This article originally appeared in:

Available online at:

This is Daniel C. Dennett's final draft before publication. It has been modified to reflect the pagination of the published version of the work.

PHILOSOPHICAL BOOKS
Edited by ANTHONY ELLIS

Volume xxxviii No. 2   April 1997

CONTENTS PAGE

IN THE BEGINNING, THERE WAS DARWIN

   Like any good preface, that of Dan Dennett's latest brainchild Darwnin
Dangerous Idea: Evolution and the Meanings of Life (Simon & Schuster, 1995. 586
pp. $30.00) sets the tone for the rest of the book: "Highly technical
arguments ... are absent here" (p. 12). What is here is nonetheless a ripping
good story, one goal of which, Dennett submits, is "to get thinkers in other
disciplines to take evolutionary thinking seriously, to show them how they
have been underestimating it, and to show them why they have been listening
to the wrong sirens" (p. 12). I would like to report that Dennett builds his
narrative around a technically flawless underlying framework of the kind of
rigorous and insightful philosophy which has made him famous, that it draws
out salient large scale features of the philosophical and biological terrain
while neither involving the reader in needless gory details nor clashing with
those details. But although this book is generally excellent, evaluated on its
own terms and with respect to its stated goals, as philosophy, technical blemishes
occasionally stand in the way of extracting positions at once precise, significant,
and internally coherent. We leave these faults aside for the moment and look
at the three main parts of the story itself and the over-arching vision
Dennett offers.
   The first section introduces the Darwinian account of design in Nature.
What Dennett calls "Darwin's dangerous idea" is this: that species and all
the design they manifest evolve by descent with modifications-and, signific-
antly, that this is a mindless algorithmic process. Why should we find that
idea dangerous? Maybe, Dennett says, it's dangerous because insofar as
natural selection is a 'substrate neutral' algorithm (or several), it can be "lifted
out of... [its] ... home base in biology" (p. 58) and applied to everything
from the birth of the cosmos through to the subtlest nuances of human
psychology. The level of the Darwinian algorithm, we are told, is the right
level of description for understanding and uniting all occurrences of design
in Nature. Dennett suggests that "all the achievements of human culture-
human language, art, religion, ethics, science itself-are themselves
artifacts . . . of the same fundamental process that developed the bacteria,
the mammals, and Homo sapiens" (p. 144). Darwinism, says Dennett, truly is
a Universal Acid: an idea which leaves nothing untouched.
   Dennett likens components of the mindless processes which build up design
to the cranes of ordinary building construction: by 'crane' he means a feature
or subroutine of a physical process which may accelerate the basic process
but which is itself entirely explicable at a lower level of physical description.
A crane par excellence is a DNA molecule which, once made by processes

81

wholly in the realm of chemistry and physics, permits whole new kinds of organisation (and organisms) to develop. Dennett contrasts cranes with the 'skyhook', a "'mind-first' force or power or process" (p. 76) which doesn't admit of an explanation in terms of lower level processes; a paradigm example is a God who creates life in some miraculous way. The difference between crane and skyhook is itself a sort of linguistic crane which prompts a further useful distinction, this time between good reductionism and greedy reductionism (pp. 80-82): good reductionists hold that everything can be explained without skyhooks, while greedy reductionists think everything can be explained without cranes. (The field might spare itself reams of wasted hedging, accusation, and rejoinder over charges of reductionism if philosophers everywhere adopted this simple convention.) Dennett jettisons greedy reductionism while embracing its mild mannered sibling enthusiastically.

  The book's second part explores Darwinian thinking within biology itself and the challenges it faces there, especially from those who would upset Dennett's explanatory cascade of cranes upon cranes by inserting skyhooks. Whenever some theorist or other has sought a skyhook to support a 'miraculous' property of minds or of bird wings or of anything else, he maintains, science has succeeded only in uncovering yet more cranes.

