

 Open access • Journal Article • DOI:10.1007/S10699-014-9349-Z

Beyond Physics? On the Prospects of Finding a Meaningful Oracle

— [Source link](#) 

Taner Edis, Maarten Boudry

Institutions: Truman State University, Ghent University

Published on: 02 Mar 2014 - Foundations of Science (Springer Netherlands)

Topics: Physicalism, Oracle and Hypercomputation

Related papers:

- [Problems with the 'physical' in physicalism](#)
- [>How to avoid implying that physicalism is true: a problem for teachers of science](#)
- [Formulating physicalism: Two suggestions](#)
- [Deprioritizing the A Priori Arguments against Physicalism](#)
- [Absolute Versions of Incompleteness Theorems](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/beyond-physics-on-the-prospects-of-finding-a-meaningful-4p7so44e24>

Beyond Physics? On the Prospects of Finding a Meaningful Oracle

Taner Edis · Maarten Boudry

Published online: 2 March 2014
© Springer Science+Business Media Dordrecht 2014

Abstract Certain enterprises at the fringes of science, such as intelligent design creationism, claim to identify phenomena that go beyond not just our present physics but any possible physical explanation. Asking what it would take for such a claim to succeed, we introduce a version of physicalism that formulates the proposition that all available data sets are best explained by combinations of “chance and necessity”—algorithmic rules and randomness. Physicalism would then be violated by the existence of oracles that produce certain kinds of noncomputable functions. Examining how a candidate for such an oracle would be evaluated leads to questions that do not admit an easy resolution. Since we lack any plausible candidate for any such oracle, however, chance-and-necessity physicalism appears very likely to be correct.

Keywords Physicalism · Chance and necessity · Computability · Supernaturalism · Intelligent design · Hypercomputation · Randomness

1 The Limits of Physics

The natural sciences, particularly mature sciences such as physics, have enjoyed great success in explaining the phenomena of our world. And this success has motivated more comprehensively naturalistic ways of thinking, casting doubt on the existence of anything beyond nature (Edis 2002, 2006; Boudry et al. 2010).

T. Edis (✉)
Department of Physics, Truman State University, 100 E Normal, Kirksville, MO 63501, USA
e-mail: edis@truman.edu

M. Boudry
Department of Philosophy and Moral Sciences, Ghent University, 9000 Ghent, Belgium
e-mail: maartenboudry@gmail.com

M. Boudry
Konrad Lorenz Institute, Vienna, Austria

The most ambitious form of such a view, physicalism, suggests that *everything* we interact with (setting aside questions about the status of causally impotent abstract objects) is in some sense reducible to physics (Melnyk 2003). Social, cultural, and biological phenomena supervene on physical processes—there can be no change in higher-order processes in our world without an underlying physical change.

Physicalist ambitions naturally raise questions about the limits of physics. The present debate often turns on whether and how minds can be realized by physical systems, such as brains and perhaps even machines. Scientists and philosophers defuse dualist intuitions, while suggesting mechanisms by which brains produce conscious experience (Dennett 1991; Damasio 2010; Churchland 2002). Arguments making connections between psychophysics and qualia (Clark 1993; Edelman 2004), physicalists think, take us closer to understanding how minds are physically realized.

The cogency of physicalism depends on what exactly it means for something to be physical. In the literature, a distinction is made between a theory-based conception, which ties the notion of physical to the properties described by certain theories, and an object-based conception, which refers to intrinsic properties of paradigmatic physical objects (Stoljar 2009).

This debate would be better served if we had a clearer idea of the limits of physics. The prospects for a more complete neuroscience firmly anchored on physics may be good, but without an ability to investigate an unambiguous failure in physical explanation, such arguments may issue too many promissory notes. Physical science will certainly advance, but how far? In addition to hopes for future science and various philosophical intuitions, it would help if we had an argument driven by *data*. Physicalism would be a clearer and perhaps stronger position if we knew how it would be vulnerable to refutation.

One complication now arises, which is known as Hempel's dilemma (Hempel 1969). We would like to explore the limits of physics per se, not just the limits of present-day theories. We need data that would not just lead to the replacement of one theoretical framework with another, but rather signal a breakdown in physical explanation construed much more broadly. In the development of classical physics, where at first a mechanical picture of physical objects and their interactions predominated, the action-at-a-distance of gravity, and initially somewhat occult-seeming phenomena such as electricity and magnetism, were assimilated into physics without too much of a strain. Modern physics significantly departs from the mechanistic materialism popularized in nineteenth century Europe (e.g., Büchner 1884), but physicalism is still an updated descendant of materialism. Our broad understanding of physics cannot depend on details such as whether fundamental objects are taken to be particles, fields, or M-branes, or whether interactions are local or nonlocal in character.

In other words, physicalists should not commit themselves to the physical theories of their time. But on the other hand, according to Hempel, if physicalism is defined to include any future revision of our current theories, it becomes impossible to evaluate. Who knows what our physical theories will look like in a hundred years? To steer a way out of Hempel's dilemma, we need a conception of physics flexible enough to cover future extensions of physics, but not so broad as to encompass anything we could ever discover by empirical means.

Our historical experience recommends such a moderately broad view. Electromagnetism did not, at first, comfortably fit in with earlier physics. Today, only cranks associate electricity and magnetism with the occult or supernatural. But the integration of magnetism in a physical framework is a contingent state of affairs, not the result of a trivial extension of the term "physical." If electromagnetism had turned out to be a mysterious life force—if it animated bodies but was otherwise fundamentally disconnected from material entities and processes—physicalism would then have become much less plausible. Likewise, if electromagnetism had

exhibited mind-like properties not further reducible to impersonal processes, physicalism would have been in trouble.

2 Supernatural Agency and “Intelligent Design”

A usefully broad definition of physical explanation can start with the commonly-drawn contrast between physics and the supernatural. Physical objects are made up of mindless components; physical processes are impersonal. Physicalists think that phenomena such as life and mind are complex processes that are assembled “bottom-up” out of lifeless and mindless physical interactions (Dennett 1995; Edis 2006). In contrast, supernatural beings are conceived as being beyond physical restraints. Tales of the supernatural are full of agents such as ghosts, gods, and demons, who are anthropomorphic in their psychology but who violate expectations about ordinary embodied agents (Boyer 2001; Atran 2002; Atran and Henrich 2010). Theistic traditions elaborate notions of agency beyond physics into a top-down, mind-first picture of the world in which instances of interesting complexity such as life and mind are due to divine design (Nasr 1989).

This is not to say that any ostensibly supernatural phenomenon would signal a breakdown of physics. For example, while parapsychologists have often claimed a scientific revolution overturning materialism (Rhine 1953; Radin 1997), it is harder to say that a substantiated parapsychological claim would go beyond any conceivable future physics. At present, attempts to link alleged psychic phenomena with quantum physics tend to degenerate into a quantum mysticism that attracts few physicists (e.g., Goswami 2001). But the notion that consciousness and quantum measurement are connected also comes in versions that, though typically judged wrong by neuroscientists, are still within the bounds of respectability—such as efforts by Hameroff and Penrose (1996). These let us imagine that *if* convincing evidence for parapsychological phenomena were to surface, a future physical theory *might* still account for them. ESP, if real, could still be due to a channel of information transfer not anticipated by today’s physics.

