

MIND AND LIFE: IS THE MATERIALIST NEO-DARWINIAN CONCEPTION OF NATURE FALSE?

Martin Zwick

Systems Science Graduate Program
Portland State University
Portland Oregon 97207-0751

<http://www.pdx.edu/sysc/research-systems-theory-and-philosophy>

zwick@pdx.edu

Abstract

A partial review of Thomas Nagel's book, *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature Is Almost Certainly False* is used to articulate some systems-theoretic ideas about the challenge of understanding subjective experience. The article accepts Nagel's view that reductionist materialism fails as an approach to this challenge, but argues that seeking an explanation of mind based on emergence is more plausible than one based on panpsychism, which Nagel favors. However, the article proposes something similar to Nagel's neutral monism by positing a hierarchy of information processes that span the domains of matter, life, and mind. As depicted in this hierarchy, subjective experience is emergent, but also continuous with informational phenomena at lower levels.

Keywords Consciousness; Emergence; Information; Materialism; Mind-body problem; Neutral monism; Panpsychism; Reductionism; Subjective experience; Systems theory

Introduction

This article uses a partial review of Thomas Nagel's book *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature Is Almost Certainly False* (2002) to present a systems-theoretic approach to the challenge of understanding subjective experience. From a systems perspective, Nagel's view of mind omits a necessary distinction between experience and information processing; his view of emergence is narrow; his view of neo-Darwinism as materialist is incomplete; and his view of the relationship between constitutive and historical explanations of life is unreasonable. These views are rejected and systems-theoretic ideas are offered in their stead. However, one argument of Nagel's is accepted after modification: informational processes, but not psyche, go "all the way down." This allows the neutral monism and emergence approaches to mind to be schematically integrated.

Some readers of this article may experience a tension between its review of Nagel's book and its presentation of systems ideas. Presenting these ideas is the primary aim of this article, but Nagel's important contributions to the philosophy of mind warrant the systematic assessment of his position that is also offered.

In what follows, a summary of Nagel's position is presented. A flaw in this position is pointed out, namely its omission of the necessary distinction between experience and information processing. After reviewing various conceptions of emergence, Nagel's objections to emergence as an explanation of mind are assessed. Nagel's position is also criticized because it understates the degree to which life is understood, and because, by ignoring the significance of information, it incorrectly characterizes neo-Darwinism as materialistic. The last section integrates these arguments by proposing that informational emergence occurs repeatedly in the domains of matter, life, and mind, but that subjective experience emerges only at the highest levels of these domains.

A Summary of Nagel's Position

Nagel asserts that psychophysical reduction – explaining mind in terms of matter – has been and will forever be a failure and that therefore, as the book's subtitle says, "the materialist neo-Darwinian conception of nature is almost certainly false." "Mind" here refers primarily to consciousness; more precisely, subjective experience. Nagel insists that since human (and animal) subjective experience is a salient property of life, the inability of "materialist neo-Darwinism" (MND) to explain experience should not be considered a minor and merely "local" flaw. He argues that extension of scientific theory to encompass consciousness will require radical change, so current theory is not merely incomplete, but "false." Nagel also discusses rationality and value realism, but these subjects are beyond the scope of this paper.

To situate his discussion in a broader philosophical context, Nagel reviews various approaches towards the relation of mind to matter; these are summarized in Table 1. He dismisses dualism out of hand. He also rejects idealism, the standard alternative to materialism, defined by him as (a) invoking God as ultimate reality (theism), or (b) taking ideas as more real than matter (absolute idealism), or (c) grounding reality in human experience (subjective idealism, e.g., logical positivism,¹ phenomenology,

constructivism). His main target, however, is reductionist materialism, the dominant form of monism, which he also refers to as scientific naturalism.² Rejecting dualism, idealism, and materialism, Nagel considers two alternatives: neutral monism and emergence. Between these, he prefers the former. This article criticizes this preference, and argues for the latter.

Table 1 Nagel’s summary of philosophical viewpoints on mind and matter

“Cardinality” = ontological cardinality. The currently dominant view is *materialism*; Nagel prefers **neutral monism**; this article advocates REDUCIBLE EMERGENCE.

CARDINALITY	VIEWPOINT	VARIETIES
Nagel rejects:		
Dualism	substance dualism	Cartesian dualism
Monism	<i>materialism</i> idealism	<i>reductionism</i> theism, absolute idealism, subjective idealism
Nagel considers:		
Monism (unspecified)	neutral monism EMERGENCE	panpsychism (reductive) irreducible, REDUCIBLE

Neutral monism is the view that underlying both the physical and the mental there is a neutral substrate that itself is neither physical nor mental. Nagel’s neutral monism, however, is different from this, and is identified by him as the same as panpsychism, the view that everything has aspects of both matter and mind.³ While mental properties may only be discernable in certain complex entities, they actually go “all the way down,” just like material properties, and higher-level mental properties derive from lower-level ones. Reductiveness thus applies to both material and mental properties. This is different from materialist reductionism, which reduces all properties to material ones, and regards mind as epiphenomenal or, in eliminative reductionism, as identical to its physical substrate.

In Nagel’s neutral monism, matter and mind are not merely parallel but interact causally in some unknown manner. Nagel explicitly refrains from attempting to offer any detailed content to his position. He says (pp. 57-8),

The sense in which they [the elements of the physical world] are [also] mental is so far exhausted by the claim that they are such as to provide a reductive account of how their appropriate combinations necessarily constitute conscious organisms of the kind we are familiar with. Any further consequences of their more-than-physical character at the microlevel remain unspecified by this abstract proposal.

So Nagel’s neutral monism or panpsychism has, in his own words (p.62), “only the form of an explanation without any content”; no content, that is, beyond the claim that it might someday provide the basis for a reductive account of subjective experience. Since Nagel repeatedly refers to neutral monism and panpsychism together, henceforth, when speaking of Nagel’s views, these terms will be treated as synonyms.

In the emergence-based alternative that Nagel also considers, mental properties only emerge at a certain level, so only some higher levels are both material and mental; lower

levels are just material. Nagel does not attribute a particular ontological cardinality to emergence (hence the entry of “unspecified” in Table 1), but it is commonly regarded by its advocates as a monistic view.

For Nagel, neutral monism and emergence each has both appeal and weaknesses. The weaknesses of monism are that it attributes to simple entities psychic properties that are unobservable, that it is hard to conceptualize whole-part relationships involving mental properties, and that it seems to require not only mind but also life to go all the way down, a violation of Ockham’s parsimony principle. The weakness of emergence is that the sudden appearance of mind without any precursors appears magical.

Nagel’s analysis includes a constitutive (synchronic) mode and a historical (diachronic) mode, and implicit in his analysis is the idea that both modes are required for a complete explanation of any phenomenon. It is useful to expand this explanatory dyad into a structure-function-history triad (Gerard 1958), where structure refers to the internal order of a system, function to its participation in external order, and history to qualitative change in both internal and external orders. In terms of this triad, constitutive analysis of mind focuses on both brain (structure) and behavior (function), while historical analysis takes up the origin of mind and its evolutionary development. The structure-function dyad can also be applied within the nervous system, where the function of a subsystem is its external relations with other subsystems, and its structure is its internal organization.

Nagel’s objections to materialist neo-Darwinism go further than denying its capacity to explain subjective experience. He doubts that evolutionary theory adequately explains the history – especially the origins – of life even in its simplest forms. While this objection to evolutionary theory is a supplementary argument that is not dependent on his primary thesis, Nagel’s doubts about the adequacy of MND to account for the origin of life reinforce his doubts about its adequacy to explain mind.