  A perfect example of a leading theorist searching for skyhooks to limit the reach of Darwin's dangerous idea-and turning up only cranes-is none other than America's "evolutionist laureate" Stephen Jay Gould. Of all the authors on whom Dennett's critical glare comes to rest in this book, none fares worse than Gould. Dennett's critique is shattering, revealing Gould's self-proclaimed 'revolutions' in Darwinian thinking as little more than mild crane-style restatements of what modern Darwinism could already accommodate. Although the near invective sometimes verges on personal (p. 306), it is a sorely needed corrective for some of the damage Gould regularly inflicts on the public's grasp of evolution.

  Part II also sees Dennett draft Conway's Game of Life into service for a crystal clear look at explanations at different levels of description and the role of cranes in making those accounts intelligible (pp. 166-175). In a world where some philosophers seem still not to have grasped the point Davidson made decades ago[1], Dennett's lucid discussion, together with his standard explanatory 'stance' talk (pp. 235-237), is a godsend (but not a skyhook!).

  The highlight of Part III, called 'Mind, Meaning, Mathematics, and Morality', is Dennett's account of meaning. Warnings in the Preface notwith-standing, this discussion verges on serious philosophical argument. Dennett locks horns with Fodor, Putnam, Dretske, Burge, Kripke, and the like and finishes the encounter with his position undeniably more intact than theirs. His evolutionary account of meaning, roughly parallel to Millikan's, is crisp and subtle, with a beautiful new thought experiment thrown in to confound Kim and others (but I wouldn't want to spoil the fun by revealing it!); lie even gives Twin Earth its "proper Darwinian funeral". Less satisfying is Chapter 17, 'Redesigning Morality', a short section-short on conclusions,

---

[1] D. Davidson, 'Radical Interpretation', Dialectica, vol. xxvii (1973), pp. 313 28.

short on connections with Darwin, and short on Dennett's usual sparkling wit and insight-which could almost have been written by someone else.

Continuing on from summary to critique, the first thing which may strike many readers is the sound of wheels being reinvented: Florian von Schilcher and Neil Tennant engaged many of the same or similar topics and drew largely parallel conclusions more than a decade earlier with Philosophy, Evolution and Human Nature(1). The omission of this book from an otherwise broad and remarkably complete bibliography is disappointing, as is the absence of evolutionarily preferable alternatives such as those of Sampson(2) (drawing on work by Herb Simon) and the Montague grammar(3) in Dennett's critical review of Chomsky (pp. 384-393) and the 'language organ'.

More curious are minor but distracting oversights in philosophy of mind and psychology. In discussing the errors of "the lowliest Skinnerian creature", Dennett remarks that, "In order to learn from them one has to be able to contemplate them" (p. 380). But unless 'contemplate' takes on a very broad meaning, or unless learning itself presupposes the capacity to contemplate-and neither usage seems consistent with Dennett's work elsewhere-this is silly. Even some of the simplest artificial neural net models can 'learn' from classical conditioning, yet few would attribute to them the capacity to 'contemplate' their errors. Earlier (p. 374), the discussion also seems to suggest that the selective advantage for early organisms with phenotypic plasticity was restricted to Skinnerian learners. But here we're missing a few steps and a few million years: plasticity in, say, a nervous system surely could confer advantage (for instance, by making possible sensory pattern recognition without 'hard-wiring') long before the phylogenetic emergence of Skinnerian learning. Whether it did or whether Skinnerian creatures appeared before die advent of nerve complexes and non-hard-wired sensory pattern recognition are empirical questions left unaddressed. Finally, Dennett adopts a puzzling approach to minds themselves, suggesting (pp. 341, 369) that they were created by the infestation of appropriate anatomical structures with memes, Dawkin's analogue of a gene for the realm of ideas and culture. Apart from the implausibility of memes actually creating minds-shouldn't the two have coevolved?-this approach accommodates only one of many alternative levels of description where useful explanations might be given, ignoring huge tracts of psychology and cognitive science for the sake of a Darwinian level of description. The view is more extreme than that in Consciousness Explained(4) and, surprisingly, smacks of a questionable brand of anti-reductionism: maintaining explanation at a high level of description to the exclusion of other interesting and perhaps more useful things which might be said at lower levels.