There are, however, other claims on the fringes of science that challenge the very possibility of physical explanation. A prominent example today comes from “intelligent design” (ID) creationism. ID proponents sometimes describe their efforts as a research program to identify the products of intelligence, including, perhaps, alien intelligences that could communicate with us. However, ID also has strong theistic overtones (Forrest and Gross 2007), and especially in their more popular and philosophical presentations (Dembski 1999; Johnson 2000), ID advocates draw a sharp distinction between what is physical and what is allegedly irreducibly intelligent and purposeful. While the ID movement has attracted attention due to its opposition to Darwinian evolution (Dembski 2004; Meyer 2009) and promotion of dubious astrophysics (Gonzalez and Richards 2004), it also is a dualist research program. The fundamental concern of ID is *intelligence*, not biology—many ID proponents target evolution because Darwinian evolution, by providing a physical explanation of functional complexity, strongly suggests that minds are part of a physical order (Edis 2004a).

Now, the kindest comment that can be made about ID is that the scientific community has found it thoroughly unconvincing (Young and Edis 2004). Even so, unlike parapsychology, ID represents an uncompromising and unambiguous opposition to physicalism. Even as a failure, the efforts to develop ID should contain useful hints to where we might find the limits of physics.

A common intuition driving ID is that the generation of *information* is beyond physics. William Dembski, the leading theorist of the ID movement, has repeatedly attempted to

mathematically specify the domain of what cannot be achieved by physical processes. The signature of intelligent design should be identifiable as a pattern in data, where this data might be a text, a DNA base sequence, or even radio communications from alien intelligences (in which case it need not directly signal a breakdown of physics).

Plausibly enough, Dembski and kindred ID proponents identify the basic characteristic of naturalistic, physical explanations as a reliance on chance and necessity. Both haphazard chance and rigidly rule-bound phenomena that manifest physical laws are fundamentally mindless. Historically, materialists have relied on chance and necessity as alternatives to divine design or other supernatural claims. Darwinian evolution is just one example; it is no accident that naturalistic classics such as Jacques Monod's *Chance and Necessity* (1971) place modern biology in this broader tradition of rejecting supernatural entities.

Physics, having jettisoned final causes, has come to rely entirely on mathematical models that conjoin lawful, rule-based behavior and, especially with the advent of quantum mechanics, randomness. So a more rigorous development of what "chance," "necessity," and their combinations stand for can lead to a variety of physicalism that is both suitably broad and has clearly defined limits.

In his early arguments for ID, Dembski presented an "explanatory filter" for first eliminating lawful processes as an explanation of a data set, then also ruling out pure chance, and concluded that if the data also fit a prespecified, meaningful pattern, it must be the result of design (Dembski 1998). This proposal was severely flawed, in particular because it did not adequately address *combinations* of chance and necessity such as Darwinian variation-and-selection. Later, Dembski tried other arguments, for example, misusing the No Free Lunch theorems to claim that "specified complexity" was beyond the reach of physical processes that did not already embody information about a target (Dembski 2002; see Häggström 2007; Perakh 2004; Young and Edis 2004).

As a research program, this version of ID is a non-starter. There is an obvious way to make the notions of chance and necessity rigorous, however, by availing ourselves of concepts developed in theoretical computer science. We know what can and cannot be accomplished by machines capable of executing algorithms and calling on a true random number generator. Nothing remotely like "specified complexity" is inaccessible to combinations of rules and randomness (Edis 2004b).

But there are still tasks that are beyond chance and necessity, and hence are plausibly beyond physics. These are found among a subset of noncomputable functions. Computer scientists often invoke an "oracle" to study decision problems, including oracles that produce noncomputable functions. For a research program seeking to establish a clear-cut example of a breakdown of physics, such noncomputable oracles—perhaps justifying the supernatural connotations of the word "oracle"—would be a good place to start. So if we are interested in the limits of physics, we can ask if anyone could conceivably succeed in finding a signature in data sets that would make us suspect something beyond physics. We can, in other words, ask how we would know if we were looking at the output from an oracle beyond physics.

3 Chance-and-Necessity Physicalism

3.1 Finite Resources

We now propose a version of physicalism that more formally expresses the claim that everything we interact with is produced by mindless processes combining chance and necessity.

For this, we present two separate claims: one about our capabilities of modeling the world, and another about the world at large.

Let us start with ourselves:

1. When doing physics, we only have access to finite computational resources.

Proposition (1) should be uncontroversial, as it only means that physicists disavow godlike powers. Nothing physicists do requires use of an infinite number of bits of memory or an infinite number of steps in a computation. Still, it needs some unpacking, since we do not want to claim that physical theories are a species of computer program.

It is easy to accept (1) in the context of experiments. Data sets are finite, and data analysis is limited to finite memories and finite numbers of processing steps. However, we usually describe physical variables with real numbers, which can contain an infinite number of bits. For example, a qubit, unlike a classical bit, can represent an infinite amount of information in the real numbers that determine a superposition state. A digital computer cannot simulate a quantum world in complete detail (Feynman 1982; Mermin 2007). But there is nothing specially quantum mechanical about real numbers: any classical variable that takes real values can also store an infinite number of bits.

Proposition (1) only demands, however, that only a finite amount of information be *accessible*. A quantum computer is not a device for processing infinite information in a finite amount of time. The mathematical models used by physicists make use of constructs such as real numbers or functions, but these models involve a finite number of discrete symbols manipulated in a finite number of steps, and they produce finite amounts of information to compare with experimental data sets.

Some examples will help. Say we have a theory that produces “ π ” as a prediction. This does not require infinite information, since we have finite algorithms that produce π to any desired precision. The resources needed to compute the 10^{80} th digit of π may be large but they are finite.

Also consider a case where infinite resources would be needed for an exact result. Some solutions to certain differential and partial differential equations are expressed by a non-computable number—call it η —that cannot be exactly computed by any finite algorithm (Pour-El and Richards 1979, 1981, 1989). In that case, to compare with experiment, a finite approximation to η would be required. The number η will then stand not for any particular algorithm, but a set of approximation algorithms, all of which demand finite resources, and only a finite number of which will be available to a physicist at any time. And the equations leading to η , and the process of approximating a solution, all involve a finite set of symbols that are manipulated in a way that demands finite resources.

3.2 Randomness

Physicists do not just compare exact predictions to experimental data. They also account for the unpredictability of individual events by making predictions about the statistical properties of data obtained about a series of events. Here, physicists rely on randomness, which is a particular kind of nonalgorithmicity.

The best definition of randomness in this context is obtained through the notion of algorithmic or program-size complexity (Kolmogorov 1965; Chaitin 1987). We can attempt to represent any finite bit sequence in a compressed form, by finding an algorithm that allows us to reduce the number of bits required to specify the sequence. But the vast majority of sequences are not compressible: we can do no better than to individually list each bit in the sequence. Incompressibility gives us a definition of randomness for finite sequences:

“algorithmic randomness.” At the infinite length limit, we obtain a rigorous definition of randomness equivalent to the notion of flipping a fair coin. The set of infinite bit sequences (which can be mapped one-to-one onto real numbers or integer functions) includes random sequences that are, in a sense, maximally nonalgorithmic (Chaitin 1987, 2001). Almost all infinite sequences are random.

Randomness in a mathematical model does not produce definite predictions for individual events. Instead, randomness denotes a set of possible outcomes, to be statistically tested through a series of trials.

