies the water molecule such that it cannot be exhaustively decomposed along eliminativist lines.³³

The second is to attribute to the molecule an individuated essence of its own, e.g. a substantial form construed as an abstract particular, an individuated essence construed as an *infimae species*, or an haecceitas construed as a Leibnizian essence. There are at least two reasons why there is pressure for naturalists to reject this solution. First, this alternative commits one to a shopping-list approach to metaphysics in which one *sui generis* entity after another is added to one’s ontology. Any self-respecting materialist should reject this bloated ontology because it is not simple and it means that in principle an ideal physics does not exhaustively carve the cosmos at its most basic joints. To be sure, there are several versions of *reduction*, but now is not the time to discuss them. However, if understood as the drive to keep one’s ontology to a minimum in light of the Grand Story, the spirit of *reduction* is at the heart of naturalism. Second, such an ontological pluralism and its voluminous list of brute facts stretch the explanatory resources of the naturalist epistemology and Grand Story beyond the breaking point.

Note carefully that such individuated essences are neither physical (they are not constituents over which physical theory quantifies) nor entailed by the Grand Story. Instead, they are metaphysical entities added to the naturalist ontology to solve distinctively philosophical problems for which science is silent. For example, to solve problems of individuation, Jackson acknowledges that it may well be the case that,

The physicalist will need to require that minimal physical duplicates of our world be ones which, in addition to being identical in respect of physical properties, laws, and relations with our world, are identical in which *haecceities* are associated with which physical properties, laws, and relations.³⁴

For Jackson, *haecceities* are entities not included in the physical description of the world. Such individuated essences were central to the Medieval Great Chain of Being, which remains the paradigm case of pluralistic shopping-list ontology, an ontology completely rejected with the rise of the mechanical philosophy.

If this is correct, then the two main naturalist alternatives for analyzing “higher level” individuals should be the eliminativist or the boundary/surface alternatives. On each alternative, “higher level” wholes are wider than and not higher than their constituent parts, and those wholes are capable of combinatorial explanation ultimately in terms of the micro-physical level in keeping with the naturalist epistemology and Grand Story. No non-structural, *sui generis* entities of which the lowest level is bereft are needed, and we have a macro-ontology in which entities are differentiated by individuated relational structures in keeping with ontological constraints (a) and (b) above.

So much for the category of individual. What about the category of property (and relation)? Are there ontological constraints for what sorts of properties a naturalist should include in the hierarchy? I believe there are, and to get at those constraints, note that as typically presented, the hierarchy entails the causal closure of the basic microphysical level along with the ontological dependence of entities and their activities at supervenient levels on entities and their activities at that basic level.

Causal closure and the related matter of top/down causation are controversial, and not all naturalists accept the former or reject the latter. But closure and a rejection of top/down causation are hard for a naturalist to avoid. As Kim nicely states, the basic naturalist argument in favor of causal closure is that if it is rejected, then

you are ipso facto rejecting the in-principle completeability of physics—that is, the possibility of a complete and comprehensive physical theory of all physical phenomena. For you would be saying that any complete

explanatory theory of the physical domain must invoke nonphysical causal agents. . . . It is safe to assume that no serious physicalist could accept such a prospect.³⁵

The so-called “completeability of physics” is not an arbitrary postulate in a naturalist worldview. It follows quite naturally once one understands the Grand Story. That Story gives the naturalist an account of how all things have come-to-be, and the Story’s account is one according to which one begins at the Big Bang with a small number of physical entities and explains the origin and behavior of everything else in terms of the laws of physics and new combinations of micro-physical entities. The Story itself gives pride of place to micro-physical entities and it is bottom/up at its core. The completeability of physics is essential to the plausibility of the naturalist creation myth.

The causal closure principle is not arbitrary from a naturalist perspective nor is it an additional postulate that naturalists are intellectually free to reject. It follows from the combinatorial mode of causal explanation and the Grand Story’s commitment to the sort of micro-macro constitution and determination at the core of the atomic theory of matter, evolutionary biology and other central theories of how things have come to be. As we shall see in later chapters, if a naturalist rejects closure he or she will have to accept *sui generis*, contingent brute facts. In turn, this undermines the claim that a naturalist worldview is superior to rivals because it can explain how all things have come to be.