The last relatively minor oversight is Dennett's shifting definition of the word 'algorithm'. Early on, Dennett removes all limits to what qualifies as

1. Routledge & Kegan Paul, 1984.
2. G. Sampson, 'Linguistic Universals as Evidence for Empiricism', Journal of Linguistics. vol. xiv (1978), pp. 129-175.
 3. R. Montague. Formal Philosophy (Yale University Press, 1974).
 4. D.C. Dennett(Little, Brown, 1991)

an algorithm, suggesting we can "treat any process at the abstract level as an algorithmic process" (p. 59). Here, Dennett clearly means to encompass every physical process of any kind. Yet this connotation can no longer be at work when he says "every algorithm is finitely specifiable" (p. 430); this rests instead on the narrower definition-adopted without warning (until p. 437)-equating the set of all algorithms with the set of all Turing Machines. Dennett compounds confusion with the common mistake(1) of assuming that "physics ... is all computable" (p. 446, emphasis original). He returns to the all-encompassing definition later in dubbing nonalgorithmic processes skyhooks and likening their appearance in products of natural selection to "the oracles on the toadstools" (p. 448).

On the narrow formal definition of 'algorithm', nonalgorithimic processes are not skyhooks. We have recently learned, for instance, that certain types of chaotic neural network with analogue components can compute a superset of the Turing-computable functions; they evolve through physical processes which no Turing algorithm can model accurately, yet they manage with neither minds nor oracles on toadstools. More importantly, such networks might well play a nontrivial role in the lives of some of natural selection's most well designed creatures.(3) At issue is the question of the most appropriate model of computation for the grand unifying algorithms Dennett envisions; we return to this question later.

So far these trifles hardly bear on Dennett's overall picture, that of a world entirely united at some one or several levels by one Darwinian algorithm or another. But two more serious complications mar that vision: a dubious extension of Darwinism   into cosmology and a mathematically incoherent representation of design and possible gene sequences. It may well be that Dennett's vision will turn out to be correct, but the case made for it in the book is not, by itself, convincing.

The first doubt centres on Dennett's handling of a view to which Darwin himself was sympathetic: the idea that a God, while not designing species directly, might nevertheless have designed the laws of Nature and set the initial state of the cosmos in exactly such a way as to enable natural selection to do so. The critic of universally applied Darwinism acknowledges the role of natural selection in the development of species but asks: how do we explain the lowest level physical laws which make natural selection itself possible? Since tiny deviations in some physical constants (such as the precisely opposite

1. A juicier example follows, but for now this will do: M. B. Pour-El and J.L. Richards, 'The Wave Equation with Computable Initial Data Such that its Unique Solution is not Computable', Advances in Mathematics, vol. xxxix (1981), pp. 215-239.
2. H. T. Siegelmann, 'Computation Beyond the Turing Limit', Science, vol. cclxviii (1995). pp. 545-548. I argued the same case on theoretical grounds in a talk at the International Congress on Computer Systems and Applied Mathematics (St Petersburg, Russia 1993) and treat the topic more extensively in Mind Out of Matter: 'Topics in the Physical Foundations of Consciousness and Cognition (Kluwer Academic, forthcoming).
3. Although Dennett covers the topic only scantily (p. 448), there is no oddity about an algorithmic process evolving creatures with nonalgorithmic components-because natural selection is not an algorithm for producing nonalgorithmic components. Evolved creatures contain protons, too, but natural selection is not an algorithm for producing them.

charges of proton and electron) could make even simple structures like atoms impossible, it seems the proponent of a universal Darwinism has a case to answer. Dennett offers two replies; neither does the trick.

The first appeals to physicist Lee Smolin's[1] speculation about applying something like natural selection to the creation of entire universes. He suggests that the spacetime singularities of black holes are the birthplaces of new universes with randomly mutated physical constants; those universes with physical constants permitting greater production of black holes would naturally foster more offspring universes. Coincidentally, making baby universes and making babies have at least one thing in common: carbon. The element is essential both to the structure of simple organic molecules and to the life cycle of stars which survive the passage through supernova and white dwarf with enough mass left to undergo gravitational collapse; Smolin's hypothesis is that the physical constants of our own universe have evolved to permit a near maximum in black hole production, with carbon as a by-product.