A sequence of bits with $1/2$ probability of a 1 or 0 is merely the beginning. To get more interesting distributions, we have to combine algorithms and randomness. For example, say we want to denote the set of sequences where the probability of a 1 is $1/4$ instead. Now imagine we had a universal computer (a device that can compute any computable function) that could also make use of a true random number generator. This machine could call a subroutine $\rho_{1/2}$ that produces a random infinite sequence. To produce the n th bit in an instance of a $1/4$ bit probability sequence, the computer would follow the instructions: “To obtain $\rho_{1/4}(n)$, call $\rho_{1/2}(2n)$ and $\rho_{1/2}(2n + 1)$. If both these calls result in a 1, output 1. Otherwise, output 0.” This set of instructions run on our augmented universal computer would produce instances of $\rho_{1/4}$, the set of sequences with bit probability $1/4$.

We can extend such a procedure to get sequences of bits with a specified probability x for a 1, which we can symbolically express and manipulate as ρ_x . And this further generalizes to sets corresponding to much more elaborate probability distributions, much as with models producing definite predictions such as η or π . Now the predictions will be statistical in nature—for example, an average—to be compared with statistics obtained from experimental data sets. Again, the model generates predictions associated with sets of approximations; the key is that each component—the mathematical model, the approximations, and the statistical predictions—relies solely on finite resources.

As a more concrete example, consider how a physicist models the formation of a sixfold symmetric snow crystal. The growth process is partly determined by rules: microscopic constraints such as the H_2O bond angle and the physical details of crystallization determine the sixfold symmetry of a snowflake growing in an idealized environment. But the growth process is sensitive to noise, or random fluctuations. Snowflakes, then, would be modeled by a combination of rules (algorithms) and randomness, specifying a set of outcomes that can be statistically compared to experimental populations of snowflakes.

Everything physicists do, provided proposition (1) holds, fits this general description. Indeed, this is true for all mathematical models in the sciences. A population geneticist models a population through constrained and iterated processes such as variation-and-selection and genetic drift—all of which are combinations of rules and randomness, and all which are compared statistically with available data. Note, however, that this does not mean that a mathematical model is directly equivalent to a particular program on a computer with an ability to call a random number generator. The randomness invoked in a model denotes sets of interchangeable sequences, not any specific random sequence. And a model is associated with a set of algorithms, only a finite subset of which is available to the modeler at any time. These algorithms produce approximate calculations when comparisons with data are needed.

Clearly, there is a lot that combinations of rules and randomness can capture. But not everything. To understand just what rules and randomness can give us, we need a “completeness theorem” (Edis 1998). *Every* infinite bit sequence can be expressed as a combination of a finite algorithm and a specific infinite random sequence. Therefore every sequence must be a member of some set designated by a combination of a finite algorithm and $\rho_{1/2}$. Because the algorithmic part is always finite in length, the description of the set in which any particular

sequence falls is always finite, as provided for by proposition (1). However, tasks that require a *specific* random sequence still demand infinite memory or computation time. For certain tasks, like testing a model statistically, it is not important which exact random sequence is instantiated. All elements of the set denoted by $\rho_{1/2}$ are interchangeable. But if we want to perform a task for which specific random sequences are not interchangeable, then this task cannot be accomplished by a combination of rules and randomness.

3.3 Non-interchangeable Randomness

Such tasks are performed by noncomputable functions, for which we can imagine that an oracle is available. However, not just any noncomputable function will serve our purpose. After all, a true random number generator, according to the standard definition, is an oracle producing a noncomputable result. An oracle that would be a candidate for a violation of physicalism would have to be a *meaningful* oracle, indicative of intelligence.

Consider the classic example of a noncomputable function, Turing's halting function $h(p, i)$. This takes a program p and its input i as its arguments, and returns $h(p, i) = 1$ if $p(i)$ halts and produces output, and $h(p, i) = 0$ if $p(i)$ fails to halt. Turing's h has some algorithmic structure—it is possible to compute $h(p, i)$ for some p and i —but no finite algorithm can compute $h(p, i)$ in general. By the completeness theorem, h must be expressible as a combination of rules and a specific random sequence. However, this specific random sequence must be known exactly in order to compute h . This requires an infinite number of bits of information.

A meaningful oracle, then, would not just produce an arbitrary partly random function, but a noncomputable function that *performs a well-defined task*. We must be able to provide a finite definition specifying the function produced by the oracle. Turing's h would be just such an oracle: we know exactly what it is defined to perform. Another example is Chaitin's Ω : a random number completely defined in terms of halting probabilities, or, equivalently, exponential Diophantine equations (Chaitin 1987)—knowing Ω requires knowing every bit in a *specific* random sequence.

An arbitrary random function, in contrast, has no finite definition specifying its exact sequence. Being able to consult a random number generator is of little help in computing $h(p, i)$ or Ω , since a specific random sequence is required to obtain h or Ω . Indeed, because there are only a countable number of definitions for meaningful oracles, almost all random functions, with probability 1, cannot be used to implement any meaningful oracle.

In that case it is meaningful oracles, and only meaningful oracles, that cannot be implemented by combinations of rules and randomness. (From now on, for brevity, by "oracle" we will mean "meaningful oracle" in this sense.) ID proponents' "specified complexity" cannot be the signature of anything beyond physics. But an oracle producing a well-defined non-computable function is exactly what ID requires: meaningful oracles cannot be assembled by mindless physical processes combining chance and necessity.

3.4 Strong and Weak Physicalism

This brings us to our second proposition in defining chance-and-necessity physicalism, a claim about the world at large. First, the strong version of the claim:

2a. No meaningful oracles exist.

This would mean that all experimental data sets can be accounted for by combinations of rules and randomness.

Proposition (2a) allows for the world at large to contain an infinite number of bits of information. The world may indeed be an extremely nonalgorithmic place. What (2a) claims is that all of this nonalgorithmicity is due to randomness. The specific random sequences that actually exist are meaningless, interchangeable, without functional significance. In other words, the world may be awash with random noise, but it does not contain anything that can be harnessed to produce $h(p, i)$, Ω , or any other specified noncomputable function.

The strong version of chance-and-necessity physicalism directs us toward oracles as possible phenomena beyond physics. However, we also need to propose a weaker version. This is because, as with the example of η discussed above, there may be physical contexts that can lead to noncomputable numbers. The literature on hypercomputation explores some possibilities, though none appear physically realizable. For example, a Malament-Hogarth space-time can allow a Turing machine infinite time for calculations along a world-line, which remains a finite time interval for other observers (Etesi and Nemeti 2002). But such a calculation would still require an infinite amount of energy, and thus is not a good candidate for a nonmiraculous oracle. There are other possibilities for “infinite time Turing machines” (Hamkins and Lewis 2000) that might *perhaps* be allowed by some kinds of physics, though the details are hard to flesh out.

In the spirit of allowing for further developments in physics and hypercomputation, however, let us say that some noncomputable η could be implemented in a physical context. This would still challenge physicalism, as the very availability of η could be construed as a miracle. Consider, once again, an intelligent design perspective. If we were to establish the existence of an oracle producing η , an ID proponent would argue that mindless physical processes could not account for the *origin* of η . After all, finite combinations of rules and randomness cannot assemble the infinite resources implicit in the exact computation of η .

Such an emphasis on origins is, in fact, quite usual for ID arguments. When ID proponents propose alleged examples of biochemical design (Behe 1996, 2007), they do not revive the notion of chemical souls—they accept that only ordinary physical processes are at work in their “molecular machines.” Instead, they argue that blind physical processes—chance and necessity, as in a Darwinian mechanism—cannot account for how such machines came into existence.

