Religious Studies (2020) **56**, 209–230 © Cambridge University Press 2018
doi:10.1017/S0034412518000392

Natural evil: the simulation solution

BARRY DAINTON

Department of Philosophy, University of Liverpool, Liverpool, L69 7WY, UK
e-mail: bdainton@liverpool.ac.uk

Thou shalt have no other gods before me.
(Exodus 20:3)

Abstract: **That natural evils are plentiful in our world is impossible (plausibly) to deny. Why would a benevolent and all-powerful God allow such an obviously imperfect world to be created? I argue that the problem of evil in this form only seems serious and intractable because our science and technology are currently at comparatively primitive levels. Civilizations whose technological capabilities are significantly more advanced than our own will be able to create artificial worlds of their own in a variety of different ways. Given the very real possibility of world-makers who are non-divine, the problem posed by natural evil is very much diminished.**

Virtual heavens and hells

If, as orthodox theists maintain, our universe was created by a Deity that is all-powerful, all-knowing, and maximally good, why is there *any* evil to be found in the world? Wouldn't an all-powerful and benevolent creator ensure that only the best of worlds are brought into being, and wouldn't the best of worlds be largely or wholly free from pain and suffering? In fact, the situation is not so straightforward. From the orthodox theistic perspective not all so-called evils are of the same kind. A great many of the death – and much of the destruction and suffering – to be found in our world are the direct consequence of freely chosen *human* actions. In creating us as subjects capable of making free choices, but permitting us to act as we freely decide, God created a world in which a good deal of evil was possible, though not necessary.

The truly intractable part of the problem of evil lies not with 'moral evils' such as murder, rape, torture, and genocide, but the so-called 'natural evils'. The suffering caused by natural disasters and diseases – in both humans and non-human animals – is obviously of immense proportions. Couldn't an all-powerful God

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

have prevented many of these innumerable natural evils? Couldn't a genuinely all-powerful benevolent being design a world where there is *far* less pointless suffering than in our world? It seems very plausible to think so.[1] Hence the fact that natural evils are so prevalent casts a dark shadow over the very existence of God as traditionally conceived. There have been many attempts to reconcile natural evils with a benevolent God over the centuries, but the problem of why they exist in such quantities has proved recalcitrant. In the eyes of many the dark shadow remains.[2]

The late Iain Banks's novel *Surface Detail* (2010) contains many of the features familiar to aficionados of his 'Culture' space operas. More relevantly to our current concerns, it also points to a solution to the problem of natural evil that is in many respects more attractive than any of the currently available alternatives.

Over the course of his novel Banks introduces us to various vast and powerful interstellar spacecraft, controlled by hyper-intelligent but often whimsical machine 'Minds'. Distributed through large tracts of the galaxy there are also the Culture's complement of human beings, with their biologically enhanced bodies and senses, equipped from birth with drug-glands which can be activated at will, and the ability to change their sex and appearance. Many of the Culture's human inhabitants choose to wear 'neural laces' which allow them to interface with machine minds, and which unobtrusively make detailed recordings of their brain states from moment to moment, allowing them to be reincarnated in the event of accidental bodily death. The principal innovation in *Surface Detail* is as memorable as it is disturbing.

In previous novels it had emerged that in the Culture universe 'virtual heavens' are commonplace. These are richly detailed virtual worlds, created and sustained by powerful artificial intelligences and carefully tailored to the preferences of individuals. The AIs responsible for these worlds do not merely generate the virtual scenery and virtual bodies the inhabitants enjoy, they also generate – via their computational activity – the minds and consciousness of the inhabitants themselves, who generally lack physical bodies of the non-virtual kind. People who opt for this path can spend their afterlives doing whatever takes their fancy, whether it be artistic creation, philosophical contemplation or 'boundless sex, adventure, sport, study, exploration, shopping [or] hunting' (Banks (2010), 127). In *Surface Detail* we make the disturbing discovery that in addition to virtual heavens, there are also *virtual hells*. These too are machine-generated virtual worlds, occupied by billions of fully conscious virtual inhabitants. But these worlds are far from heavenly. Their inhabitants are all condemned to lives of intense pain and sanity-threatening suffering, for vast expanses of subjective time.