The proposal is unsatisfying on two different levels. The first trouble is that natural selection demands not only differential production but also heritability, a likeness between physical laws at work in parent and offspring universes. Smolin offers no theory to account for heritability or even mutation itself; he merely explores what the ramifications might be if there were such a theory. He does note that "if the random changes are too large or too infrequent, then although natural selection will operate it is improbable that any structure significantly far from thermal equilibrium will emerge".[2] Later he warns, "I want to emphasize that the proposal made in this paper is extremely speculative ... it cannot be taken very seriously unless a detailed scenario and mechanism, based on known physics, can be developed to explain the value of a particular parameter which is falsifiable by some combination of experiment and theory".[3]

The second trouble with Dennett's appeal to Smolin concerns more directly the objection it is meant to defuse. Dennett notes (personal communication) that he needn't solve all the problems of Smolin's view to disarm the naive rhetorical question of 'what else could it be?' with respect to the involvement of intelligence in setting physical laws. But as for what it could have been, it is, after all, logically possible that the laws were set by a profoundly powerful but manifestly stupid little green Creature Who is a He or a She. The critic of Darwin's universality might endorse Smolin's proposal while still demanding an explanation of how it is that the rest of the physical laws-such as those governing the heritability and mutation of physical constants-came to be so conducive to the selection of constants themselves ultimately so amenable to life.[4]

1. L. Smolin, 'Did the Universe Evolve?', Classical and Quantum Quartly, vol. ix (1992), pp. 123-191.
2. L. Smolin, p. 177.
3. L. Smolin, p. 189.
4. Dennett does suggest (personal communication) that these 'background' laws might, following Malcbranche (see p. 184 of the text) be like eternal necessary truths of mathematics, requiring no further explanation.

Dennett's second reply to the standard objection echoes Hume's Philo in *Dialogues Concerning Natural Religion*. Run the creation and annihiliation of universes for eternity, he tells us, and we're bound to pass through each possible arrangement-including our own-not just once but an infinite number of times. Dennett extends the basic idea to natural laws themselves simply by supposing that physical constants change randomly between each death and birth of the cosmos.(1) The view, however, ignores a mathematical snarl, one twist of which is that the story amounts to iterative (countably infinite) universe testing in a space populated with continuous (uncountably infinite) possibilities. In other words, on the assumption that physical constants are described by real numbers (rather than, say, rational numbers, a restriction to which would itself presumably demand explanation), the infinity of universes which could be tested in this way is smaller than the infinity of all possible universes. It is just plain wrong that testing universes for an eternity would allow us to pass through all possible universes even once, let alone an infinite number of times.(2)

Despite concluding with musings about our "self-creating universe" (p. 185) and dubbing the entire topic of the chapter the "very foundation of Darwinian theory" (pp. 310-311), Dennett's applications of Darwin's dangerous idea to cosmology are overextended cranes which fail to support the "deflationary" replies he tries to build with them. Darwin's Universal Acid might well reach this far, but showing it will take more tools than the present book provides.

The second fly in Dennett's universal Darwinian ointment is a mathematically incoherent representation of design and possible gene sequences. Dennett's "Library of Mendel" (LoM), modelled on the Library of Babel of Jorge Luis Borges,(3) contains all possible nucleotide sequences (or, rather, descriptions of them) of length three billion (p. 112)-about that of the human genome. This set of genome descriptions, we are told (p. 143), is a model of Universal Design Space, the space in which the algorithms of Darwinian evolution trace their paths and in which the Tree of Life (ToL) unfolds (pp. 133, 143, 520). Dennett wants the LoM and Universal Design Space as a home lor all occurrences of design in Nature, from the human appendix to the Marche au Supplice of Berlioz and the latest methods for laying bricks; in his "Central Salvo", he maintains that 'there is only one Design Space, and everything actual in it is united with everything else" (p. 135, emphasis original). If there is any particular foundation on which Dennett intends his notion of universal Darwinism to rest, this is it. But the idea and the seductive vision which accompanies it are deeply flawed.