In the sections that follow, a critique of Nagel’s position is used to articulate an opposing emergence-based systems-theoretic view on the challenge of understanding subjective experience. In this opposing view, however, one aspect of Nagel’s position is retained in modified form.

The Informational is Not the Same as the Experiential

At the very outset, there is a problem with Nagel’s use of words like “mind” and “mental.” Most of the time he seems to use these words to refer to subjective experience, but sometimes they appear to refer merely to *informational* processes occurring in organisms, without regard to whether these processes are or are not experienced. The word “mental” can thus mean experiential or informational (or both), but these two meanings must be differentiated, since they are the very basis for the distinction between the “hard problem of consciousness” (Chalmers 1996), namely explaining the fact of subjective experience, and the not-so-hard problem, namely understanding biological information processing.

“Information” here refers to what can be analyzed with Shannon’s (1949) theory, which was originally called a theory of communication but which applies to information more generally. As defined by this theory, information is not material, despite its being

instantiated on matter-energy markers, because it is quantified in non-physical units (bits), i.e., units that do not involve grams, ergs, and so on. Technically, information is a reduction of uncertainty, which Shannon called “entropy,” borrowing the term from statistical mechanics. If p_i is the probability that a system is in state i , then the Shannon entropy (uncertainty) of the system is $H = -\sum p_i \log_2 p_i$. Entropy in statistical mechanics, given by $S = -k \sum p_i \ln p_i$, involves the same sum over probabilities, but this sum, scaled to natural logarithms, is multiplied by Boltzmann’s constant, k , which has physical units: ergs per degree Kelvin. Having such units is an indication that S is physical; H , lacking such units, is not. The second law of thermodynamics applies to thermodynamic and statistical mechanical entropy, but not to Shannon entropy: in isolated systems, S must increase over time or stay the same, but in general there is no such law for H .

Information applies not only to such contexts as communication, measurement, and computation, but even more deeply to the fundamental notion of constraint, which is central to the basic idea of relation. So information can be used to characterize and to measure the organization of any static structure or dynamic process.⁴ This is further explained in the next section, where its relevance to emergence is considered.

Nagel’s omission of a distinction between the informational and the experiential obscures the fact that informational processes (including those occurring in nervous systems) pose no fundamental challenge to scientific understanding. Functionalist theories in cognitive science, even those that attempt to explain consciousness (Dennett 1991), are primarily if not exclusively informational. The experiential aspect is left unexplained, so the hard problem of consciousness remains unsolved. While it is reasonable for Nagel to argue that psycho-physical reduction has failed so far for the hard problem of consciousness, it would be incorrect for him to assert that it has also failed to yield an understanding of cognitive information processing. While he doesn’t make this assertion explicitly, by not differentiating in his idea of the “mental” between the informational and the experiential, he implies that the former is as mysterious as the latter. It isn’t. Through advances in neurobiology, we have a growing understanding of biological information processing.

It is hard to see how panpsychism could accommodate the difference between the informational and the experiential and yet remain purely reductive. Nagel repeatedly (and synonymously) refers to the “protomental” or “protopsychic” but does not define these. If protomental means only the informational, to explain experience panpsychism would need a notion of emergence, which is what reductive panpsychism is supposed to obviate.⁵ If, however, the protomental includes the experiential, the relationship between the informational and the experiential needs to be explained. These alternatives might be called weak and strong panpsychism, respectively. The tension between them cannot be resolved by saying that the experiential is the inner aspect of phenomena whose outer aspect is informational, since in panpsychism outer aspects of phenomena are material, not informational. If panpsychism is expanded into an experiential-informational-material triad, this scheme would be vulnerable to the criticism that the material and the informational are plainly apparent in many phenomena, but the experiential is apparent only in complex biological organisms.

That *material* aspects of biological phenomena go all the way down is unproblematic for everyone (except idealists). That *informational* aspects of biological phenomena go all the way down is novel but plausible; in fact, one might argue that whatever scientific appeal panpsychism might have is its suggestion that something nonmaterial exists at all levels. But that *experiential* aspects of biological phenomena go all the way down⁶ is an extraordinary claim that lacks not only the extraordinary evidence obligatory for such claims, but any scientific evidence at all. One might thus abandon experiential panpsychism, but preserve a pan-informational scheme,⁷ where information goes all the way down but experience is a high-level emergent. This position is advocated below, but first Nagel's objections to emergence must be considered, and to take these up, it is necessary to have an overview of some of the conceptions of emergence that are available in the philosophical and scientific literatures.

Various Conceptions of Emergence

Emergence involves novelty having greater or lesser "irreducibility," and is thus often characterized as being strong or weak. O'Connor and Wong (2012), without using these particular labels, classify the British emergentists of the 1920s, such as Broad, Alexander, and Morgan, as holding the strong emergence position; Chalmers (2006) agrees. These emergentists differ somewhat in the strength of the type of emergence that they advocate. O'Connor and Wong (p. 14) write that

For Mill and Broad, emergence involves the appearance of primitive high-level causal interactions ... *additional* [italics added] to those of the more fundamental levels. Alexander, by contrast, is committed only to the appearance of novel qualities and associated high-level causal patterns that cannot be directly expressed in terms of the more fundamental entities and principles. But these patterns do not supplement, much less supersede, fundamental interactions.

Broad's view is thus stronger than Alexander's. Despite this difference, for Alexander as for Mill and Broad "emergent qualities were metaphysically primitive"⁸ (O'Connor & Wong 2012). Emergence becomes even stronger in Morgan (1923) when he writes of an "increasing richness ... in substance throughout the stages of evolutionary advance." If emergent entities are different in their very substance, this is strong emergence indeed.

Mayr (1982) defined emergence as "the appearance of new characteristics in wholes that cannot be deduced from the most complete knowledge of the parts, taken separately or in other partial combinations." In this, he echoed Broad (1925, p.61), who considers a whole having constituents A, B, and C in relation R, and defines emergence to mean that "the characteristic properties of the whole R(A,B,C) cannot, even in theory, be deduced from the most complete knowledge of the properties of A, B, and C in isolation or in other wholes which are not of the form R(A,B,C)." Wholes not of the form of R(A,B,C) include partial combinations of constituents, i.e., decompositions of R that organize one, two, or even all three pairs of constituents, but do not organize all constituents in a single relation. Write R as ABC, and its minimal decomposition, which includes all three pairs of constituents, as AB:AC:BC (the colon means "and"). Information-theoretically, ABC can embody a constraint that is absent in AB:AC:BC and in any simpler decomposition. Because of this additional constraint, R can have properties not present in or deducible

from any “partial combination” of constituents. This is a rigorous interpretation of “the whole is more than the sum of its parts,” an idea central to emergence; and this is what is meant above in the statement that “information” encompasses “organization.”⁹ The significance of holism was recognized by the earliest writers on emergence: properties of wholes that go beyond those of their parts are what Mill (1843) meant by “heteropathic” as opposed to “homopathic” effects, and what Lewes (1875) meant by “emergent” as opposed to “resultant” effects.¹⁰

Such holism can give meaning to Broad’s view of high-level interactions being additional to low-level interactions. It is also a mereological interpretation of “downward causation,” a feature often used to characterize strong emergence, as the effect of a whole on its parts that goes beyond the interactions in all subsets of parts. Downward causation, however, in the philosophy of mind is often interpreted as the influence of mind on matter. This interpretation is not mereological, since it is not clear how the mind can be viewed as a whole that has the body as its parts. Mereological downward causation can operate wholly in the domain of matter, where the question of how such causation can occur is easy to state. By contrast, a downward causation of mind on matter is a different idea, which requires considerable explanation. The two should not be conflated.