Despite its many its imperfections our world is far from being the worst conceivable world, and – happily – Banks's hells are only fictional. However, so far as natural evils are concerned, Banks's virtual hells provide a useful clue as to their origins and causes in a divinely created cosmos. Perhaps the worlds where plentiful natural evils can be found are, without exception, *virtual worlds*, i.e.

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

sub-universes within the broader cosmos created not by God (directly), but by the inhabitants – human or otherwise – of technologically advanced societies. The advanced technology in question enables its users to create and sustain fully conscious beings, and control – to some extent at least – the character of the experiences of these beings. Since our world clearly contains very considerable quantities of natural evils, the natural conclusion to draw for the theistically inclined is that *this* world is in fact a virtual world, and that its imperfections are a direct consequence of the choices made by the beings who established the operating parameters of the simulation we find ourselves in.

Dealing with the problem of natural evils by (in effect) reducing natural evils to moral evils is a familiar one in theistic circles. The problem hitherto has been finding a remotely plausible explanation for *how* natural evils derive from the freely chosen actions of beings other than God. In an attempt to do just this Plantinga draws our attention to St Augustine's suggestion that the natural evils which cannot be attributed to God's punishment are the result of the actions of 'significantly free but non-human persons' in the form of Satan and his cohorts. (Plantinga (1974), 58)[3] In a similar vein, C. S. Lewis also appeals to Satanic corruption, suggesting that it is

> a reasonable supposition, that some mighty created power had already been at work for ill on the material universe, or the solar system, or at least, the planet Earth . . . This hypothesis is not introduced as a general 'explanation of evil': it only gives a wider application to the principle that evil comes from the abuse of freedom. (Lewis (1944), 77)

In response, Tooley dryly remarks that 'though it is possible that earthquakes, hurricanes, cancer, and the predation of animals are all caused by malevolent supernatural beings, the probability that this is so is extremely low' (Tooley (2015), 7.2). Tooley's doubts are surely warranted. Since the relevant scriptures tell us comparatively little about Satan (or his demonic cohorts), many contemporary Christians are reluctant to fall back on Satanic explanations of anything, let alone the totality of natural evils. More importantly, the progress of the natural sciences over the past three centuries has been such that there is very little room left for supernatural interventions in the natural order. In St Augustine's day, the notion that earthquakes or hurricanes – and perhaps some diseases – were due to the interventions of supernatural agents was by no means entirely implausible, given the absence of compelling alternative explanations. But these days the situation is altogether different. We know that earthquakes are the products of tectonic movements in the Earth's crust, and hurricanes are the natural by-products of the circulating gases and oceanic water currents which constitute the Earth's climate systems. Bolts of lightning (terrifying though they can be) are electrical discharges, not manifestations of the displeasure of demi-gods. Diseases too are entirely natural affairs: with infections being caused by micro-organisms, cancers by malfunctions in genetic transcription, and so forth. There is no need to go on: without obvious exceptions, it

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

looks very much as though all the natural evils we encounter in this world derive solely from the basic laws governing the material contents of our universe.[4] For better or worse, there is little or no room left for widespread demonic interventions.

In contrast, the *simulation solution* (as I will call it) to the problem of natural evil makes no appeal whatsoever to anything supernatural. All we are required to believe is that a technologically advanced future society (a) has the ability to create virtual worlds containing fully sentient beings, and (b) it decides to make use of this ability. The question 'Why is there so much suffering and disease in our world?' has the same answer as the question 'Why is there so much violence in the *Grand Theft Auto* series of video games?' In both cases the fault lies with the designers of the virtual world in question. In selecting the world-generation parameters as they did, the programmers of our world ensured that our world is as plagued by natural evils as it is.