As presented in the book, the LoM is not equivalent to the geneticist's sequence hypercube (the point space of all genomes of a fixed length), and it

---

1. Dennett cites physicist John Archibald Wheeler for concurrence. Wheeler is a brilliant physicist; but in the context of his cosmological speculations it might be worthwhile recalling two other endorsements: the 'participatory universe' (the existence of which depends upon acts of observation) and the idea that a present measurement of a quantum system reaches back in time to make actual the path of history which made possible the observation itself.
2. Indeed, this holds even for universe re-creations with identical physical laws, since a particle's quantum mechanical position operator, for instance, has a continuous eigenspectrum.
3. 'The Library of Babel', in Labyrinths: Selected Stories and Other Writings (New Directories, 1962).

does not, contra Dennett (p. 161) preserve Hamming distance, defined as the number of loci at which two genomes differ. Of course, Dennett may concede this (personal communication) while simply re-ordering the volumes in the LoM in a new way, on different library 'floors' in billions of dimensions, so it does mimic sequence space for a genome of a given fixed length.

And if we were desperately fond of the Library of Mendel idea, we could massage it to accommodate genomes of length less than three billion-perhaps by adding a new nucleotide called 'blank' which could bring the shorter ones up to standard length-although this entirely ruins Hamming distances and the analogy with real sequence space.' Even still, the LoM does not, as Dennett claims, contain all possible genomes (p. 111) or even the smaller Library of Babel (p. 135), which Dennett renders as the set of all one million-character books. Dennett mistakenly supposes (p. 109) that both spaces contain an infinite number of volumes simply because we could represent any sequence of any length by re-using books as a basic alphabet. Remarkably enough, he concedes (p. 109) what is equivalent to the observation that if am alphabet is all we want, we can represent infinitely many books using the head and tail of a coin in my pocket. Yet the coin in my pocket doesn't contain all those books or descriptions of them; using heads and tails, it contains exactly two strings and a two-character alphabet with which the rest could be specilied. The LoM is no different: Dennett specifically populates it with descriptions of genomes" (p. 111, emphasis original), yet, owing to the noncompressibility of the vast majority of arbitrary strings,2 the number of such descriptions which can be stuffed into the LoM, the Library of Babel, or the coin in my pocket is strictly finite. While it is vast, the Library of Mendel contains only a finite subset of the infinite set of all finite-length books. No such set of books can serve as home to all occurrences of design (or, rather, to all descriptions of such occurrences, since comparatively few occurrences of design are actually books).

Problems build. Contrary to Dennett, the Tree of Life does not grow in the LoM. The ordinary ToL, as Dennett describes it in Chapter 4, is a dendrogram showing time on the vertical axis and arbitrary orderings on the horizontal axis. Searching for a sympathetic reading, we might write in a full book-length nucleotide sequence next to each organism in the dendrogram; we'd see everything liom multi-part chunks of RNA for viruses, if we included them, right up through to full-blown human genomes. But this dendrogram doesn't fit in the LoM-or the ordinary sequence hypercube, for that matter. The LoM has no time axis, it contains only genomes of fixed length. It is

1. Will the real genetic hyperspace please stand up? See pp. 92-100 of a book Dennett cites frequently: M. Eigen, Steps Towards Life (Oxford University Press, 1992).
 2. For complementary approaches to the relevant complexity and information theory, see Bennett, Chaitin, or Mulhauser: C. Bennett, 'How to Define Complexity in Physics, and Why', in W.H. Zurek (ed.), Complexity, Entropy, and the Pysics of Information (Addison-Wesley, 1990), pp. 137-148; G. Chaitin, 'Algorithmic Information Theory', IBM Journal of Research and Development, 21 (1977), pp. 350-359, 496; G. R. Mulhauser, 'To Simulate or Not to Simulate: A Problem of Minimising Functional Logical Depth', in F. Moran, A. Moreno, J.J. Merelo, P.Charon (eds.), Lectures Notes in Artificial Intelligence 929: Advances in Artificial Life (Springer-Verlag,1995), p.p. 530 514.

unclear whether Dennett intends his ToL to differ from the ordinary one or whether he means his rendition of sequence space to differ, but their marriage in the text is incoherent.