There are two different ways to understand causal closure: (1) Every physical event has a physical cause. (2) No physical event has a non-physical cause. The Grand Story, as Kim correctly notes, implies a complete and comprehensive theory of the physical domain and, indeed, the cosmos and everything that occurs within it. Thus, the Grand Story provides a reason for preferring (2) to (1).

Besides closure, a related issue for deciding what sorts or properties should populate the hierarchy is the problem of top/down causation. I believe there is severe intellectual pressure that follows from the nature of naturalism itself for rejecting top/down causation for genuinely emergent *sui generis* properties. Moreover, the only way to save top/down causation is to reduce it to outside/in causation that occurs with respect to structural wholes at the same level as their parts via causal feedback. I also think that the price to be paid for retaining causal laws in the special sciences is to disallow emergent properties and allow only microphysically based structural properties constituted by microphysical parts, properties and relationships. If this is right, it follows that an adequate treatment of these desiderata (to preserve “top/down” causation and causal laws in the special sciences) entails that a naturalist ontology constituted by the standard mereological hierarchy can countenance structural wholes in the category of individual and structural supervenient properties in the category of

property. But note that it cannot countenance genuine emergent properties, especially causally active emergent properties. All emergent properties, if such there be, must be epiphenomenal.

Before proceeding, I want to clarify the difference between emergent and structural properties and supervenience in a way that is apt for what follows. An emergent property is a unique, new kind of property different from those that characterize its subvenient base. Accordingly, emergent supervenience is the view that the supervenient property is a simple, intrinsically characterizable, novel property different from and not composed of the parts, properties, relations, and events at the subvenient level. We may characterize “novel” as follows:

Property P is a *novel* emergent property of some particular x at level I_n just in case P is an emergent property, x exemplifies P, and there are no determinates P' of the same determinable D as P such that some particular at a level below I_n exemplifies P or P'.³⁶

By contrast, a structural property is one that is constituted by the parts, properties, relations, and events at the subvenient level. A structural property is identical to a configurational pattern among the subvenient entities. It is not a new kind of property; it is a new pattern, a new configuration of subvenient entities. In addition, many philosophers would characterize emergent and structural supervenience as causal and constitutive, respectively. Since I am contrasting emergent and structurally supervenient properties, I will use the notion of an emergent property as simply being a novel, *sui generis* property.

For two reasons, if we assume that mental properties are genuinely emergent *sui generis* qualities, then given the mereological hierarchy and its disallowance of top/down causation, the existence of emergent mental properties presents at least two problems for naturalism. First, for those naturalists who accept a causal criterion of existence, emergent mental properties are epiphenomenal and, thus, do not exist. One is then faced with a dilemma: Either one accepts phenomenal consciousness, which construes emergent mental properties along familiar lines as what-it-is-like to be such and such and rejects causal closure *or* one retains closure and rejects phenomenal consciousness because it is epiphenomenal. In my view, the latter option is the correct one for a naturalist to take. Nevertheless, irrespective of whether I am right, in subsequent chapters, we will be examining only versions of naturalism that accept emergent mental properties. For naturalists of this stripe, the problem of epiphenomenalism must be addressed in an adequate way.

Second, it is obvious that mental states are causal factors in our behavior. It is hard to see how knowledge and agency can be salvaged if this is denied. Indeed, if an analysis of mental states entails epiphenomenalism, this is widely recognized as fodder for a *reductio* against that analysis. This

is why most naturalists think that the only way to save mental causation is in one way or another to identify it with the physical.

More than any other naturalist, Jaegwon Kim has pressed the problem of top/down mental causation for naturalism.³⁷ Kim correctly notes that the problem of mental causation arises from the very nature of physicalism itself, and not from a Cartesian view of mental substance and, indeed, mental causation is a difficulty in the category of property every bit as much as in the category of substance.