Strong emergence is sometimes called ontological emergence, as opposed to weak emergence, which is called epistemological, i.e., observer-dependent, emergence (Chalmers 2006).¹¹ O’Connor and Wong define the latter as characterizing phenomena that cannot be predicted (Rosen 1985) or cannot even be represented in lower-level terms, but to which no special ontological status is accorded. Downward causation, dismissed by Bedau (1997) as “magical,” is excluded, so weak emergence, defined by Bedau as emergence demonstrable only by simulation,¹² is “metaphysically innocent.” Chalmers has a yet more liberal idea of weak emergence, which he defines as involving phenomena that are “*unexpected* [italics in the original] given the principles governing the low-level domain.” Being unexpected is weaker than being unpredictable which is weaker than being incapable of representation, so epistemological emergence, like ontological emergence, has stronger and weaker versions. What cannot be represented qualifies unequivocally as both irreducible and novel. What cannot be predicted but is derivable by simulation is irreducible in a technical sense (the absence of a closed form solution) but is not completely devoid of explanatory value; and it is less novel than what cannot be represented. What is merely unexpected is not irreducible at all. The requirement of irreducibility can in fact be completely jettisoned, and phenomena can be considered emergent solely by virtue of their novelty – novelty being the sine qua non of emergence – even if they have been given – after first being observed – a complete explanation, and even if this explanation is theory-based and not simulation-based.¹³ For example, macroscopic properties of a gas are emergent from the micro-dynamics of the molecules that make up the gas, but the thermodynamics of gases can be reduced to statistical mechanics. As Bunge argues (2010, p.77), “Explained emergence is still emergence. In other words, epistemology cannot erase ontology because the universe is objectively stratified.” Novelty does not require irreducibility.

This last conception of emergence makes it ubiquitous and perhaps prosaic. One might prefer, as Chalmers (2006) does, to reserve use of the word for rare phenomena, but many thinkers adopt the broader notion. Morgan (1923, p.1) writes that

Under what I here call emergent evolution stress is laid on this incoming of the new. Salient examples are afforded in the advent of life, in the advent of mind, and in the advent of reflective thought. But in the physical world emergence is no less exemplified in the advent of each new kind of atom, and of each new kind of molecule. It is beyond the wit of man to number the instances of emergence.

Morowitz (2002) offers 28 examples of emergence, and Bunge (2010) regards emergence as central to scientific metaphysics. Pines (2014, p.22) suggests that emergence is a “unifying theme for 21st century science,” and writes, “Indeed, we live in an emergent universe in which it is difficult, if not impossible, to identify any existing interesting scientific problem ... that is not emergent.”

The systems-theoretic orientation adopts this broad notion of emergence: a system is a set of elements and a set of relations that link elements via their attributes, where the system is a higher-level element, and elements are lower-level systems. (Broad’s R, discussed above, is a simplified notion of system, in which relations organize elements directly rather than through attributes, and multiple relations between elements are subsumed in a single net relation.) This recursive definition of system, expanded to speak also of its environment, reflects a view of the world as inherently stratified, ontologically and epistemologically. The relations that organize a system confer upon it certain attributes. These attributes can be novel, i.e., they can differ from the attributes of the elements and from the attributes of subsystems less than the whole. If so, they are said to emerge from the organizing relations of the system. Emergence, in this sense, is a very general aspect of the world. Given this definition, emergence is not the opposite of reduction. They are two sides of the same coin: emergence is what we see looking up from the lower level; reduction is what we see looking down from the higher level (Fromm 2004).

Nagel’s Objections to Emergence

Nagel thinks that a neutral monism-based theory of experience is more likely than an emergence-based theory. His reasons for rejecting emergence are a little obscure, not surprisingly since we don’t actually have either an emergence-based theory or a neutral monism-based theory of consciousness to assess, but his position seems to be that subjective experience is extremely novel, and that an emergence explanation of such great novelty would necessarily be irreducible, so it would have no explanatory value. Nagel writes (pp. 55-56),

If emergence is the whole truth, it implies that mental states are present in the organism as a whole, or in its central nervous system, without any grounding in the elements that constitute the organism, except for the physical character of those elements that permits them to be arranged in the complex form...that connects the physical with the mental...It is hard to give up the assumption that whatever is true of the complex must be explained by what is true of the element.

And,

That such purely physical elements, when combined in a certain way, should necessarily produce a state of the whole that is not constituted out of the properties and relations of the parts still seems like magic... such emergence, even if systematic, remains fundamentally inexplicable.

Understanding Nagel's viewpoint poses some difficulty, since sometimes it seems that he thinks that wholes can only acquire properties by inheriting them from their parts, which would be a fallacy of decomposition,¹⁴ but sometimes he seems to give proper weight to the relations of the parts. The first inclination is illustrated in the above sentence about what it is hard to give up. It is also implicit in Nagel's embrace of panpsychism, which sees the key to the mentality of an organismic whole in the hypothesized mentality of its parts; this is why he repeatedly acknowledges concern about the problem of conceiving mental whole-part relations. Overall, the book conveys the impression that Nagel believes that mental wholes require mental parts.

Yet, as the above quote shows, Nagel sometimes invokes relations explicitly, and he also seems to accept emergence in some contexts. He writes (p. 56),

That does not mean that new phenomena cannot emerge at higher levels, but the hope is that they can be analyzed through the character and interactions of their more elementary components. Such *harmless emergence* [italics added] is standardly illustrated by the example of liquidity, which depends on the interactions of the molecules that compose the liquid. But the emergence of the mental at certain levels of biological complexity is not like this. According to the emergent position now being considered, consciousness is something completely new.

Presumably, this "harmless emergence" is either weak emergence or the broad notion of emergence discussed above. Nagel doesn't believe that a theory of subjective experience based on this type of emergence could account for the great novelty of consciousness, yet he gives no argument to defend this belief. He is not impressed by the novelty of liquidity that is explained by "harmless emergence." More seriously, he does not realize – or does not accept the fact – that the novelty of *life* itself is explained by this very same "harmless emergence." Nagel never entertains the possibility that interactions of parts might engender mental properties, especially when grounded not only in material processes but in informational processes as well.

So one response to Nagel's rejection of emergence is simply to deny his premise that emergence in its broadest sense – an emergence fully compatible with reduction – is inherently incapable of explaining subjective experience. One might, however, ask how such an emergent/reductive explanation of mind would differ from the reductiveness of the panpsychism that Nagel favors. The answer is that the reducibility that can be joined consistently with emergence would not attribute experience to the lowest levels of mind or life. The realm of experience would be local. It would not cast new light on the basic properties of life or matter. By contrast, a neutral monism explanation of experience

would attribute experience to lower levels, and would impact our understanding of life and matter.

Or, one might accept Nagel's claim that explaining consciousness requires a form of emergence that is deeply irreducible. Nagel's objection to such strong emergence is not unusual; Bedau and many others concur that strong emergence seems magical, especially if it entails downward causation. If all that can be said about strong emergence is that it manifests supervenience, that macro-properties presume changes in but are still not reducible to micro-properties, then such emergence is indeed magical and the associated notion of supervenience is scientifically empty. An ontology that posits an emergence of mind from matter about which nothing more can be said is dualistic, regardless of any claims it makes to the contrary, and a dualistic ontology cannot offer a satisfactory scientific explanation of subjective experience.