Dennett suggests (personal communication) that we can easily understand how the ToL fits in the Library of Mendel by just looking at the genotypes-perhaps in the way suggested above-but that getting to Design Space demands we look at phenotypes. This hints that Design Space is made of points just like sequence space, except that now the points are grown-ups. How to organise such a space or measure distance is an open question. (Put some ocelots on one floor, some jays on the next, long-legged birds through the door on the left. . .?) Elsewhere in the text (p. 125), Dennett seems to equate Design Space with the *fitness landscape*, constructed by giving each point in a given length genetic hypercube a number indicating an organism's relative fitness with respect to given selective conditions. This interpretation of Design Space, together with a tenuous equation of design with fitness allows some reasonableness to bubble up out of Dennett's often repeated image of evolutionary cranes doing "lifting in Design Space" so as to increase amounts of design. But the problem then becomes that a given fitness landscape cannot serve as a remotely 'universal' space in which to discuss the ToL, design, or fitness in general.

I have searched carefully for a sympathetic reading of the Library of Mendel, Tree of Life, and Universal Design Space, together with some relationship between them which matches what Dennett wants to say, while discharging the inconsistencies. The best I can manage is this: everything 'designed' is related by descent with some product of natural selection. 'Universal' design space is the merely ordered set of arbitrary-length descriptions of actual physical objects in spacetime, a set in which distance metrics are patently irrelevant both to the course of evolution and to changes in fitness. But this is unremarkable; it doesn't compare with Dennett's grand vision of the cranes of a giant algorithm relentlessly lifting the actual into progressively higher niches in a Universal Design Space. That grand vision, as it stands, is incoherent, and it remains to be seen whether any useful formal framework can support it.

Diagnosing a probable source of the quandary isn't too difficult. Dennett wants an account of Darwinian algorithms operating at several levels in what is, at a lower level, a dynamical system (a cosmos of particles evolving through four dimensional spacetime under the laws of physics). The trouble comes from trying to stuff a multi-level algorithmic description based on the Turing model of computation back into a dynamical systems framework. Dennett wants the 'trajectories' and 'lifting' characteristic of a dynamical system, yet neither point spaces with a Hamming metric nor-even worse-merely ordered sets make appropriate homes for dynamical systems. In itself, this isn't an insurmountable difficulty. After all, my Macintosh lurches and jerks discontinuously through the hypercube for its memory and CPU state in a thoroughly non-dynamical fashion; yet projecting activity from particular points in that hypercube into lower dimensions can turn up very nice dynamical patterns (such as the image formed on its screen when the computer

88

calculates trajectories of bouncing balls). Projecting actions of giant algorithms operating on a merely ordered set of genome descriptions into appropriate lower dimensional spaces might likewise turn up, for example, interesting patterns of population dynamics. These patterns are eminently useful, pictured in the right spaces-but not in anything resembling Dennett's Universal Design Space.

What difference does it make? The difficulty isn't just that Dennett's vision of grand algorithms in Universal Design Space is a helpful but replaceable little device for reaching some conclusion in one of many independent ways. The trouble is that this vision largely is the conclusion: Dennett wants to show that this is the right way to understand Darwin's dangerous idea and the world's storehouse of design. But the complications of the real world defy tidy and all-encompassing summary; Dennett offers a beautiful but imaginary vision. The alternative and much messier view endorses Dennett's contention that the influence of natural selection is inescapable while denying its complete monopoly over relevant explanation-it doesn't demand that every crane be a Darwinian one. The alternative view relies on cranes working at levels of description from particle physics all the way up, often under a Turing model of computation but not exclusively so.

In any case, as we noted at the outset, Dennett never pretended to offer a rigorous and highly technical framework. That the book fails to provide one thus isn't a fair criticism. So long as its few limitations are kept in mind, Darwin's Dangerous Idea makes an excellent introduction to the subject, and I sincerely hope it succeeds in its goal of awakening thinkers in far-flung fields to the rich philosophical terrain of modern Darwinism.

THE UNIVERSITY OF GLASGOW   GREGORY R. MULHAUSER

REPLY TO MULHAUSER

It is a pleasure to respond to such clear, constructive criticism. I will comment on five minor points raised by Mulhauser, and then turn to his major criticisms.