An ID argument based on a physically implemented oracle would, in fact, be analogous to a design argument based on cosmological fine-tuning. Instead of insisting that the universe operates by nonphysical means, fine-tuning arguments suggest that we need to go beyond physics to account for the origins of unlikely features of the universe. Presently, fine-tuning arguments are not very promising (Stenger 2011; see also Barnes 2011; Stenger 2012). Physicists have had considerable experience with fine-tuning problems, both in cosmological and smaller-scale physics. Often, fine-tuned physical parameters signal that some new physics is on the horizon; it is hardly a sign that the limits of physics are near. The existence of an oracle producing a meaningful, well-defined η , however, could anchor a fine-tuning argument that cannot be undermined by future physics, since no finite combination of rules and randomness could produce η .

There needs, however, to be more to a successful ID argument than just the production of a function such as η , as our definition of a meaningful oracle suggests. Randomness, after all, refers to brute facts—what we have to list as given, since we are unable to assimilate it into an overall algorithmic pattern. As other critics of ID have noted, going beyond the limits of physics requires that one comes up with an explanation that has more substance than “the Designer just did it” (Boudry and Leuridan 2011). If the existence of an oracle for η made sense in the context of agency and purpose—if η accomplished the sort of computational task that an intelligent Designer analogous to human designers could be expected to have

an interest in—then the ID project could claim some success. Otherwise, η would just be a random feature of the universe.

Some oracles, if they were to exist, would be much more promising than others in this regard. For example, the function performed by Turing's h , or an oracle to provide perfect compression of data, comes up naturally in the context of computation. Such oracles would be very useful for an intelligent agent. And since such oracles are naturally meaningful in the context of intelligent activities, it would be hard to see them as mindless, random phenomena, as mere accidents of fundamental physics. By contrast, they could easily be integrated into an explanation of origins that refers to an intelligent agent able to command infinite resources.

Therefore we also present a weak version of the claim of chance-and-necessity physicalism:

- 2b. If any oracles exist, they have no deeper explanation than that they exist as random facts: accidental features of physics.

A necessary condition both for violating the strong (2a) and weak (2b) versions of chance-and-necessity physicalism is the existence of meaningful oracles. Testing the boundaries of physics, then, must involve a search for oracles. We should look for data that is best modeled by a function that performs a well-defined task, but is not implementable by combining rules and arbitrary, interchangeable random numbers.

4 Detecting Oracles?

4.1 Compressibility

To test either version of chance-and-necessity physicalism, we need to know how to recognize a meaningful oracle. And we can only test an oracle claim against a finite data set obtained through experiment. There are an infinite number of functions that fit every finite sequence.

The problem is difficult, but not hopeless. We begin with how we recognize randomness in physics. After all, a random function is also not computable by any algorithm. And how we recognize randomness is in any case an important question for a chance-and-necessity version of physicalism.

An obvious starting point is the algorithmic randomness of finite data sets. In other words, we look at the data and try to find some pattern in what we observe. Any pattern can be expressed as an algorithmic structure within the data that results in compression. Incompressible parts of the data indicate that the best model is a combination of the rules that specify the found pattern and the randomness that produces incompressibility.

There are some technical complications here. First, algorithmic randomness is itself noncomputable—we have no way of ensuring that in all cases where we think the data is incompressible, this is in fact the case. Determining compressibility is a matter of using a set of finite resource-using, computable approximations, in keeping with proposition (1). In other words, we use the limited means at our disposal to find algorithmic patterns as best as we can. Where we fail, we can do no better than list the data bit by bit. This incompressibility is exactly what algorithmic randomness is; therefore, we tentatively model what cannot compress as random. We might later improve our means of finding patterns, in which case we may revisit what we thought was random and find otherwise. Nonetheless, especially with large data sets, randomness is a good bet even though our means for data compression are always limited, since the number of incompressible sequences grows to be overwhelmingly large when compared to compressible sequences. Indeed, in the infinite limit, nonrandom sequences are a zero-measure set.

A second complication is that algorithmic randomness is always relative to a particular universal computer U . Any finite data sequence that is incompressible is always compressible with respect to a redefined universal computer U' . All we have to do is to incorporate the sequence into the instruction set of U' , in which case the data becomes trivially compressible. So our modeling efforts have to take such complications into account. For example, it would be illegitimate to adopt U' after the data is taken, since then the modeling apparatus would not be independent of the data. In the long run, our background theories that frame the data, along with the rest of our modeling apparatus, are subject to change. Confidence in our modeling ability requires that this change should be adaptive. For our purposes, we assume that the mutual correction between theory and experiment that produces convergence on some remarkably solid results in mature sciences like physics illustrates a successful way around this particular complication.

4.2 Chaotic Systems Versus True Randomness

With these complications taken into account, the role of randomness in a physical model is to account for data that do not conform to an algorithmic pattern—where individual events remain unpredictable, but large populations of events allow the definition of statistical variables that are predictable.

The best example of randomness in physics is quantum mechanics. The theoretical apparatus of quantum mechanics predicts probability distributions, not individual events. In our nearly century-long experience with quantum physics, we have obtained no data exhibiting any pattern that allowed us to improve upon the merely statistical predictions produced by quantum mechanics. So, for example, the exact time of decay of any individual radioactive nucleus is completely unpredictable as far as we know. All we can do is calculate probability distributions or approximately calculate probabilities according to the resources we can devote to the calculation. But if we have a large population of identical nuclei, a condition approximated very well with macroscopic samples, we can make extremely accurate predictions about statistical variables such as the half-life of a sample.

It bears emphasizing that the randomness exhibited in quantum mechanics is, as far as we can tell, true randomness. For example, quantum unpredictability is different from the unpredictability exhibited by chaotic classical systems that are sensitive to initial conditions. In a chaotic system with positive Kolmogorov-Sinai entropy, the predictive utility of existing information about a system decreases rapidly (Gaspard 1992). However, in such systems it is possible, in principle, to extend predictive abilities indefinitely by obtaining increasingly precise information about the initial state. Quantum systems do not allow for any such extension: the probability distribution is all we can predict, regardless of any amount of extra data collected.

Debates about the interpretation of quantum mechanics are also beside the point in this context. Even if we adopted, for example, a Bohmian hidden variable interpretation, information about these variables would be inaccessible. Such interpretations displace the randomness from the dynamics to the boundary conditions, but they have no effect on the randomness in the predictions generated by quantum theory.

4.3 Signs of Oracles

Now we can use this experience with randomness to suggest ways to investigate the existence of an oracle producing a well-defined noncomputable function. The fact that we are dealing with finite data sets should not matter. The signature of our oracle would be finding that bits

in the incompressible sector of a function description are not arbitrary or interchangeable: that they can be associated with a well-defined task performed by an independently described oracle.

We cannot test an oracle claim by comparing data to the output of a known oracle. We can, however, construct finite, algorithmic approximations. Therefore we propose that if we were to find that a data set consistently conforms to the known output of an oracle, and that this conformity continues upon extensions of the data set, we would have some evidence for the existence of an oracle.

Imagine that a new religious movement were to arise, led by a prophet who claims to perform a miracle appropriate to the computer age. When presented with a computer program, p , and its input, i , the prophet goes into a very brief trance and promptly announces whether $p(i)$ halts. In other words, the prophet claims to produce Turing's $h(p, i)$.