Since the parameters in question ensure that natural evils are commonplace in our world, these parameters would not have been selected by the God of orthodox theism. But the Earthly descendants of today's video-game programmers will, I take it, be neither all-powerful, nor all-knowing, nor maximally good. Consequently, in designing *their* simulations they will not be constrained in the way God was, and the (virtual) worlds they create may very well contain plentiful natural evils – just as many gratuitous evils as we find in our world.

This solution to the problem of natural evils was advanced in a survey of the implications of Nick Bostrom's simulation argument in Dainton (2002), but I did not develop it any further. More recently it has been discussed in more detail by Johnson (2011). Johnson concludes that while the simulation solution indeed provides the theist with the best solution to the problem posed by natural evil, it also leaves the theist in a very unpalatable position, for if we adopt the solution, we are acknowledging that our universe was most likely designed and created by computer programmers rather than God. I agree that many theists will not find this prospect a very agreeable one, at least initially, but whether or not they can find it ultimately acceptable will depend on a variety of factors, not least how serious a problem natural evil poses for them, and the availability of alternative solutions they deem to be plausible.[5]

My focus here will be on the viability of the simulation solution itself. I suspect that many theists will be inclined to dismiss it out of hand, for a variety of different reasons. I will argue that many of these reasons are not good reasons, and so conclude that the solution should be regarded as a serious contender, irrespective of how surprising or unpalatable some may initially find it.

Minds and machines

Although the simulation solution evidently has merits when compared with other responses to the problem of natural evils, it also has some potential

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

weaknesses, and so can be challenged on several fronts. Let's start with a response along these lines:

> You are overlooking one important thing. For the simulation solution to be viable, it must be possible for computer-like machines to produce and sustain conscious minds. But while science fiction writers might take this hypothesis seriously, the rest of us needn't and shouldn't. Since there's no reason to think it is possible to artificially create *one* conscious mind in this fashion, let alone millions or billions, there's no reason to take the simulation solution seriously.

While many religious believers may find this a very natural line to take, simply to dismiss the simulation solution solely on this basis would be wrong. To make matters a little more precise, the envisaged objection is targeting the following type of virtual reality:

> *Type-1 Simulations*: Ordinary human consciousness – the sort of experience you and I are currently enjoying – is not essentially dependent upon human brains. Advances in technology will make it possible for programmable machines to produce streams of consciousness and precisely control their character. These simulation devices could, if they were appropriately programmed, generate streams of consciousness of the kind you and I are currently enjoying. A sufficiently powerful simulation system could sustain very large numbers of virtual subjects, in a shared virtual world. In some Type-1 simulations the conscious subjects act with the same freedoms as are enjoyed by non-virtual subjects; in others their choices are determined by the program, but this is not obvious – the subjects in question have the impression of being able to make entirely free choices, just as we do.

In assessing the challenge to the possibility of Type-1 simulations several issues come into play, some more fundamental than others. One obvious issue concerns computational power. The computational complexity of the envisaged machines far surpasses anything that currently exists. Are such devices really possible? The second – arguably deeper – issue is whether any digital computer, irrespective of its power, could generate consciousness.

Let's start by considering the complexity issue. As anyone who encountered recent generations of video games will vividly appreciate, computers already have the ability to sustain 3-D virtual environments that are highly impressive in both their graphical realism and sheer scale. Even so, containing as it does around a hundred billion neurons, each of which can be connected to thousands (or tens of thousands) of other neurons, the complexity of the brain is undeniably immense. No doubt many of us will have encountered the claim that the human brain is *the* most complex structure in the known universe. If so, it may seem implausible to suppose that this mind-numbing complexity – and hence the workings of an average human brain – could be replicated in the silicon circuits of a computer, no matter how advanced. In fact, if current estimations of a human brain's computational capacities are accurate (i.e. 10^{14}–10^{17} processing operations per second), then some of today's supercomputers already have a similar computational power, and the next generation of *exaflop* machines, predicted to arrive in 2020, will be capable of 10^{18} operations per second – more than enough for brain