(1) I am grateful for the information about von Schilcher and Tennant's book, which unaccountably has not crossed my path before. I am not the only author to have overlooked it, and a cursory reading shows that it does anticipate points in many more recent books, including my own. And Mulhauser is right that Sampson and Montague would have been valuable allies had I developed my critique of Chomsky further-but that is a tangential issue in an overlong book.

(2) Do Skinnerian creatures "contemplate" their past errors? Only, as Mulhauser says, "on a very broad meaning", which is all that I intended: there must be some surviving trace of any past error so that the system has something relevant to adjust. The way I put it was indeed misleading, and I have taken steps to clarify the issue in my new book, Kinds of Minds (Basic Books and Weidenfeld and Nicholson, 1996), in which I also have more to say about the coevolution of memes and minds.

89

(3) My understanding of "algorithm" is not, I think, shifting but fixed throughout the book: the set of all Turing machines. Mulhauser suggests that this cannot be my meaning, since I say we can "treat any physical process at the abstract level as an algorithmic process". I don't see the conflict, since even if we grant for the sake of argument that there are physical processes the fine details of which are non-computable, this does not preclude their being treated at some more molar level as algorithmic, which is all that I was claiming. Mulhauser, however, says there is a more radical challenge arising from very recent work by Siegelmann and himself. I was wrong, he claims, to draw the crane/skyhook boundary at Turing computability, since there turn out to be chaotic neural networks with analogue components whose powers reach beyond Turing-computatibility and whose means cannot be "accurately" modelled by any Turing algorithm. I recommended that we consider the attractive prospect that the physical universe is Turing-comput-able, but I certainly didn't claim to offer a proof. Mulhauser offers grounds for any who wish to resist this assumption, even as a working hypothesis, and we shall see. If this new result stands up to scrutiny (of a sort beyond my own competence to conduct), and if it turns out moreover that these chaotic neural networks are a discernible part of our universe-the physical universe, not just the mathematician's universe of logical possibility, I will be obliged to move the boundary, as he notes. But recall all the hapless attempts to brandish Godel's Theorem as a disproof of materialism about the mind. Until we are given an argument that shows in detail why we must model some portion of the physical world as a chaotic neural net (on pain of manifest inability to predict and explain), we do well to drag our feet on the issue. Mathematical proofs about what is possible or impossible in these areas sometimes hinge on assumptions that make them less interesting and important than they first appear.

(4) I am way out of my depth when the topic turns to the computability of physics, and if I fall for the "common mistake" of assuming that all standard physics is computable, it is Roger Penrose who has led me astray, since I took the point from him. (It turns out that I also (mis-)took a point from Richard Feynman. I quoted with approval (p. 360) his rhapsodic commentary on potassium in the brain, but it is based on an elementary mistake, according to a recent letter to Nature. That's the risk I run when I rely on experts in other fields, as I often do.)

(5) I think Mulhauser misses my point about the logical problem with an algorithmic process spawning a non-algorithmic sub-process. Mulhauser sees that this is not an inescapable contradiction, since there is no telling what byproducts may happen to be found by-or even created in the course of operations of-any particular concrete instantiation of an algorithm. That was my point: any such non-algorithmic sub-process must be considered something external (even if it is a physical byproduct, somehow, of the larger process), found and then incorporated into the larger process. Otherwise it would be, in effect, a non-algorithmic subsystem of the supersystem, contra-dicting the premise that the supersystem was algorithmic.

Now to the major issues raised. Mulhauser is not persuaded by my Darwinian deflation of God the Lawgiver as the source of the order needed if evolution by natural selection is to unfold. He is right that Smolin's proposal does not (yet) include any account of heritability, and is, as Smolin himself acknowledges, highly speculative. And yes, the critic of Darwin's universality might indeed hold out for a further explanation of how the rest of the physical laws (the regularities which Smolin's speculations take for granted) came to have their apt values. All I claim to have done is to remove the aura of inescapability from the anti-Darwinian challenge. One is under no pressure to grant that there must be a Lawgiver, since it could all just be the product of random variation over aeons (or even eternity, in John Archibald Wheeler's version). Mulhauser's cautionary footnote about Wheeler's untrammelled imagination is well-taken, by the way. My own infrequent conversations with Wheeler have left me gasping at his blithe tolerance of theoretical extrava-ganve, not to say weirdness, but as I point out, there are no humdrum facts of cosmology.