Kim's supervenience argument (a.k.a. the exclusion argument) purportedly shows that, given the irreducibility of the mental, there can be no mental causation in a world that is fundamentally physical, and according to Kim, this raises serious problems regarding cognition and agency, two features of our lives that are hard to give up. The supervenience argument, says Kim, may be construed to show that mental causation is inconsistent with the conjunction of four theses: (1) closure; (2) exclusion (no overdetermination); (3) supervenience (not construed simply as property covariance, but taken to entail dependence and synchronic determination); and (4) mental irreducibility. The fundamental idea of the supervenience argument is that "vertical determination excludes horizontal causation." To see this, Kim invites us to consider two physical events, p and p^* , along with two mental events, m and m^* such that (1) m and m^* supervene on p and p^* respectively (where supervenience includes the notion of dependence and determination, even if this is not taken to be efficient causality) and (2) p causes p^* .

The argument proceeds in two stages. Stage 1: Focus on m to m^* causation. Since m^* obtains in virtue of p^* , if m is going to cause m^* it must do so by causing p^* . Stage 2: Kim offers two different ways to complete the argument. Completion 1: Assuming causal closure and exclusion (no causal overdetermination), p will be the cause of p^* and there is no room for m to be involved in bringing about p^* . We have m and m^* supervening on p and p^* , respectively, and p causing p^* , nothing more and nothing less. Completion 2: Granting that m causes m^* by causing p^* , it follows that m causes p^* . By irreducibility, we have $m \neq p$. So m and p cause m^* . By exclusion and closure, m is ruled out and p is selected as the only cause for p^* . Completion 2 avoids reference to supervenience. On either way of completing stage two, we have $m \rightarrow m^*$ and $m \rightarrow p^*$ giving way to $p \rightarrow p^*$.

Some have not been persuaded by Kim's argument, though I am not among them.³⁸ More importantly, we have already examined reasons for why a naturalist should accept closure and reject top/down causation that follow naturally from the naturalist epistemic attitude and Grand Story. I will argue below that there are additional reasons for accepting closure and rejecting top/down causation if we limit our focus to emergent properties. Within this limitation, I believe that there are strong reasons to hold that top/down causation is disallowed by a naturalist view that entails the standard mereological model. If this is so, then mental causation can

obtain only if the mental is not emergent but, rather, in some way or another identified with the physical. However, there have been counterexamples offered that allegedly show that top/down causation is real and consistent with the standard hierarchy. Roger Sperry offered a paradigm case of such a counterexample.³⁹ According to Sperry, there can be top/down causation without disrupting or intervening in the causal relations or micro-interactions at the elementary level. Such top/down causation does not disrupt the laws or behavior of entities at that level. Here is his example:

A molecule within the rolling wheel, for example, though retaining its usual inter-molecular relations within the wheel, is at the same time, from the standpoint of an outside observer, being carried through particular patterns in space and time determined by the over-all properties of the wheel as a whole. There need be no “reconfiguring” of molecules relative to each other *within the wheel itself*. However, *relative to the rest of the world* the result is a major “reconfiguring” of the space-time trajectories of all components in the wheel’s infrastructure.⁴⁰

Unfortunately, for two reasons this analogy fails as an example of real top/down causation of emergent properties as we are assuming mental properties to be. Note first that the “over-all properties of the wheel as a whole” are not emergent properties caused by and on top of subvenient entities. They are mere structural complexes constitutively supervenient “on,” that is, constituted by base entities. What we have is a wider relational context of molecules than would obtain as inter-molecular relations among a small set of the wheel’s constituents if the larger context were annihilated. But the wheel and such a set of molecules are at the same level. Regarding the wheel’s effect on some specific molecule, we have outside/in causation, not top/down. Second, the “reconfigured” space-time trajectory of some specific molecule is not an emergent property at all, at least not in the way we are treating mental properties. Mental properties are emergent in the sense that they are genuinely new *kinds* of properties that in no way characterize the base level. However, the “reconfigured” trajectory is merely a new combination of spatial and temporal properties that already characterize the base. The wheel case is merely an example of outside-in causal interaction due to wider structural relations at the base level. I believe similar problems beset other alleged cases of top/down causation.