Nagel's objection to strong emergence might, however, be countered if one could explain *why* reduction is impossible, and if such an explanation can be given, it would provide content to the notion of supervenience. For example, reduction might be impossible in principle because of mathematical undecidability (Baas & Emmeche 1997) or because of spontaneous symmetry breaking (Anderson 1972). Or, irreducibility could be seen as not mysterious if it is defined as non-decomposability, in the sense discussed above. Such irreducibility might also be considered a possible interpretation of mereological downward causation.

Consciousness, Not Life, is Yet Unexplained

Nagel not only conflates the informational and experiential in his notion of mind, but he also conflates mind and life in his suggestion that the mental quality that goes all the way down and that explains consciousness in complex organisms might also (at least partially) explain life in simple organisms. He offers no hint, however, about what mental quality could possibly accomplish both purposes. His panpsychism, the view that everything has mind, here shades into pan-vitalism (hylozoism), the view that everything is alive. This isn't surprising: mind is commonly assumed to supervene on life, so if mind goes all the way down, life might well also.

It should be noted that the assumption that mind depends on life is not universally held. The field of artificial intelligence (AI), for example, is implicitly based in the belief that mind does *not* supervene on life, and can thus be instantiated in artificial entities that are not alive by most definitions of life: these entities are allopoietic, not autopoietic (they are constructed by external processes rather than constructing themselves via matter-energy resources gained by metabolism); their structure and behavior are not coded for or regulated by genomes transmitted by heredity; populations of these entities do not evolve by natural selection; and so on. So to the degree that one thinks that AI implements mind, one can reasonably argue that mind does not depend on life.

But here the distinction between the informational and the experiential is critical. AI implements only the informational. Few scientists would accept the proposition that current AI instantiations are accompanied by experience. Nonetheless, the assumption of machine sentience is ubiquitous, especially in the cinema, reflecting the advances that

have been made in Turing-test mimicry of human intelligence. Even consciousness researchers do not always distinguish between the experiential and the informational. Koch and Tonioni (2011) offer a test for consciousness that is a variation on Turing tests for intelligence; since all such tests are functional, i.e., external, it is hard to see how they could test for the presence or absence of internal experience. Admittedly, a test for experience *must* be functional, since experience is private by definition. But the functions that are tested must be extensive and diverse. The Koch-Tonioni test (which asks questions of the sort, “What is wrong with this picture?”) is far too simple.

Nagel links the hard problem of consciousness to what he sees as the inadequacy of the neo-Darwinian account of life. He writes (p. 8), “Mind, as a development of life ... casts its shadow back over the entire process.” He goes further, saying that our understanding of life is suspect even on its own terms, because our *historical* understanding of life is inadequate (p. 5): “The more details we learn about the chemical basis of life and the intricacy of the genetic code, the more unbelievable the standard historical account becomes.” Nagel doubts the adequacy of MND to explain the diversity of life forms that exist¹⁵ and (especially) also the origin of life.¹⁶ This alleged inadequacy of our historical account of life somehow undermines also our constitutive account of life. But Nagel does not say what exactly in our constitutive understanding is undermined.

Much can be said in response to Nagel’s views about life. First and foremost, the main processes of life, viewed constitutively as opposed to historically, are *not* a mystery. They are well understood. Pan-vitalism gains no support from any scientific fact. There is no scientific basis for a reductive theory that attributes protolife to constituents of living systems, to constituents of these constituents, and so on. Even if one allows, for the sake of argument, that some constituents or their precursors, e.g., ribozymes, exhibit protolife, one would not make this attribution for constituents even one level down, since nucleotides or amino acids are not living in any sense of the word. Life plainly does *not* go all the way down. Those who think otherwise typically confuse energy or dynamics with life. If life does not go all the way down, it is likely that mind – specifically, experience – also does not go all the way down.

The constitutive explanation of life that is accepted today is emergence based, and this emergence is not irreducible. Quite the opposite: we explain how biological phenomena result from components of living systems, their relations, and the interactions of systems with their environments. Nagel says (p. 87) that his imagined neutral monism explanation of experience “is modeled on the physical reductionism encouraged by molecular biology, but with an expanded metaphysical basis.” This is ironic, since the reductiveness of molecular biology supports and complements emergence, understood broadly in the systems-theoretic way discussed above.

Nagel’s view that a historical explanation of life requires a convincing account of origins, and that a constitutive explanation of life requires an adequate historical account, are unreasonable. Origins are *always* difficult to explain. The origin of life on earth is one of many mysteries. The origins of the universe and its galaxies and the solar system and the human species and human language are also obscure. Nagel is correct in insisting that a full understanding of any phenomenon requires constitutive *and* historical explanations,

but he errs in thinking that these explanations are totally interdependent, and that, within historical explanations, accounts of origin and development are similarly inseparable. Absence of an account of origin does *not* undermine explanation of what happened after origin. Incompleteness of historical explanation does *not* invalidate constitutive explanation. Historical explanations are always incomplete; it is remarkable that they can be offered at all, since phenomena that occurred in the past not only cannot be experimented upon; they cannot even be observed directly or fully.

One is certainly warranted in calling neo-Darwinism incomplete, but incompleteness is the rule and not the exception in scientific theories, so if being incomplete means being false, much of science would have to be considered false. This may be acceptable – science takes pride in fallibility – but it is preferable to reserve the word “false” for assertions that are erroneous for reasons other than incompleteness. For example, the theory that held that proteins, not nucleic acids, are the carriers of genetic information, and the theory that declared that stretches of DNA that don’t code for proteins are junk, were both proven false, and not because of incompleteness. Admittedly, these are easy cases, and the relationship between incompleteness and falsity is more complicated. One might, for example, ask: at high velocities, has Newtonian mechanics been falsified by special relativity, or does the fact that the two theories are identical at low velocities render Newtonian mechanics merely incomplete, or, to put it positively, a good approximation to truth in a well-defined, albeit limited, domain?

Resolving such issues in general is clearly outside the scope of this paper, but the dramatic title of Nagel’s book does call for explicit comment. A great deal in biology has become comprehensible via neo-Darwinism. Dobzhansky (1973) puts this even more strongly: “Nothing in biology makes sense except in the light of evolution.” It seems likely that what Nagel has in mind in imagining MND proven “false” is something like the status of Newtonian mechanics after the revolutions of quantum theory and relativity. In the domains of its past successes, Newtonian mechanics still stands, but in other domains quantum and relativity theories reveal it to be incorrect in important respects; in these respects it is appropriate to call it “false.” A neo-Darwinism that was overturned in the future would probably still be regarded as true, or at least adequate, in the domains of its successes. But if some future neutral monist theory of mind explained subjective experience, *and* altered our constitutive understanding of life, *and* introduced subjective experience as an important factor in evolution even for simple organisms, *and* provided new understanding about the origins of life, this indeed would be a major revolution in science; in such a hypothetical future, it would be accurate to say that MND was proven “false.”

On the other hand, if subjective experience were explained as an instance of emergence, the effects of this accomplishment on science would most likely be easily assimilated. Its effects would only be local, and not go all the way down. They would not impact our understanding either of life or of matter. Such a scenario would confirm Simon’s (1981) view that reality is vertically “partially decomposable,” meaning that adjacent levels of reality are usually (though not always) relatively insulated from one another. Of course, it is possible that subjective experience will never be explained by any theory. But neuroscience is young, and science has only just begun to take on the challenge of

understanding consciousness. Despair over our current lack of understanding is premature, and arguments from impatience are unwarranted.