   I in turn am   not persuaded by Mulhauser's objection based on the assumption of a mathematically continuous space of physical possibilities, but I am glad he raised it, since it brings into the open an issue I have long been puzzled by. Is nature really continuous? For instance, can we make sense of the idea that the value of some physical variable is actually an irrational number rather than any rational approximation thereof? You can't do physics without using equations that call for real numbers, but so far as I can see, this goes no distance towards establishing Mulhauser's strong thesis of continuity, with its implication of more than countably infinite possible universes to test. Even in the manifestly digital, quantised, and only countably infinite world of Conway's Lile, real numbers are needed to express such important constants as the 'speed of light'-the fastest velocity of propagation of change across the Life plane-which occurs along the diagonal and hence is faster by a factor of the square root of 2 than any change propagated from cell to cell along rows or columns at maximal speed (if we preserve the standard convention that Life cells are unit squares in a Euclidean space). This geometrical fact obviously carries no implications about continuity of the actual workings of the Life world. Does the mathematics used to describe the physics of the real world have a different status? I don't know. I am not sure what it would mean to say that we need either real numbers or noncomputable functions to describe the world. There may be arguments in theoretical physics that show why the continuities we observe in nature are Real continuities, but I have yet to be introduced to them. I will be very interested to consider any further instruction or argument Mulhauser can offer on this score.

   The larger problem he sees with my book is my "mathematically incoherent" representation of a single Design Space. He is right that there are different ways one might try to systematise the ensemble of possibilities I wished to discuss into a multi-dimensional space, and I did not settle on one way and stick to it. Now does this make my account "deeply flawed" or just annoyingly inconstant in its presentation? Some of the problems Mulhauser sees can be

brushed aside, I think. I don't claim, as he says, that the LoM contains all possible genomes; I say (in scare-quotes) that it contains "'all possible genomes'"-alluding by those scare-quotes to the point I had already made about the Library of Babel, which doesn't contain all possible books either (having none in Arabic or Chinese, for instance, but only translations of them), a negligible shortcoming for my purposes. His points about the difference between infinite and merely Vast spaces of possibilities are, so far as I can see, elaborations of points I made, not contradictions thereof. I would say the same thing about his comments on the "open question" of how to organise a space of phenotypes and how to incorporate a fitness dimension into the space; it is for the sorts of reasons he cites that I myself expressed doubts about developing the boundaries between different grades of possibility in any rigorous way (pp. 103-7). And he is right that given the various ways I invite my readers to visualise the space, my appeal to Hamming distance as an intuitive measure of difference is out of order, so I should drop it.

His effort to find a consistent interpretation of my remarks on Design Space leads him to an "unremarkable" destination which was indeed, pretty much all that I had in mind for the time being. We agree on the main point. As he puts it, "everything 'designed' is related by descent to some product of natural selection". Provided we understand 'descent' in all its varieties to an unmysterious natural relation, this is not a trivial claim. If one of the fruits on the Tree of Life is the beaver's dam, another is the Aswan Dam, and the processes that led to the design of each are related: the later, more sophisticated phenomena are both descendants of and composed of the very elements that account for the earlier, simpler phenomena.

I proposed Design Space as the single arena in which all this has happened. Mulhauser wishes my vision of Design Space was more mathematically precise and cogent, and so do 1. It would be wonderful to have a powerful scientific theory instead of just a philosophical image, but that is beyond me. He points to some of the difficulties that lie in the path of those who (I hope) will try their hands at this difficult set of problems. I welcome attempts like his to push what is impressionistic or metaphorical in my story into either clarity or oblivion. More of the former than the latter, if I'm lucky, but there is still a lot of conceptual work to be done.[1]

TUFTS UNIVERSITY   DANIEL DENNETT

1. I am indebted to Seymour Papert for illuminating discussion of Mulhauser's review and my reply.