Not being able to compute $h(p, i)$, we would have to check the prophet's proclamations against approximations to h . We can, for example, scan p to find certain infinite loops that would prevent halting when fed with input i . If such loops exist, then we know that in this case, $h(p, i) = 0$. We can also implement $h_n(p, i)$, which runs $p(i)$ for n steps and reports whether it halts during this time. If $h_n(p, i) = 1$ for any finite n , then $h(p, i) = 1$ also. In other words, we have an always finite but also extensible bag of algorithmic tricks to get a handle on $h(p, i)$.

Now, say that with our best efforts to test the prophet, we always find that the prophetic output corresponds to $h(p, i)$ when we know the value of $h(p, i)$. This is analogous to detecting randomness, which comes down to our best efforts not discovering any pattern that allows us to compress part of our data. We think that if our best efforts were to show conformity to an oracle, the claim that an oracle exists should at least become plausible. Depending on how we assess the quality and comprehensiveness of "our best efforts," the existence of an oracle may even become the most reasonable explanation for the data.

We can also imagine oracles producing noncomputable functions that are more useful than h . Since algorithmic randomness is not computable, we need an oracle to tell with certainty if a given string of symbols can be further compressed without loss of information. So it would be very useful if, for example, ID proponents seeking evidence of design at the molecular level were to happen upon a biological computer that could produce such a noncomputable result. Imagine that ID-sympathetic researchers were to discover a bacterium that, given an input RNA string, outputs a compressed form of that string that is minimal in length relative to some universal computer U . We could test such a bacterial oracle by presenting it with sequences that look random according to some simple tests but are in fact highly compressible. For example, suppose we take 100 000 bits from the binary expansion of π , starting with the millionth bit. For good measure, we write the new string backward, change all the bit sequences '0001' of our string into '1000,' and use some other similarly simple algorithms to further scramble the bits in the sequence. The resulting string will look haphazard, but it will in fact be highly compressible. If the bacterium is a genuine oracle, it should be able to reverse-engineer the compressed description of our string, or give an even shorter description if possible.

Now, let us say that the bacterium consistently comes up with correct answers. It even discovers further compressions that no one had anticipated. In that case, our best efforts at testing the bacterium should make us suspect that we have encountered a genuine oracle. Indeed, we would probably consider our experience to be compelling evidence of a supernatural phenomenon.

5 Rival Physical Models

5.1 Pseudorandomness

There is a possible objection to this proposal. We need a computable approximation to a noncomputable function to test whether the oracle producing that function exists. But this test cannot distinguish between any of the infinite computable functions that match this data set. Among these rival models, there is the same approximation that we are using for the test, plus an infinite number of better approximations that still draw on finite resources. How can we say that our data set is produced by a true oracle, rather than by an advanced approximation that is still implemented by physical means?

In some respects, this objection can be deflected by an answer similar to what applies to randomness. Any finite incompressible data set can also be generated by a finite approximation to randomness—a pseudorandom number generator. But unless a specific pattern is found that points to pseudorandomness rather than true randomness, some combination of rules and true randomness will be the better model.

We can show this by an argument that, interestingly enough, draws on some of the same intuitions that animate ID. First, we should note that a pseudorandom number generator, since it is an algorithm, requires a specific structure. In this sense, which we shall make more precise, a pseudorandom number generator is more complex than a true random number generator, and hence more in need of an explanation. Second, again like ID proponents, we can say that in our experience, pseudorandom number generators are associated with design by intelligent agents—computer programmers. In the case of incompressible and haphazard finite data produced by, for example, quantum systems, we know of no such agents. In that case, both the initial implausibility of a pseudorandom generator and the lack of support from our background knowledge make it a much weaker option than true randomness.

In *His Master's Voice*, the science fiction novel by Stanislaw Lem (1983), a group of scientists confront a possible message from outer space: a signal that is incompressible for a long stretch of bits, but which then repeats itself exactly, over and over again. In other words, it is a pseudorandom sequence. A simple signal, periodically repeated, might have a straightforward physical explanation, as with the discovery of pulsars. But this signal is very complex.

The scientists adopt the hypothesis that this is a meaningful message, which is highly compressed, and broadcast repeatedly to the stars as a kind of cosmic beacon. They are frustrated in their attempts to understand the “message,” and also encounter speculative but purely astrophysical hypotheses that do not require an intelligent source. (As with many Lem stories, the *lack* of a scientific resolution is a major theme of the novel.) Their reasoning process, however, is familiar: a pseudorandom phenomenon would be notably odd, in need of explanation in a way that random noise would not be. Our background knowledge suggests an intelligent source as a plausible model, but we also lack any other evidence for intelligent agents responsible for the data. In *His Master's Voice*, the scientists are left in a state where no hypothesis seems to be a clear favorite. If, however, there were no signature of pseudorandomness—a pattern such as the repetition—then true randomness would uncontroversially win over the scientists.

5.2 Program-Size Complexity

We can make this reasoning more formal. We follow [Elsberry and Shallit \(2004, 2011\)](#), who in their critiques of ID point out that instead of Dembski's “specified complexity,” a more appropriate complexity is algorithmic or program-size complexity. In our context, however,

this complexity measure needs a modification. Complexity, after all, is a cost minimum. If we are trying to reproduce our data set bit-by-bit, and program size is the cost we are trying to minimize, algorithmic complexity is appropriate. But when we are modeling data with combinations of rules and randomness, we do not care about reproducing the incompressible sector of our data bit-by-bit—particular random sequences are interchangeable, and only overall statistical properties are important. So the cost of invoking randomness in a model should be small. For example, the cost of using the symbol “ $\rho_{1/2}$ ” in a description of data should be finite and bounded, and for large data sets, this cost will be much less than the program size cost for specifying the random bits individually.

The completeness theorem is proved by showing that all bit sequences can be described with a finite algorithmic part, followed by a trailing algorithmically random sequence that serves as a lookup-table for individual bit values. In the infinite sequence limit, each sequence still has a minimal-length algorithmic part that is always finite in length (Edis 1998). We therefore propose that the appropriate complexity here is the program size for the minimal-length algorithmic part *alone*. This produces a complexity measure where sets of infinite sequences have finite complexity, provided that the specific random sequence that defines the set is irrelevant. When it is relevant, as with meaningful oracles, the complexity becomes infinite. And the complexity of arbitrary random sequences is quite low, corresponding to the notion that there is not much to say about such sequences other than that they are random.

Program-size complexity measures are naturally associated with probability (Solomonoff 1964), capturing the intuition that less complex sequences are more probable. For convenience, let us restrict our attention to universal computers programmed through syntax-free languages (Chaitin 1987), which makes sure that every bit sequence corresponds to an executable program. In that case, the natural probability measure is that each infinite bit sequence—every one interpreted as a program—is equally probable. Low complexity then means that the data can be described by a short minimal-length program and a low-cost, interchangeable trailing random sector. The probability of a set of sequences described by such a minimal program is then correspondingly high.

With all this in place, we can understand how randomness should be favored over pseudorandom alternatives when modeling physical data from, for example, experiments on quantum systems. Randomness is initially more plausible, for reasons that, technicalities aside, come down to the observation that there are far more incompressible sequences than compressible ones. And the additional background knowledge that we have—itsself deriving from modeling of other data sets—make pseudorandom alternatives even less plausible. Our experience with pseudorandomness prompts us to look for mechanisms responsible for generating the more complex pseudorandom output, and no candidate is available.