There is also a problem with emergent causal powers and laws in the special sciences. To get at this problem, let us begin by examining Kim’s treatment of the generalization problem.⁴¹ According to Kim, if we grant property/event dualism, then given causal closure and the rejection of overdetermination, mental top/down causation cannot occur. He also argues that causal closure is not needed for the rejection of mental causation. If we reject the notion that there cannot be two sufficient efficient causes for some physical event that is, in fact, caused, and grant for the

sake of argument that there is a mental cause for some physical event, then we have two competing causal stories. Why? Given the Grand Story and the assumption that the physical event in focus was caused, there is a clear story about antecedent physical causes for the event. However, there is no room for the higher mental story to be efficacious given that the supervenience of the higher story on the micro-physical story entails an ontological dependence and determination of the former on the latter. Thus, eschewal of top/down mental causation follows merely from supervenience (taken to include the dependence and determination of the emergent mental event on its physical base) and a rejection of overdetermination (and from the irreducibility of the mental).

But now we face the generalization problem: Does the causal impotency of higher-level mental properties/events threaten other higher-level properties in the special sciences, e.g. chemical, geological, biological properties? Kim offers a response to the generalization problem that he takes to be adequate for saving the causal powers in focus in the special sciences. His solution begins by rejecting the idea that the mind/brain supervenience relationship which renders the mental epiphenomenal is the same as the relationship between higher-level properties in special sciences like geology and lower-level micro-physical properties. Kim draws a distinction between different levels up the hierarchy and different orders within a single level. Now, if we focus on the supervenience of the mental on brain states, we are actually depicting higher-order mental properties as structural functional properties (for Kim, concepts) that are realized by brain states. In this case, the realization relation is the converse of the supervenience relation and we have either mere conceptual supervenience, structural supervenience or both.

So understood, the realization relation is different from the micro-macro relation, and the realization relation does not track up or down the levels of the hierarchy as does the micro-macro relation. The realization relation stays within one level and higher-order mental properties are at the same level as their physical realizers. Given that the supervenience relation between mental and brain properties and its converse (the realization relation) render mental properties causally impotent, says Kim, this problem is one of higher and lower orders within one level. Thus, it does not generalize to the macro-micro relation that connects different levels up the hierarchy and that is the relevant relation between special science properties and microphysical ones. For an object at a macro-level to have a micro-based property such as being water is for the property to be identical to a set of parts having their properties and standing in certain relations to each other. Being ten kilograms is a micro-based property of a table, says Kim, and it is causally efficacious. In the same way, micro-based properties of earthquakes are such that earthquakes cause things to occur.

Is Kim's solution to the generalization problem successful? If the generalization problem is understood in the specific way Kim frames it, it may

well be. The distinction between intra-level higher order functional properties and lower order realizers on the one hand and inter-level micro-based properties along with the micro-macro relation that (allegedly) tracks up the hierarchy is a clear one. Given this distinction, the generalization from mental property causal impotence to properties in various special sciences is blocked.

However, in the context of *sui generis* mental properties and epiphenomenalism, Kim's solution comes at a price. It should be clear that his employment of micro-based properties entails that they are structural and not *sui generis* emergent properties. As such, they are exhaustively decomposable into parts, properties and relations at the subvenient level. These micro-physical constituents constitute micro-based properties. So construed, they may have "new" causal powers in the sense of additive sums of constituent powers or due to a new spatial shape resulting from a new arrangement of subvenient entities, but there are no new *kinds* of causal powers.

Moreover, as I have argued above, the macro-micro relation does not travel up anything. Rather, it ventures outward at the same level, including wider and wider relational structures. If the properties of the special sciences are emergent, then for the reasons we have investigated, they are epiphenomenal. *It is the hierarchy itself, along with the ontological dependency and determination of higher-level novel emergent properties on their subvenient bases that rules out top/down causation. And the hierarchy is not arbitrary for naturalism. The ontology flows out of the Grand Story, which, in turn, is certified by the hard sciences.* Novel emergent properties are like shadows produced by a flashlight. Moreover, if one accepts physical closure and rejects causal overdetermination, we have further reasons for rejecting the top/down causation of novel emergent properties.