Life is Not Merely Material but Also Informational

The solution we have to the constitutive problem of life suggests that the constitutive problem of mind might also be solved by reducible emergence. It also suggests that a successful solution will not be purely materialist, since our constitutive view of life is in fact not purely materialist. Nagel's characterization of neo-Darwinism as materialist is incorrect, however conventional it is. Life is based not only on materiality, but also on informational processes. That autopoiesis, replication, adaptation, and evolution can be conceptualized without reference to a *specific* materiality is the core insight of the field of artificial life (Langton 1997): properties of life could in principle be realized via different forms of matter. But even if life necessarily requires the materiality it has on Earth, it would still be an error to regard life as purely material, since its signature properties are formal. Life requires autopoiesis (Maturana & Varela 1980), and such self-organization is a formal property, although in concrete systems it manifests in material structures; the purely formal aspects of self-organization are explored, for example, in artificial life studies (e.g., Bagley 1992, Fontana 1992). Development, viewed as informationally (genomically) programmed or as self-organization (Laland et al 2014), has a formal aspect. The principles of evolution are also formal, as shown by the fact that these principles are applied to nonbiological, even nonphysical, contexts. For example, the genetic algorithm utilizes the evolutionary principles of replication, variation, and selection to perform multi-solution global optimization; this algorithm is mathematical (computational), not physical. Understanding life requires ideas of information to explain genetic determination, positive and negative feedback to explain biological regulation, and utility and fitness to explain agency and evolution. Unlike matter and energy, information and utility are nonphysical. The nonphysicality of Shannon information is noted above. Utility, in von Neumann and Morgenstern's (1944) game/decision theory, is also expressed in nonphysical units: utils, just like bits, do not have dimensions of mass or energy. Fitness, a concept that is absolutely necessary for Darwinism, is merely one type of utility; this is clear in evolutionary game theory (Maynard-Smith 1982). Cybernetic control mechanisms in molecular biology, elegantly first popularized by Monod (1971), are also inherently formal arrangements. The formal includes more than the informational, but the informational is a central part of it.

Neo-Darwinism thus relies on "stuff-free" theories (Bunge 1973) that are not purely materialist, so the M in MND is inappropriate. This is not merely an epistemological observation; it is also an ontological one: genetic information, feedback control, evolutionary fitness, and so on, are formal, not material, properties. Both epistemologically and ontologically, the material and the immaterial are equals. Referring to the latter as "form," matter and form are joined together, as Aristotle held: just as form is always materially instantiated, so matter always manifests with form. There is no justification for privileging matter over form by calling their fusion "materialist," taking form for granted. One could as well call our ontology "formist," thus taking matter for granted. Why matter continues to be privileged is best left to historians and sociologists of science to explain. They would probably understand this as

a reaction against idealism and dualism and also as the result of an older victory of atomists over Pythagoreans. The emergence after World War II of systems theory and cybernetics and the renewal of this research program in the sciences of complexity partially corrected this imbalance. (Before this development, the fundamentals of natural science were matter and energy, but afterwards information and utility were added as new – and nonphysical – fundamental categories.) However, in mainstream science, materialism still prevails as the default label for scientific ontology. Even ontologies that posit emergence often come with a materialist label. For example, Broad (1925) wrote of “emergent materialism,” and today Bunge (2010) calls his ontology materialist¹⁷ despite the fact that he endorses emergence and the systems perspective. Cybernetics was seen as a neo-mechanism, mechanism often being a surrogate term for materialism. But the transdisciplinary scope of the systems and complexity sciences derives from their being stuff-free, i.e., formal, Pythagorean.

Mind is One of Several Informational Emergents

Life is now understood as an instance of emergence that encompasses both the material and the informational. Mind is also both material and informational, but unless life is conflated with mind, only mind has, in sufficiently complex organisms, an additional aspect that is experiential. One can regard experience as an emergent novelty, which is, however, continuous with other informational aspects of living and nonliving matter. This continuity could be described as a hierarchy of levels of information, as shown in Table 2. This hierarchy is introduced to convey three ideas: (1) that types of information vary over the domains of matter, life, and mind¹⁸; (2) that experience occurs only at the highest levels of the hierarchy; (3) that unconscious cognitive processes, which occur just below these levels, bridge the informational and the experiential.

Table 2 Levels of information

The list of system types is borrowed with modification from Boulding (1956). The realm of experience is restricted to only the top two shaded levels.

SYSTEM TYPE	INFORMATION	PHENOMENON	DOMAIN
(8) humans	neurological	self-awareness	↑↑↑ mind
(7) complex animals	neurological	experience	
(6) simple animals	neurological	unconscious cognition	
(5) genetic-societal (e.g., plants)	meta-genomic	differentiation	 life
(4) genetic (cells)	genomic	reproduction	
(3) autopoietic systems	network	self-organization	 matter
(2) control mechanisms	control	regulation	
(1b) clockworks (dynamic systems)	algorithmic	organization	
(1a) frameworks (static systems)	form	organization	

In this hierarchy, emergence is exemplified in differences between levels, but the table might be viewed as compatible with a kind of neutral monism in the similarities between the levels. Something *does* go all the way down but it is information, not experience. Psyche, understood as subjective experience, does not. Psyche emerges.

In the scheme of this table, subjective experience emerges at level (7); its fuller development in a consciousness that encompasses a sense of self, not identical with the

contents of experience, emerges at (8). Level (6) is transitional, being the first level where information is neurological. Information processing here is not experiential, as this level refers to the cognition of very simple animals. Information processing at this level is also a component of cognition in complex animals having subjective experience; that is, levels (7) and (8) include but go beyond level (6). From the perspective of this hierarchy, the human mind encompasses all three top levels: some mental processes are unconscious (6), some are experienced but not associated with a sense of self (7), and some include self-awareness (8). This triadic scheme is inspired by Bennett (1961), who called these levels automatic (6), sensitive (7), and conscious (8), respectively. In Kahneman's (2011) distinction between types of thinking, fast thinking, which is associative, exemplifies the automatic; slow thinking, capable of being logically ordered, occurs at the sensitive and/or conscious levels. Ordinary perception might be considered to occur at level (7), subliminal perception below it at level (6), and apperception (perception associated with a sense of self) above it at level (8).

The difference between what is unconscious and what is conscious is partially contingent or fluid: some mental processes seem to be permanently and fully unconscious, but others can rise to the level of experience. If neural correlates of consciousness could be reliably established, one would then try to understand these correlations in the hope of explaining how the experiential (7) emerges from the non-experiential (6). We cannot imagine now what an explanation of experience might look like, but we do have an idea of what empirical knowledge would be valuable to have, upon which a successful theory might conceivably be built.¹⁹ Correlations themselves would not solve the hard problem of consciousness, but they would focus attention on precisely what needs to be explained.