5.3 Hoaxes and Pseudo-Oracles

Now let us apply a similar form of reasoning to an oracle claim, such as ID proponents happening upon a bacterial data compressor. First, let us specify some more background knowledge that would eliminate some obvious alternative explanations, such as rival bio-engineers having constructed the bacterium and passed it onto the pro-ID lab as a hoax. In fact, let us make sure that something very unusual and beyond our present understanding of physics is taking place.

Therefore, say we discover that data-compressing behavior is not confined to a single bacterial species. In fact, we find that many naturally occurring species have this capability, provided that a small change is made in a noncoding segment of DNA that happens to be shared by many species. No ordinary biochemical mechanism can be identified for the data

compression process—physicists calculate that the process seems very sensitive to quantum fluctuations and express surprise that the bacteria consistently produce the same output for any given input. Moreover, the output is always produced instantaneously, as far as we can tell within the precision of our best instruments. A computable approximation to perfect data compression would take time, since its mechanism would have to execute a series of steps. (Even consulting a lookup table of pre-calculated solutions takes time if the table is large enough.) Given the difficulty of data compression, we expect the approximation to take *many* steps for many inputs. And finally, with every new and improved compression algorithm we discover, we find that the bacterial data compressor had already and effortlessly reached the correct result.

In such a situation, it could still be possible that an extremely powerful space alien was having fun by hoaxing us primitive humans. Such an alien would have to have physical knowledge far beyond ours; perhaps they would have mastered a subquantum physics to which our quantum mechanics is just an approximation. And naturally, the alien would be in possession of much more advanced data compression capabilities than what we are able to manage. Which fantastic possibility, then, seems more likely—a true oracle, or the product of an alien hoaxer that is still operating within the realm of physics, albeit a physics extended far beyond our present knowledge?

Some of our background knowledge counts further against the alien: the bacterial data compressor is our *only* evidence concerning its existence, and what we know about our corner of the universe gives us no confidence in the possible existence of such an alien. Nonetheless, a determined physicalist might still favor the alien hoaxer hypothesis. After all, we have no experience with oracles either. All we have encountered are physically implemented finite approximations, designed by physical creatures originating through physical processes such as Darwinian evolution. An oracle may well seem like an intrinsically more outlandish claim.

To express this intuition, we can again refer to complexity and probability. A determined physicalist can stick to her guns, because a meaningful oracle producing a function such as a perfect data compressor requires an infinite program size. It has infinite complexity, since the random sector in its minimal description has to be known exactly—it cannot be set aside as interchangeable and unimportant and assigned a low cost. Accordingly, its initial probability is zero. It is not strictly impossible, since a zero-measure set need not be empty. But it remains the case that no matter how unsupportive our background knowledge is, a finite oracle-approximation, like an alien hoax, is always more likely.

6 Room for Oracles

6.1 Determined Physicalism

The argument for always favoring a physical alternative depends on the initial improbability of an oracle. Our background knowledge, such as that in the outlandish scenario of the bacterial data compressor, may lead to ambiguous conclusions. But presumably this background could be even more outlandish, making anything even remotely connected to physics as we understand it even less plausible. And since oracles that produce noncomputable functions are not self-contradictory, setting their initial probability to zero is questionable.

Structurally, the determined physicalist position is similar to ID arguments, or creationist improbability arguments in general, so it might inspire some suspicion. There is, however, nothing wrong with using a complexity measure to single out complex phenomena worthy of attention. The mistake in ID arguments is to assume that low initial probability rules out

the viability of physical processes combining rules and randomness—such as Darwinian evolution—for assembling instances of complexity. If there is any substance to ID’s design-detection procedures, it is that they unwittingly detect Darwinian processes at work (Edis 2001). If a determined physicalist can make use of intuitions usually associated with antiphysicalist notions of design, this is an interesting irony, but no more.

The best way to cast doubt on the determined physicalist’s preference for physical oracle-approximations is to abandon ID-like intuitions and directly challenge the notions of complexity and probability in use. Why should oracles be infinitely complex? There is some precedent for associating omnipotent supernatural agents with simplicity. Whether modern versions of such views (e.g., Swinburne 1996) are convincing is highly debatable, and it is not clear how such considerations would apply to oracles. But it is also true that notions of complexity that assign infinite complexity to oracles are developed in the context of implementing functions by physical means. Yes, *physically* implementing an oracle is almost certainly impossible, as it demands infinite resources. But that is another way of saying that oracles are beyond physics—it would beg the question to therefore assign an oracle zero probability.

This is not to say that we should completely ignore the determined physicalist position. Even though “simplicity” in science is fraught with ambiguity (Sober 1994), it is often a useful tool. Like any method in science, it finds a pragmatic rationale in a track record of success. As a rough guide to the initial plausibility of a hypothesis, considerations of complexity can be helpful. If we ever run into a candidate for an oracle, the claim that it really is an oracle rather than a physical approximation should face a heavy burden of proof.

In that case, if we were to face a situation like that of the prophet computing h or the bacterium perfectly compressing data, we would find ourselves living in a Stanislaw Lem story—confronting data that did not just demand a revolutionary adjustment in our scientific knowledge, but data that challenged our conceptions of how to do science.

6.2 Miracles Versus Oracles

Therefore instead of proposing a clear-cut solution, we express a more qualitative judgment. The scenarios of a prophet or a bacterium implementing an oracle have important similarities to other miracle claims. We can draw on the history of belief in the miraculous and philosophical debates concerning miracles and the testing of supernatural hypotheses (e.g., Fishman 2009).

Imagine that ID researchers were to follow up their hunch that there is no such thing as junk DNA, and find a segment of noncoding DNA common to all life. Then multiple labs run samples of this segment through an X-ray diffraction experiment, and all find “© Yahweh 4004 BC” appearing on their images. Should we then conclude that there is an alien but physical intelligence named Yahweh that has been controlling life on Earth (Boudry and Leuridan 2011)?

Say we were to find a Lourdes that actually worked, and worked spectacularly. A noticeable fraction of pilgrims are healed, to the extent of visibly growing lost limbs. And this only works for devout Catholics, not for anyone else. Would we then say that some hitherto-unknown physics must have been at work, with a “health field” coming into a strong resonance with Catholic brain-patterns?

In Anatole France’s *Le Jardin d’Epicure* (1895), the French writer visits Lourdes with a companion. Upon seeing all the crutches and cranes lying there as evidence of miraculous healing, France’s friend remarks: “A single wooden leg would have been quite more convincing.” But France himself argued that, rather than rashly proclaiming a miracle, a “real philosopher” would simply infer that the human body has as yet undiscovered regenerative

properties, like those of the lizard's tail. Maybe so, but we could press a determined physicalist like France with even more outlandish scenarios. Would he always refuse to consider supernatural explanations?

At some point, perhaps if anything like a Yahweh signature or a super-Lourdes were to happen, we should at least give supernatural hypotheses serious consideration. There could well be something beyond physics in such cases, if they were real. Cosmic designers, miraculous healings, psychic powers, ghosts, gods—these are all standard elements of the supernaturalist imagination. Whether in the context of organized religions, tribal belief systems, or individualist spiritualities, most people have expected some such phenomena to be real. Physicalists, or those who hold to any version of naturalism, have not. All this should carry some weight. It is all very well to say and continue to say that further scientific work *might* assimilate ostensibly miraculous occurrences into a physical order. But it does not appear quite legitimate for physicalists to enjoy an all-purpose excuse, any more than a determined creationist should be able to get away with claiming that the fossils must have been created by Satan to tempt the faithful. Denying the very possibility of oracles producing a meaningful noncomputable function would eventually turn physicalism into another cosmic conspiracy theory.