Before we summarize our discussion of the naturalist ontology, I want to mention one final constraint on the sorts of properties it should include. If we limit ourselves to macro-properties, an appropriate limitation because consciousness is a macro-feature (except for certain versions of panpsychism—see chapter six), then the following principle seems to be *prima facie* justified:

Principle of Naturalist Exemplification (PNE): $(x) Px \rightarrow Ex$

P stands for any property whatever and E stands for the property of being extended. Moreover, x ranges over and only over property-instances. Elsewhere I have defended a constituent ontology in which property-instances are complex entities, and I shall merely assume this ontology here.⁴² According to this ontology, when some concrete particular e exemplifies a property P , then the-having-of- P -by- e is a property-instance that is modally distinct from both P and e . X is neither identical to P nor e . So

understood, property-instances are certain sorts of states of affairs and, moreover, if the instantiation of P by e is temporal, then the property-instance becomes an event.

Note that P and e are constituents of x. If we focus on paradigm cases that satisfy PNE, it becomes reasonable to hold that the spatial extension of x is grounded in, obtains in virtue of the spatial extension of e. For example, when an apple is red, the-having-of-red-by-the-apple is a property-instance spread out through the extended region occupied by the apple. It is in virtue of the apple's extension that the particular instance of red is extended. This may be seen, for example, by noting that it is because the apple has a particular shape that its instance of red has that shape as well.

PNE says that if a property in the naturalist ontology is to be exemplified, then a necessary condition is that both the concrete particular that exemplifies P and the property-instance that results have spatial extension.

PNE seems to capture nicely the wide range of properties in macro-physics, chemistry, geology, neuro-science, and so forth. It could be objected that PNE fails because certain entities, e.g. some quantum entities are or the point particles of Roger Boscovich were unextended, and provide sufficient counter examples to PNE. I do not think this objection works. Regarding quantum entities, there are at least eight different empirically equivalent philosophical models of quantum reality and, at this stage, it is irresponsible to make dogmatic claims about the ontology of the quantum level.⁴³ Moreover, since I have limited PNE to the macro-level, we may set aside the quantum world for our purposes. Regarding entities such as Boscovich's point particles, rather than conclude that they are counter examples to PNE, their lack of spatial dimensionality may be taken as a *reductio* against them. And, indeed, this is how the history of physics ran. Boscovich's point particles fit more easily into a spiritualist ontology, e.g. Berkeley's, than in a straightforward version of materialism, and like action at a distance, they were rejected.

There is a debate about whether individual mental states such as pains and thoughts are extended. I cannot enter that debate here. Nevertheless, based on PNE, if it turns out that mental states are not extended, then PNE banishes them and their constituent properties from a naturalist ontology. In this case, PNE counts against any naturalist ontology that quantifies over emergent mental properties.

It is time to summarize what a naturalist ontology should look like.⁴⁴ In the category of individual, if we reject an eliminativist strategy, then all wholes "above" the microphysical level are structural, relational entities constituted by the parts, properties and relations at the microphysical level. Such wholes stand in a constituent/whole relation to these microphysical entities and are actually wider entities at the basic level. This is true whether we adopt an eliminativist line for such wholes (and accept atomic simples and various arrangements of them) or add some sort of boundary or surface.

Regarding the category of property, consider the following:

Emergence₀: New features that can be deduced from base (e.g. fractal patterns).

Emergence₁: Ordinary structural properties (e.g. being water, solidity)

Emergence_{2a}: *Sui generis*, simple, intrinsically characterizable, new kinds of properties relative to base that are also epiphenomenal (e.g. being painful construed epiphenomenally).

Emergence_{2b}: *Sui generis*, simple, intrinsically characterizable, new kinds of properties relative to base with new causal powers construed as passive liabilities (e.g. being painful understood as having top/down causal liabilities).

Emergence_{2c}: *Sui generis*, simple, intrinsically characterizable, new kinds of properties with active power (e.g. being active power that characterizes most versions of agent causation).

Emergence₃: An emergent, suitably unified mental subject or I with active power.

Clearly, emergence₀ and emergence₁ fit nicely in the mereological hierarchy and conform to the naturalist epistemology (e.g. combinatorial explanation) and Grand Story. But emergence_{2a} through emergence₃ should be disallowed for reasons we have already investigated, e.g. they resist functionalization. It would seem that all a naturalist could do with them is simply to label them as contingent brute facts and assert that they are not a problem for the naturalist. We will look at different attempts to handle some of these sorts of properties in subsequent chapters. But we have already examined reasons to be highly suspicious of a naturalist view that accepts one or more of these sorts of properties and also claims that naturalism is explanatorily and epistemically superior to alternative worldviews.