The lowest level of the table shows information in its most general manifestation (1a), namely "form" (spatial information) measured in information theory by Shannon entropy, and the process version of information (1b) developed in the algorithmic information theory of Kolmogorov (1965) and Chaitin (1975). Level (2) information occurs in control systems, where it is concentrated and explicit, e.g., the set point of a thermostat. At level (3), self-constructing but nonliving systems (such as flames, eddies, or tornadoes, autopoietic in a narrow sense) depend on a network of interactions that construct and maintain the organization of the system; information here is dispersed throughout the network. In cells, concentrated genomic information at level (4) governs metabolic autopoiesis, reproduction, and adaptation; and reproduction generates populations that have history, i.e., undergo evolution.²⁰ In cellular metabolism, enzymes are catalysts at level (1) specified by genomic information, catalytic information being the spatial organization of an active site, which has static and dynamic aspects. Cells also include feedback and feedforward control mechanisms, assigned here to level (2). In multicellular plants and simple animals, genomic information is subject to meta-level control (5) by chemical messengers, e.g., hormones; in differentiation, there is informationally guided self-organization. In the animal kingdom, organisms implement informational regulation neurologically (6), and at some point, complex neural systems allow the emergence of subjective experience (7) and self-awareness (8).

Table 2 hypothesizes that neurological information processing occurs at levels (6), (7), and (8), but that subjective experience emerges only at (7) and (8). It is thus of interest to

comment here on Tonioni's (2008) proposal, which resembles a similar notion by Chalmers (1996), on the informational correlates of consciousness. Tonioni holds that consciousness is "integrative" information processing; more specifically, he proposes that if (a) information processes are holistic (in the sense of holism discussed above as it relates to emergence) and if (b) their associated dynamic states have high information, then consciousness will be present. Tonioni's theory is usually aligned with panpsychism, but if these two conditions are only satisfied in complex neurological systems, this theory also posits emergence. Three reservations about this theory: First, testing for the satisfaction of the two conditions is difficult in practice. Second, the theory doesn't explain why subjective experience should be the consequence of these conditions being satisfied, so it doesn't propose a solution of the hard problem of consciousness. Third, the theory implies that mind can be present even if life is absent, perhaps a difficult proposition to swallow. On the positive side, however, the theory is testable in principle, if high information but holistic computation can be implemented, and if convincing tests for consciousness can be devised. Tonioni's conditions might also be viewed as only necessary but not sufficient for consciousness to occur.

Table 2 is only a working sketch, and the reader is invited to improve upon it. The levels are roughly ordered, going upwards, from ubiquitous to less and less prevalent, from the general to the more and more particular.²¹ Each emergence adds new informational phenomena to those already existing, so a system that exhibits informational phenomena at one level typically also exhibits informational phenomena at lower levels. Levels are grouped into domains, and the domains of matter, life, and mind are similarly nested: life emerges from matter, and mind emerges from life. Parts of the table reflect a historical (evolutionary) sequence.²² Still, since there is no universal character to all instances of emergence (Bunge 2010), there is no simple organizing principle that governs this entire set of levels. Some inter-level relations are mereological (compositional), linking parts and wholes; for example, humans and other animals as well as plants are composed of cells. Other inter-level relations are regulatory; for example, neurological processes control meta-genomic processes, which control genomic processes. Aside from the fact that the multiple instances of emergence in the table are all not of the same type, the table is also incomplete: it needs at least a level (9) that covers sociocultural, e.g., language-based, phenomena. The table exemplifies the "crude look at a whole" that Gell-Mann (1994) regards as a valuable contribution that can be made by the complexity sciences.

This view of mind, which combines emergence with a modified neutral monism, might be called "continuous emergence," where "continuous" has two meanings. It refers to emergent phenomena that are continuous with (that resemble) lower-level phenomena in being informational, and it asserts continuous (repeated) occurrences of emergence resulting in a hierarchy of levels, each similar to and different from the level below it. What goes all the way down is information and not psyche, and higher-level informational phenomena do not actually reduce to lower-level informational phenomena, since these emergents are not all of the same type.

Summary

To summarize: this paper agrees with Nagel in rejecting materialist reductionism as an explanation of mind. In arguing against reductive neutral monism and for emergence as the best approach to subjective experience, it disagrees with Nagel. In offering a pan-information but not panpsychism view in which information but not experience goes all the way down, it modifies Nagel's position.

Taking up the question posed by Nagel's book title of whether MND is false, the position of this paper is (a) that neo-Darwinism is in fact not purely materialist, and (b) that it is not likely to be proven false in the future, because an emergence-based explanation of mind is more likely than a neutral monist explanation of mind. The precedent of our emergence-based explanation of life gives us reason to expect an emergence-based explanation of mind. If life emerged from matter and mind emerged from life, and if we focus only on the constitutive question and not also the historical one, we note that both life and mind were once mysterious, but now only the mystery of mind remains. The solution of the first mystery suggests that solution of the second is conceivable. We cannot imagine the contents of such a solution, but we can imagine the relationship such a solution might have with the rest of scientific knowledge. The kind of solution that Nagel predicts would revolutionize nearly all of science. A different possibility, one which is argued here as more likely, is a scientific revolution that is local, and that doesn't alter our constitutive explanation of life or our understanding of matter. Of course, there are those who believe that explanation of subjective experience will forever elude us, but it is hard to see how one can defend a prophecy about what science will *never* be able to do. Some people who hold this belief even claim to demonstrate its truth by philosophical argument. From the perspective of a working scientist, such alleged demonstrations are not convincing.

Continuous emergence takes subjective experience to be absent in simple systems. The emergence of experience, however, is one of multiple emergent informational phenomena that range over the domains of matter, life, and mind. This justifies a pan-information ontology, but not the specific panpsychism ontology that Nagel favors but does not explicate. But a pan-information ontology does not solve the mystery of subjective experience. To solve this mystery, new scientific discoveries will be required.

Acknowledgements

The author thanks Hugo du Coudray, Kirke Wolfe, and John Hammond for helpful discussions of an earlier draft of this paper and of Nagel's book. He is also greatly indebted to the editor and an associate editor of this journal and the two reviewers for their exceptionally thorough and invaluable critique of that draft. The deficiencies that remain are of course only the responsibility of the author.

References

- Alexander S (1920) *Space, time, and deity*. Macmillan, London
- Anderson PW (1972) More is different. *Science* 177:393-396
- Baas NA, Emmeche C (1997) On emergence and explanation. *Intellectica* 1997/2(25):67-83
- Bagley R, Farmer JD (1992) Spontaneous emergence of a metabolism. In: Langton CG, Taylor C, Farmer JD, Rasmussen S (eds) *Artificial life II*. Addison-Wesley, Redwood City, pp 93-140
- Bedau M (1997) Weak emergence. In: Tomberlin J (ed) *Philosophical perspectives: mind, causation, and world*, vol 11. Blackwell, Malden, pp 375-399
- Bennett, JG (1961) *The dramatic universe, vol 2: the foundations of moral philosophy*. Hodder & Stoughton, London
- Boulding, K (1956) General systems theory - the skeleton of science. *Management Science* 2:197-208. In: Buckley, W (ed) (1968) *Modern systems research for the behavioral scientist*. Aldine, Chicago, pp 3-10
- Broad, CD (1925) *The mind and Its place in nature*. Kegan Paul, Trench, Trubner & Co, London
- Bunge, M (1973) *Method, model, and matter*. D Reidel, Boston
- Bunge, M (2010) *Matter and mind: a philosophical inquiry*, vol 287 *Boston studies in the philosophy of science*. Springer, New York
- Chaitin, G (1975) Randomness and mathematical proof. *Scientific American* 232(5):47-52
- Chalmers, DJ (1996) *The conscious mind: in search of a fundamental theory*. Oxford University Press, Oxford
- Chalmers, DJ (2006) Strong and weak emergence. In Clayton P, Davies P (eds) *The re-emergence of emergence*. Oxford University Press, New York
- Clayton P (2004) *Mind and emergence*. Oxford University Press, Oxford
- Dennett D (1991) *Consciousness explained*. Lane A (ed) Penguin Press, London
- Dobzhansky, T (1973) Nothing in biology makes sense except in the light of evolution." *The American Biology Teacher* 35:125-129
- Fontana, W (1992) Algorithmic chemistry. In: Langton et al (eds) *Artificial Life II*, pp. 159-210.
- Fromm, J (2004) *The emergence of complexity*. Kassel University Press, Kassel
- Gell-Mann, M (1994) *The quark and the jaguar: adventures in the simple and the complex*. W. H. Freeman & Co, New York
- Gerard, R (1958) Concepts and principles of biology. *Behavioral Science* 3:95-102