Oracles are not part of any supernaturalist tradition. Nonetheless, they fit. An oracle such as that for $h(p, i)$ is defined in terms of the function it performs, which cannot be implemented by physical means. It requires infinite resources, which hints at omnipotence. If an oracle for $h(p, i)$ were to exist, it would imply that the function it implements—the *meaning* of the oracle—did not have to be realized physically. In fact, an oracle would fit very well with the supernaturalist intuition that meaning and mind transcends mere physics. Therefore, for oracles as well as for more conventional supernatural claims, we suggest that *always* preferring a physical alternative regardless of the evidence is dubious. There can be evidence for a genuine oracle, and compared to conventional supernatural claims, such evidence would be a more direct sign of phenomena beyond physics.

7 Without Oracles

Exploring the possibility of oracles leads us to a number of interesting conclusions.

Chance-and-necessity physicalism: Focusing on combinations of rules and randomness allows us to formulate a new version of physicalism, capturing the notion that all we observe can be accounted for through chance and necessity. This version of physicalism highlights our ability to mathematically model nature as a motivation for thinking that *some* kind of physics accounts for everything.

Chance-and-necessity physicalism bears a family resemblance to other formulations of physicalism or naturalism. It does, however, promote different emphases, particularly on randomness. For example, physicalism often demands an analysis of causality (e.g., [Melnyk 2003](#)), but our version instead directs attention to questions about pattern and lack of pattern. This is appropriate, since in modern physics, macroscopic causality is no longer fundamental, but emerges from the random substrate of quantum mechanics.

Even though chance-and-necessity physicalism to some extent bypasses debates over the reducibility of minds to physics, it does support more familiar versions of physicalism in this regard. After all, one possibility to consider is that human minds can perform oracular feats. It has often been suggested that humans are not algorithmic, and there are recurrent arguments that Gödel's Incompleteness Theorem shows that there is more to human minds than machines ([Lucas 1961](#); [Penrose 1994](#)). However, while it is reasonable to think that no rigid algorithm can capture human thought, the nonalgorithmic flexibility and openness to

novelty humans enjoy can be fully accounted for by combinations of rules and randomness, without meaningful oracles (Edis 1998).

Chance-and-necessity physicalism regards minds as physical phenomena. It emphasizes, however, not debates over consciousness but the prospects of achieving machine intelligence.

We know what would unambiguously lie beyond physics: oracles producing meaningful noncomputable functions. In the strong version of chance-and-necessity physicalism, no such oracles exist. In the weak version, only accidental, unexplainable oracles are allowed. Therefore the plausibility of physicalism would be most directly tested by attempts to produce oracles, especially oracles that can be plausibly related to mind-like activities such as computation.

No one has launched a search for such oracles. We know of none—none associated with human minds, and none elsewhere. If we were to find a candidate, we would still face unresolved questions about exactly when the evidence would support a true oracle as opposed to a finite approximation. Nonetheless, looking for oracles would be the most direct way to launch a data-driven challenge to physicalism and not just our current understanding of physics.

This is not to say that a search for oracles would be the only way to produce evidence relevant to the truth of physicalism. A “© Yahweh 4004 BC” discovery would surely be also relevant. However, like parapsychology, bare intelligent design always invites the question of whether the ostensibly supernatural is truly beyond physics. After all, the intelligent designer may be a physical entity in a larger physical world, who created a kind of sub-world analogous to a computer simulation. This is a common theme in science fiction (Lem 1999; Banks 2004), which prompts interesting philosophical questions about the concept of the supernatural (Boudry et al. 2010). While some evidence of cosmic design is necessary to support the notion of a supernatural designer analogous to a “Great Programmer” (Edis 2002), it still is not sufficient to directly address the question of whether such a scenario must lie beyond the limits of physics. For that, we need oracles.

Our current knowledge strongly indicates that the required oracles do not exist. To some extent, our discussion about the difficulties involved in detecting an oracle is moot. After all, we have no candidates in sight, and very little prospect of getting any. While there has been interest in the possibility of hypercomputation, none of the proposals look physically realizable (Davis 2004). We do not know where to start looking for an oracle that might signify something beyond physics.

But our discussion of oracles should improve our understanding of what exactly it means to endorse physicalism, what kind of evidence an anti-physicalist should look for, and why the current evidence against physicalism is so underwhelming. By exploring what would constitute an unambiguous failure of physicalism—a discovery of a meaningful oracle—we present a challenge to the opponents of physicalism. Since the prospects for this challenge being met seems remote, the case for physicalism must therefore be considered stronger.

Acknowledgments The authors would like to thank Peter Verdée for sharing his expertise on computability and oracles, and Yon Fishman as well as two anonymous reviewers for their insightful comments on an earlier version of the manuscript. The research of the second author was supported by the Research Foundation Flanders (FWO).

References

Atran, S. (2002). *In gods we trust: The evolutionary landscape of religion*. New York: Oxford University Press.

- Atran, S., & Henrich, J. (2010). The evolution of religion: How cognitive by-products, adaptive learning heuristics, ritual displays, and group competition generate deep commitments to prosocial religions. *Biological Theory*, 5(1), 18–30.
- Banks, I. M. (2004). *The algebraist*. London: Orbit.
- Barnes, L. (2011). *The fine-tuning of the universe for intelligent life*. arXiv:1112.464v1 [physics.hist-ph]. Accessed 8 Nov 2013.
- Behe, M. J. (1996). *Darwin's black box: The biochemical challenge to evolution*. New York: Free Press.
- Behe, M. J. (2007). *The edge of evolution: The search for the limits of Darwinism*. New York: Free Press.
- Boudry, M., Blancke, S., & Braeckman, J. (2010). How not to attack intelligent design creationism: Philosophical misconceptions about methodological naturalism. *Foundations of Science*, 15(3), 227–244.
- Boudry, M., & Leuridan, B. (2011). Where the design argument goes wrong: Auxiliary assumptions and unification. *Philosophy of Science*, 78(4), 558–578.
- Boyer, P. (2001). *Religion explained: The evolutionary origins of religious thought*. New York: Basic.
- Büchner, L. (1884). *Force and matter, or, principles of the natural order of the universe. With a system of morality based thereupon*. Translated from the 15th German edition; 4th English edition. London: Asher and Co.
- Chaitin, G. J. (1987). *Algorithmic information theory*. Cambridge: Cambridge University Press.
- Chaitin, G. J. (2001). *Exploring randomness*. London: Springer.
- Churchland, P. S. (2002). *Brain-wise: Studies in neurophilosophy*. Cambridge: MIT Press.
- Clark, A. (1993). *Sensory qualities*. Oxford: Clarendon.
- Damasio, A. R. (2010). *Self comes to mind: Constructing the conscious brain*. New York: Pantheon Books.
- Davis, M. (2004). The myth of hypercomputation. In C. Teuscher (Ed.), *Alan Turing: Life and legacy of a great thinker* (pp. 195–212). New York: Springer.
- Dembski, W. A. (1998). *The design inference: Eliminating chance through small probabilities*. New York: Cambridge University Press.
- Dembski, W. A. (1999). *Intelligent design: The bridge between science and theology*. Downers Grove: Intervarsity Press.
- Dembski, W. A. (2002). *No free lunch: Why specified complexity cannot be purchased without intelligence*. Lanham: Rowman & Littlefield.
- Dembski, W. A. (2004). *The design revolution: Answering the toughest questions about intelligent design*. Downers Grove: IVP Books.
- Dennett, D. C. (1991). *Consciousness explained*. Boston: Little, Brown and Company.
- Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon & Schuster.
- Edelman, G. M. (2004). *Wider than the sky: The phenomenal gift of consciousness*. New Haven: Yale University Press.
- Edis, T. (1998). How Gödel's theorem supports the possibility of machine intelligence. *Minds and Machines*, 8, 251–262.
- Edis, T. (2001). Darwin in mind: 'intelligent design' meets artificial intelligence. *The Skeptical Inquirer*, 25(2), 35–39.
- Edis, T. (2002). *The Ghost in the universe: God in light of modern science*. Amherst: Prometheus.
- Edis, T. (2004a). Grand themes, narrow constituency. In M. Young, & T. Edis (Eds.), *Why intelligent design fails: A scientific critique of the new creationism* (pp. 9–19). New Brunswick: Rutgers University Press.
- Edis, T. (2004b). Chance and necessity—and intelligent design? In M. Young, & T. Edis (Eds.), *Why intelligent design fails: A scientific critique of the new creationism* (pp. 139–152). New Brunswick: Rutgers University Press.
- Edis, T. (2006). *Science and nonbelief*. Westport: Greenwood Press.
- Elsberry, W., & Shallit, J. (2004). Playing games with probability: Dembski's complex specified information. In Young, & Edis (pp. 121–138).
- Elsberry, W., & Shallit, J. (2011). Information theory, evolutionary computation, and Dembski's 'complex specified information'. *Synthese*, 178, 237–270.
- Etesi, G., & Nemeti, I. (2002). Non-Turing computations via Malament–Hogarth space-times. *International Journal of Theoretical Physics*, 41, 341–370.
- Feynman, R. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6/7), 467–488.
- Fishman, Y. I. (2009). Can science test supernatural worldviews? *Science & Education*, 18, 813–837.
- Forrest, B., & Gross, P. R. (2007). *Creationism's Trojan horse: The wedge of intelligent design* (2nd ed.). New York: Oxford University Press.
- France, A. (1895). *Le Jardin d'Epicure*. Paris: Calmann Lévy.