Moreover, there is an increasingly heavy burden of proof on a naturalist ontology as one moves from emergence_{2a} to emergence₃. All types of emergence fall prey to previous arguments against emergent entities. Emergence_{2a} requires less justification than stronger forms of emergence because it does not require a rejection of closure. Emergence_{2b} is subject to these arguments and additional difficulties with top/down causation and causal closure. But relative to emergence_{2c} and emergence₃ it has the advantage of exhibiting the same sort of causal power—passive liability subject to law—that characterizes causal particulars at the microphysical level.

22 *The epistemic backdrop*

Emergence_{2c} has all the problems exemplified by emergence_{2b} and it also suffers from having a completely unique sort of causal power—active power—different from the causal powers that range throughout the naturalist ontology outside of agent causal events. Emergence₃ shares all of the difficulties with emergence_{2c} and it also suffers from two further facts not easily accommodated in the naturalist ontology if they are taken as irreducible and uneliminable facts about the world: First, the indexical fact associated with “I.” Second, difficulties with explaining how one can get a sort of primitive, substantial unity in which its various inseparable parts/faculties are internally related to the substantial subject from a mereological aggregate constituted by a structural arrangement of separable parts that stand in external relations to each other and their mereological whole.

D. M. Armstrong as a paradigm case naturalist

I have argued that the naturalist epistemology and Grand Story constrain the naturalist ontology and justify strong naturalism and a rejection of emergent entities. It may be worth noting that many naturalists who keep a steady eye on broader epistemological and metaphysical issues reach the same conclusion. More than anyone else, D. M. Armstrong had clearly reflected on this topic and he concludes that all entities in the naturalist ontology must be: 1) spatially located; 2) entities knowledge of which conform to an externalist causal epistemology; 3) capable of entering into causal relations; and 4) entities whose existence can be given a natural scientific causal explanation according to the Grand Story.

To illustrate these points in Armstrong, the following statement is an example where he uses a naturalist externalist epistemology to settle issues in ontology:

If any entities outside this [spatio-temporal] realm are postulated, but it is stipulated further that they have no manner of causal action upon the particulars in this realm, then there is no compelling reason to postulate them.⁴⁵

In this context, Armstrong is claiming that the only way something can interact with natural entities—including cognitive processes to be objects of knowledge—is by way of causation.

Armstrong also employs the Grand Story as a criterion for an acceptable naturalist ontology:

I suppose that if the principles involved [in analyzing and explaining the origin of or processes of change in things within the single all-embracing spatio-temporal system which is reality] were completely different from the current principles of physics, in particular if they

involved appeal to mental entities, such as purposes, we might then count the analysis as a falsification of Naturalism.⁴⁶

Elsewhere, he uses the Grand Story, coupled with scientism and epistemic simplicity to justify a strong physicalist analysis of mental entities, along with a *reduction* of secondary qualities to micro-physical entities.⁴⁷ Regarding secondary qualities, Armstrong claims that while the naturalist should hold that secondary qualities are objective features of the external world, she

cannot put them back into the world *alongside* and in *addition* to the properties that contemporary science attributes to physical objects. . . . There really is no place in the physical world for such extra properties. What we must say, rather, is that these properties are respectable, but *micro-physical*, properties of objects, surfaces, and so on. . . . The general idea is to find *micro-physical correlates* for the secondary qualities of physical objects and events and then to identify the qualities with these physical correlates.⁴⁸

Further, as a naturalist, Armstrong explicitly rejects internal relations because they cannot be spatio-temporally located and thus, are dis-analogous with other entities in the Grand Story. Their lack of spatial location also means they cannot enter into physical causal relations with the brain, which is a necessary condition to be an object of knowledge or justified belief in a naturalized externalist epistemology.⁴⁹ Armstrong clearly grasps the inner logic of naturalism.