- Haberman, J (2014) Samuel Alexander's theism and its relation to Jewish thought. *Journal of Jewish Studies* LXV:368-389
- Hartshorne, C (1972) *Whitehead's philosophy: selected essays 1935-1970*. University of Nebraska Press, Lincoln
- Kahneman, D (2011) *Thinking, fast and slow*. Farrar, Straus & Giroux, New York
- Kauffman, S (1993) *The origins of order: self-organization and selection in evolution*. Oxford University Press, New York
- Koch C, T G (2011) A test for consciousness. *Scientific American* 304(6):44-47
- Kolmogorov, A (1965) Three approaches to the quantitative definition of information. *Problems Inform. Transmission* 1:1-7
- Krippendorff, K (1986) *Information theory: structural models for qualitative data*. Sage, Beverly Hills
- Laland, K et al (2014) Does evolutionary theory need a rethink? *Nature* 514:161-164
- Langton, C, ed. (1997) *Artificial life: an overview*. MIT Press, Cambridge
- Lewes, GH (1875). *Problems of life and mind*, vol 2. Kegan Paul, Trench, Turbner, & Co, London
- Livingston, C (1993) *Knot theory*. Math. Assoc. Amer., Washington
- Marvin, W (1912) *A first book in metaphysics*. Macmillan, New York
- Maturana, HR & Varela, FJ (1980) *Autopoiesis and cognition*, Boston studies in the philosophy of science vol 42. D Reidel, Dordrecht
- Maynard-Smith, J (1982) *Evolution and the theory of games*. Cambridge University Press, Cambridge
- Mayr, E (1982) *The growth of biological thought*. Harvard University Press, Cambridge
- Mill, JS (1843) *System of logic*. Longmans, Green, Reader, & Dyer, London
- Monod, J (1971) *Chance and necessity: an essay on the natural philosophy of modern biology*. Alfred A. Knopf, New York
- Morgan, CL (1923) *Emergent evolution*. Henry Holt & Co, New York
- Morowitz, HJ (2002) *The emergence of everything: how the world became complex*. Oxford University Press, New York
- O'Connor, T & Wong, HY (2012) Emergent properties. *Stanford Encyclopedia of Philosophy*, Stanford.
<http://plato.stanford.edu/archives/spr2012/entries/properties-emergent/>
- Pines, D (2014) Emergence, a unifying theme for 21st century science. *Santa Fe Institute Bulletin* 28:22-29
- Rosen, R (1985) *Anticipatory systems*. Pergamon Press, Oxford

- Sayre, KM (1976) *Cybernetics and the philosophy of mind*. Humanities Press, Atlantic Highlands
- Seager, W & Allen-Hermanson, S (2013). Panpsychism. *Stanford Encyclopedia of Philosophy*, Stanford.
<http://plato.stanford.edu/archives/fall2013/entries/panpsychism/>
- Shannon, CE & Weaver, W (1949) *The mathematical theory of communication*. University of Illinois Press, Urbana
- Simon, HA (1981) *The sciences of the artificial*. MIT Press, Cambridge
- Stubenberg, L (2014) Neutral monism. *Stanford Encyclopedia of Philosophy*, Stanford
<http://plato.stanford.edu/archives/fall2014/entries/neutral-monism/>
- Szabo, C & Teo, YM (2013) Post-mortem analysis of emergent behavior in complex simulation models. SIGSIM-PADS 2013, Montreal
- Teilhard de Chardin, P (1959) *The phenomenon of man*. Harper & Row, New York
- Tononi, G (2008) Consciousness and integrated information: a provisional manifesto. *Biol. Bull.* 215:216-242.
- von Neumann, J & Morgenstern, O (1944) *Theory of games and economic behavior*. Wiley, New York
- Whitehead, AN (1929) *Process and reality: an essay in cosmology*. Macmillan, New York
- Zwick M (2001) Wholes and parts in general systems methodology. In: Wagner G (ed) *The character concept in evolutionary biology*. Academic Press, San Diego, pp 237-256

¹ The list of idealist philosophies is provided by Nagel. About the inclusion of logical positivism, Nagel says (p. 37) that the logical positivists “analyzed the physical world as a construction out of sense data.”

² Materialism per se is not necessarily reductionist, and may be joined to an antireductionist or systems view that recognizes ontological stratification and emergence. This, for example, is the philosophical position of Mario Bunge (2010). One might say that materialism + reductionism = physicalism, the position that everything reduces to the realities described by physics. What Nagel means by materialism is really physicalism, even though he accords some epistemological autonomy to biology when he speaks of materialist neo-Darwinism. Nagel also uses “naturalism” or “scientific naturalism” as synonymous with materialist reductionism. This use of “naturalism” is not conventional, since naturalism is not normally regarded as the same as materialism. Naturalism allows certain nonmaterialist positions: it excludes supernaturalism, of course, but it does not exclude nontheistic idealism or neutral monism. Also, it does not exclude an emergence-based view that gives equal status to matter and information.

³ Saying that everything is *both* matter and mind differs from saying that underlying everything is a substrate that is *neither* matter nor mind, although Spinoza says both: underlying everything is substance, which is neutral but has attributes of both extension (matter) and thought (mind). Spinoza’s view can be considered a neutral monism and a panpsychism (Seager and Allen-Hermanson 2013). In this view, matter and mind do not interact, but are parallel alternative perspectives; the duality of attributes – there are more but we know only these two – is simply asserted. Dual attribute approaches deny that there *is* a hard problem of consciousness; like ordinary Cartesian dualism, they aren’t scientifically generative. Versions of neutral monism that identify experience as the underlying neutral substrate (see Note 6) also avoid the problem of explaining experience, since they take it as a primitive. For a survey of versions of neutral monism, and criticisms of them, see Stubenberg (2014).

⁴ See Note 7, which discusses Sayre’s ontology based in information theory, and the discussion below of Tonioni’s ideas about consciousness, which are also based in Shannon’s theory.

⁵ One can, however, imagine joining panpsychism with emergence, and the philosophy of Whitehead (1929) might be regarded as such a synthesis. Seager and Allen-Hermanson (2013) write, “[Whitehead’s] panpsychism arises from the idea that the elementary events that make up the world (which he called *occasions*) partake of mentality in some – often extremely attenuated – sense, metaphorically expressed in terms of the mentalistic notions of creativity, spontaneity and perception.” But Hartshorne (1972), writing about Whitehead’s philosophy, notes its “admission not merely of emergence, but of emergent or creative synthesis as the very principle of process and reality.” There is uncertainty about exactly what Whitehead means by the mentalist properties that he attributes to elementary events. If these are just informational, his account – written 20 years before Shannon’s theory – might be compatible with the position of this paper. Whitehead was influenced by the work of Alexander (Haberman 2014).