- Gaspard, P. (1992). Diffusion, effusion, and chaotic scattering: An exactly solvable liouvillian dynamics. *Journal of Statistical Physics*, 68, 673–747.
- Gonzalez, G., & Richards, J. W. (2004). *The privileged planet: How our place in the cosmos is designed for discovery*. Washington: Regnery.
- Goswami, A. (2001). *Physics of the soul: The quantum book of living, dying, reincarnation and immortality*. Charlottesville: Hampton Roads.
- Häggsström, O. (2007). Intelligent design and the NFL theorems. *Biology and Philosophy*, 22(2), 217–230.
- Hameroff, S., & Penrose, R. (1996). Conscious events as orchestrated space-time selections. *Journal of Consciousness Studies*, 3(1), 36–53.
- Hamkins, J. D., & Lewis, A. (2000). Infinite time Turing machines. *Journal of Symbolic Logic*, 65(2), 567–604.
- Hempel, C. (1969). Reduction: Ontological and linguistic facets. In S. Morgenbesser, et al. (Eds.), *Essays in honor of Ernest Nagel* (pp. 179–199). New York: St Martin's Press.
- Johnson, P. E. (2000). *The wedge of truth: Splitting the foundations of naturalism*. Downers Grove: InterVarsity Press.
- Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. *Problems Information Transmission*, 1, 1–7.
- Lem, S. (1983). *His master's voice*. San Diego: Harcourt Brace Jovanovich.
- Lem, S. (1999). *A perfect vacuum*. Evanston: Northwestern University Press.
- Lucas, J. R. (1961). Minds, machines, and Gödel. *Philosophy*, 36, 112–127.
- Melnyk, A. (2003). *A physicalist manifesto: Thoroughly modern materialism*. New York: Cambridge University Press.
- Mermin, N. D. (2007). *Quantum computer science: An introduction*. Cambridge: Cambridge University Press.
- Meyer, S. C. (2009). *Signature in the Cell: DNA and the evidence for intelligent design*. New York: HarperOne.
- Monod, J. (1971). *Chance and necessity: An essay on the natural philosophy of modern biology*. New York: Knopf.
- Nasr, S. H. (1989). *Knowledge and the sacred*. Albany: State University of New York Press.
- Penrose, R. (1994). *Shadows of the mind: A search for the missing science of consciousness*. Oxford: Oxford University Press.
- Perakh, M. (2004). *Unintelligent design*. Amherst: Prometheus.
- Pour-El, M. B., & Richards, I. (1979). A computable ordinary differential equation which possesses no computable solution. *Annals of Mathematical Logic*, 17, 61–90.
- Pour-El, M. B., & Richards, I. (1981). The wave equation with computable initial data such that its unique solutions is not computable. *Advances in Mathematics*, 39, 215–239.
- Pour-El, M. B., & Richards, J. I. (1989). *Computability in analysis and physics*. Berlin: Springer.
- Radin, D. (1997). *The conscious universe: The scientific truth of psychic phenomena*. New York: HarperEdge.
- Rhine, J. B. (1953). *New world of the mind*. New York: William Sloane.
- Sober, E. (1994). Let's razor Ockham's Razor. In E. Sober (Ed.), *From a biological point of view: Essays in evolutionary philosophy* (pp. 136–157). Cambridge: Cambridge University Press.
- Solomonoff, R. (1964). A formal theory of inductive inference, part I. *Information and Control*, 7(1), 1–22; A formal theory of inductive inference, part II. *Information and Control*, 7(2), 224–254.
- Stenger, V. J. (2011). *The fallacy of fine-tuning: Why the universe is not designed for us*. Amherst: Prometheus Books.
- Stenger, V. J. (2012). Defending the fallacy of fine-tuning. [arXiv:1202.4359](https://arxiv.org/abs/1202.4359) [physics.pop-ph]. Accessed 8 November 2013.
- Stoljar, D. (2009). Physicalism. *Stanford Encyclopedia of Philosophy*. plato.stanford.edu/entries/physicalism/. Accessed 8 November 2013.
- Swinburne, R. (1996). *Is there a god?*. Oxford: Oxford University Press.
- Young, M., & Edis, T. (Eds.). (2004). *Why intelligent design fails: A scientific critique of the new creationism*. New Brunswick: Rutgers University Press.

Taner Edis (1967) is professor of physics at Truman State University, where he also pursues his interests in the philosophy of science and controversies involving science and supernatural claims. His books include *An Illusion Of Harmony: Science And Religion In Islam* (2007).

Maarten Boudry (1984) is a postdoctoral fellow of the Flemish Fund for Scientific Research (FWO) at Ghent University. In 2011, he defended his dissertation on the epistemic structures of pseudoscience, consisting of a collection of papers that have been published in *Philosophy of Science*, *Philosophia*, *Quarterly Review of Biology*, *Science & Education* and *Philosophical Psychology*. He is co-editor of a volume on the *Philosophy of Pseudoscience* (2013), together with Massimo Pigliucci. His current research deals with the

problem of irrationality from the viewpoint of evolutionary epistemology. Other research interests include the conflict between science and religion, methodological naturalism and skepticism. Together with Johan Braeckman, he also published a Dutch book on critical thinking (2011), aimed at a wider audience.