Serious metaphysics, simplicity and emergent properties

Frank Jackson begins his attempt to develop a naturalistic account of the mental by contrasting two very different approaches to metaphysics. The first he calls serious metaphysics. Serious metaphysics is not content to draw up large pluralistic lists of *sui generis* entities. Advocates of serious metaphysics tend to approach the discipline with a prior epistemic commitment of some sort. This commitment functions as a criterion of knowledge or justified belief for quantifying over some entity and, thus, serious metaphysics usually goes hand in hand with epistemological methodism. For naturalists, this methodism expresses the various aspects of the naturalist epistemic attitude described earlier. Accordingly, serious metaphysics is primarily *explanatory* and not *descriptive* metaphysics. Thus, advocates seek to account for all entities in terms of a limited number of basic entities and in this way serious metaphysics is inherently *reductionistic*. For naturalists, these entities will constitute those at the core of the Grand Story: A property/event/object *x* exists iff it is contained within (truth functionally entailed by) the Grand Story.

The second perspective we may call a “shopping-list” approach whose primary goal is a careful description and categorial analysis of reality. Advocates of this approach tend to employ epistemological particularism, and it is no accident that Roderick Chisholm is the paradigm case of epistemological particularism and shopping-list metaphysics.⁵⁰

Jackson claims that the scientific naturalist will prefer serious metaphysics and I think he is right about this. His naturalist approach to metaphysics expresses a certain form of the principle of simplicity and provides material content for that principle of simplicity most suited for a philosophical naturalist. To see this, let us compare two versions of the principle of simplicity, an epistemic and ontological version, respectively:

Simplicity_E: Entities must not be multiplied beyond necessity.

Simplicity_O: Our ontology or preferred theory about the world should be simple.

Of course, there are various ways to state each principle, but these will do for our purposes. Simplicity_E may not be easy to apply (one rival may be simple in one respect and the other in a different respect; one rival may be simpler and the other may be more empirically accurate), but its rationale is fairly straightforward. All things being equal, if a simpler theory does the epistemic job, then the more complicated theory has baggage that serves no important epistemic function. Ontological simplicity is quite different from epistemic simplicity and some philosophers conflate the two principles. For example, Kim rightly advocates epistemological simplicity for the same reason just mentioned. But he then passes over into ontological simplicity, apparently without noticing the equivocation. After embracing “entities must not be multiplied beyond necessity,” he urges with no justification or further explanation that “we expect our basic laws to be reasonably simple, and we expect to explain complex phenomena by combining and iteratively applying these simple laws.”⁵¹

Clearly, ontological simplicity does not follow from epistemic simplicity. In fact, it sometimes happens that progress in an area of science entails adopting a more complicated ontology even though both the simpler and more complicated ontologies are epistemically simple. The shift from the simpler ideal gas equation to the more complicated van der Waals equation is a case in point. That said, I believe that the naturalist should adopt both principles of simplicity, and Kim and Jackson give the reason why. Each refers to the Grand Story (which, in turn, is justified by the naturalist epistemology) which is inherently *reductionistic*.

Moreover, if naturalism is to retain its claim to have epistemic/explanatory superiority over its rivals, then its employment of the Grand Story must be done in such a way that entities that cannot be identified with

some structural combination of fundamental microphysical entities must be eliminated. Kim and Jackson both understand this, and while Jackson seeks to carry out this way of understanding the location project, Kim recently abandoned it a few years ago.⁵² Still, Kim's appeal to ontological simplicity, every bit as much as Jackson's, provides a representative naturalist employment of the principle.

Moreover, their characterization of it provides a way of transforming the merely formal principle Simplicity_O into a related version with material content. For Kim, we begin with simple, basic laws—and presumably microphysical particulars governed by them—and allow more complex entities into one's ontology only if they are subject to combinatorial modes of explanation that involve the iterative application of the basic laws. Similarly, Jackson says one should start with the Grand Story and allow entities into one's ontology only if they are entailed by that ontology.⁵³ For Jackson, this means accepting only structural entities that are emergent₀ or emergent₁. Expressed in terms of the appropriate naturalist material principle of simplicity, we have

Simplicity_{ON} : Our ontology or preferred theory about the world should be simple in the sense that it contains the microphysical particulars, properties, relations and laws of an ideal physics or whose existence can be explained by the naturalist epistemology (e.g. combinatorial modes of explanation, are capable of exhaustive description from the third-person perspective) applied to the microphysical entities that constitute the Grand Story.