⁶ In many versions of neutral monism, the neutral substrate is identified as “experience” per se, regarded as not being tied either to subject (ego) or to object (matter), or the substrate is identified as sense data which is closely related to experience but detached from an experiencer. For a review of various thinkers in this tradition, including Mach, James, and Russell, see Stubenberg (2014). Understandably, this view is often accused of being a disguised idealism, but, from the perspective of this article, a more pertinent objection is that it is precisely the fact of experience – the hard problem of consciousness – that lacks scientific explanation. Stubenberg recalls the wry advice of William James: “Whatever you are totally ignorant of assert to be the explanation of everything else.” James was referring to “soul,” but scientifically speaking, we are (nearly) totally ignorant about experience, so neutral monists who take experience as ontologically primitive are boldly following James’ recommendation. Yet, from an ordinary human vantage point, nothing is more familiar than experience. What one is “totally ignorant of” and what is fundamental depend on whether one’s standpoint is anthropocentric or world-centric. Proper science is world-centric. To take experience or sense data as ontologically primitive is to mistake epistemology for ontology.

⁷ A pan-informational neutral monism which might be regarded as a kind of panpsychism (Stubenberg 2014, pp.13-14) was proposed by Sayre (1976), who writes, “an ontology of informational states is adequate for an explanation of the phenomena of mind, as distinct from an ontology of physical events....It is a reasonable conjecture that an ontology of information is similarly basic to the physical sciences.” Sayre later notes, “The neutral monism I advocate holds that the fundamental principle to which both mind and matter are reducible is not a substance in any sense (Aristotelian, Cartesian, whatever), but is rather [a] structure of a sort that can only be represented mathematically. This structure is what information theorists...call “information.” Sayre is correct in arguing that information is relevant to both matter and mind, but informational ideas are currently insufficient for explaining experience. However, Tonioni’s (2008) theory, discussed below, is precisely an attempt to use these ideas to approach such an explanation.

⁸ If Alexander truly regarded high-level properties as ontologically primitive, yet also as not supplementing fundamental interactions, one is permitted to doubt the coherence of his position.

⁹ These considerations are operationalized with hypergraphs and information theory (Krippendorff 1986; Zwick 2001). For the *maximal* decomposition of ABC, namely A:B:C, where the constituents are taken in isolation of one another, mutual information or “transmission” $T(A:B:C) = H(A:B:C) - H(ABC)$, which is a difference of Shannon entropies, is the organization lost in this simple sum of parts. Equivalently, it is the organization present in ABC. $T(AB:AC:BC) = H(AB:AC:BC) - H(ABC)$ measures the organization that is lost in the *minimal* decomposition of R which includes all three pairs of constituents; it is the inherently triadic and non-decomposable part of the constraint in ABC.

¹⁰ Holistic effects are not exotic; they’re more simply known as “interaction effects.” But when holism is at its most extreme, there can be interaction effects without any main effects. In the limit of maximal holism, any decomposition destroys all the order in the system. This is illustrated graphically by Borromean rings (Livingston 1993; Zwick 2001). For the three-constituent system used here as an example, extreme holism gives $T(AB:AC:BC) = T(A:B:C)$; that is, even decomposing ABC minimally, removing only its inherently triadic constraint, amounts to a total decomposition of ABC into its separate constituents.

¹¹ Ontology and epistemology should be conceptually distinguished (see Note 6), but they are not completely separable. In the nonlinear dynamics of chaos, future states of a system are determined but not predictable because of sensitive dependence on initial conditions. It is unclear whether this unpredictability should be regarded as ontological or epistemological, since although predictability might seem to be an epistemological issue, the impossibility of long-term prediction is an in-principle limitation that stems from the necessarily finite precision of any measurement. What ultimately emerges in a concrete chaotic system is not reliably ascertainable by simulation, and thus might be regarded as strong emergence.

¹² O’Connor and Wong (2012) regard Bedau’s (1997) notion of weak emergence as epistemological, but Bedau himself claims that phenomena that in principle are explainable only by simulation exemplify ontological emergence. He argues that the need for simulation follows from “the system’s underlying microdynamic *whether or not we know anything about this* [italics added].” This illustrates the complexity of interpretations of emergence, and the difficulty of assigning individual thinkers to different positions.

¹³ Phenomena demonstrable only by simulation can be further explained (see, e.g., Szabo & Teo 2013).

¹⁴ Fallacies of composition assume that properties of parts must necessarily be present in the wholes that they constitute; fallacies of *decomposition* assume that the properties of wholes must be present in the parts of which they are constituted. These are both not only logical fallacies but also empirical fallacies.

¹⁵ Nagel gives no reason to think that genomic change rates over history are insufficient to account for the diversity of life. He might better have focused on other problems that afflict a narrowly conceived neo-Darwinism; for example, its exclusively functional-historical character and its need for structural augmentations (Kauffman 1993).

¹⁶ The origin of life was never intended to be encompassed in Darwin's theory, but neo-Darwinism includes molecular biology, from which one might expect an account of the origin of life.

¹⁷ The awkwardness of Bunge's use of the materialism label is exemplified in his chapter titles of "physical matter," "chemical matter," "living matter," "thinking matter," "social matter," and "artificial matter." These would have been more appropriately called physical systems, chemical systems, etc.

¹⁸ The three-domain framework of matter, life, and mind was central to Alexander's (1920) ontology and also to Morgan's (1923). Alexander had a fourth domain, deity, which he said emerges from mind. This will not be addressed here, except to note that Boulding's (1956) hierarchy, from which Table 2 is adapted, has a "transcendental" level at its highest point. Alexander also had a level beneath matter, namely integrated space-time, which was the fundamental basis for the existence of relations.

¹⁹ Nagel's assertion that psycho-physical reduction will forever be a failure is a statement of his belief that this second step of understanding the relevant empirical data will forever elude us, even if the first step of acquiring this data is accomplished.

²⁰ An aside on levels (3) and (4): The assertion that genomic information governs metabolic autopoiesis differs from the perspective of Maturana and Varela (1980), who formulated the idea of autopoiesis. For these authors, autopoiesis is self-organization that generates its own boundaries, in which information is dispersed throughout the network, rather than being concentrated in a few constituents. In their view also, "genetic information" is observer-dependent and not an objective feature of living systems, but this view would be rejected by nearly all biologists. As defined here, an autopoietic system is one that is closed organizationally but open to a matter-energy flux, where the internal processes utilize the flux for self-production, and where the partition between system and environment does not require an externally provided boundary. By this definition, some nonbiological systems, such as flames, exhibit autopoiesis, so this property is at most proto-life. "Life" is defined here as encompassing not only metabolic autopoiesis, but also centralization of organizational information, which codes for metabolic catalysts and confers upon the system a degree of individuality and the possibility of reproduction that allows populational evolution. This definition is offered as a statement of the core features of life as a general phenomenon. It is *not* intended as a definition of the minimal requirements that must be satisfied for a *particular* system to be considered "alive." An animal that is sterile is no less alive and individual organisms do not evolve, but life in general requires reproduction and evolution.

²¹ A scheme of this type by Marvin (1912), where lower-level entities are more ubiquitous than higher-level entities, influenced Morgan (Clayton 2004).

²² Since the instantiation of information in matter, life, and mind appeared sequentially in history, this raises the question of teleology: is there, as Teilhard (1959) thought, a tendency in cosmological evolution towards the distillation, refinement, and concentration of information? Nagel also asks us to consider the possibility of teleological laws that supplement causal laws, though such additional laws would not be needed if one posits a strong version of the anthropic principle, where basic physical constants of the universe are mysteriously "tuned" to allow for and even favor the emergence of life and mind.