Simplicity_{ON} would seem to rule out from a naturalist ontology entities that are emergent₂ or emergent₃.

A realist view of causation and emergent properties

We have seen reasons for adopting a *prima facie* burden of proof on any naturalist ontology that includes emergent entities. If such entities are accepted, then a naturalist would owe us a causal account of their coming-to-be. In closing this chapter, it is important to get before us certain constraints on such an account. In chapters to follow, we shall look at naturalist views that seek to conform to or disregard these constraints. But these constraints seem *prima facie* justified because they follow naturally from the naturalist epistemology, Grand Story and other aspects of the naturalist ontology.

Regarding emergent properties, though some demur, at least five reasons have been proffered for the claim that causal explanations in the natural sciences exhibit a kind of causal necessity, that on a typical realist construal of natural science, physical causal explanations must show—usually by citing a mechanism—why an effect must follow given the relevant causal conditions:

- (1) Causal necessitation unpacks the deepest, core realist notion of causation, namely, causal production according to which a cause “brings about” or “produces” its effect.
- (2) Causal necessitation fits the paradigm cases of causal explanation (e.g. macro-solidity/impenetrability in terms of micro-lattice structures, repulsive forces; mass proportions in chemical reactions in terms of atomic models of atoms/molecules, bonding orbitals, energy stability, charge distribution) central to the core theories (e.g. the atomic theory of matter) that constitute a naturalist worldview and in terms of which it is purported to have explanatory superiority to rival worldviews.
- (3) Causal necessitation provides a way of distinguishing accidental generalizations or coincidences from true causal laws or sequences.
- (4) Causal necessitation supports the derivation of counterfactuals (if that chunk of gold had been placed in aqua regia, then it would have dissolved) from causal laws (gold dissolves in aqua regia).
- (5) Causal necessitation clarifies the direction of causality and rules out the attempt to explain a cause by its effect.

Three points of clarification are needed about causal necessity and the reasons for it. First, minimally, the sort of modality involved may be taken as physical necessity, a form of necessity that runs throughout possible worlds relevantly physically similar to our actual world (e.g. in having the same physical particulars, properties, relations and/or laws). Second, strong conceivability is the test that is used to judge causal necessitation (given the lattice structures and so forth of two macro-objects impenetrable with respect to each other, it is strongly inconceivable that one could penetrate the other).

Finally, principles (3)–(5) have sometimes been offered as additions to a covering law form of explanation to provide an adequate natural scientific causal explanation. Strictly speaking, a covering law “explanation” is just a description of what needs to be explained and not an explanation. However, by adding a causal model that underwrites it and that exhibits causal necessitation, the total package provides explanations for both what and why the phenomena are as they are. For brevity’s sake, below I will talk as if a covering law explanation is, in fact, an explanation, but it should be understood that when I speak of a covering law explanation I include in it an underwriting causal model.

In this chapter, we have examined limitations on a naturalist ontology that follow from naturalism itself taken as a worldview epistemically/explanatorily superior to its rivals. Let N stand for the truth of naturalism and Emergence_{2a} . In the terms of epistemic appraisal proffered by Chisholm, it seems that, $\neg (N \ \& \ \text{Emergence}_{2a})$ is at least *epistemically in the clear* where a proposition is *epistemically in the clear* provided only that subject S is not more justified in withholding that proposition than in

believing it. Alternatively, it is at least *reasonable to disbelieve* (N & Emergence_{2a}) (S is not more justified in withholding that proposition than in disbelieving it).⁵⁴

However, there are additional limits placed on the naturalist ontology when a plausible rival worldview is brought into the picture. As Timothy O'Connor points out, emergent properties, especially mental properties, must be shown to arise by way of causal necessitation from a micro-physical base if we are to "render emergent phenomena naturalistically explicable."⁵⁵ Among his reasons for this claim is the assertion that if the link between micro-base and emergent properties is a contingent one, then the only explanation for the existence and constancy of the link is a theist explanation.⁵⁶ O'Connor's claim seems to me to be correct, and to probe this matter further, we turn to an examination of the theistic argument for God's existence from consciousness.