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

simulation. Since further dramatic advances in computer technology are envisaged, it looks like brute computational power is not a serious obstacle to Type-1 simulations.[6] It is of course true that we are still a long way from knowing how to program a computer so as to enable it to reproduce the information processing that occurs in our brains. But it may well turn out that this difficulty is also far from insuperable. After some years in the doldrums, over the past decade or so the machine learning field has seen some remarkable advances, and many experts consider it likely that computer-based intelligence will rival (or surpass) human intelligence before very long – see Domingo (2015) for a survey of several promising approaches currently being explored.[7]

Hardware and software aside, there is of course a deeper problem. Even if future computers can replicate or exceed the typical human brain when it comes to information processing, there nonetheless remains a crucial question which has to be addressed: are computers the right kind of thing to sustain *conscious minds* in existence? Or is there *more to* consciousness than the computational processes of the sort digital computers can undertake? As is all too familiar, contemporary philosophers of mind remain hopelessly divided on this question.

Defenders of the possibility of machine consciousness are easy enough to find. If orthodox reductive functionalism is true, all aspects of mentality – consciousness included – could in principle be replicated on a non-biological substrate, provided it has the right causal organization. According to the non-reductive functionalism explored in Chalmers (1996), conscious states are distinct from and irreducible to physical states, but nonetheless nomologically correlated with patterns of information processing. On this view, a computer which perfectly replicated the information processing that has taken place in your brain over the past half hour would have sustained a stream of consciousness that is subjectively indistinguishable from your own. Tononi (2008) has recently been developing this sort of approach in a more systematic direction. However, there are also still-influential positions on the nature of the mental and the physical which are implacably opposed to any possibility of computer-sustained consciousness. Cartesian-style substance dualism is perhaps the most obvious of these views, but there are others. Searle's version of materialism closely links human consciousness to the biological processes occurring in our brains, and hence Searle is resolutely opposed to the possibility that our consciousness could be duplicated digitally. If we turn to the writings of contemporary philosophers of religion we find a similar diversity of opinion. Different forms of substance dualism are very much alive: see Foster (1991), Plantinga (2006) or Zimmerman (2004). Not only is the Thomistic tradition by no means dead – see Feser (2006, 2014) and Madden (2013) – there have been functionalist interpretations of the Aristotelian doctrine of forms, see Putnam (1975), Nussbaum (1978), and Goldberg (2004). There is certainly no shortage of contemporary Christian materialists – e.g. Rudder Baker (1995, 2007, 2011), van Inwagen (1995), or Merricks (2001).

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

Confronted by this continuing controversy a natural first thought is that the situation with regard to the simulation solution is straightforward.[8] If you subscribe firmly and confidently to a view of the mind–body relationship which *permits* Type-1 simulations, then for you at least, accepting the simulation solution in this guise is a genuine option. In contrast, if you are completely convinced that computer consciousness is *not* a genuine possibility, you will obviously be equally convinced that the simulation solution is not available. If you are undecided on the mind-machine issue, you must adopt an agnostic stance with regard to the simulation solution: it *may* turn out to be a viable response to the problem of natural evils, but equally, it may not. In fact, the situation is not so clear-cut, and for two very different reasons.

First of all, beliefs do not exist in isolation, they are systemic, and beliefs about which one is absolutely certain can exert a particularly powerful influence on one's overall belief system. Let's suppose you are initially sceptical about the claim that technology could advance to the levels required for Type-1 simulations. Let's also suppose that you are certain that God both exists and is maximally good. Given the prevalence of natural evils in the world, your system of beliefs contains a serious conflict. You have absolute confidence that God exists, but how can this be reconciled with the existence of natural evils? You now become acquainted with the simulation solution for the first time, and you find it provides a far better explanation for natural evil than any of the alternatives – or at least it would do, if it weren't for your doubts concerning the possibility of Type-1 simulations. Given the tension in your belief system, it looks very much as though the time has come to reassess these doubts. Revising your previous stance is perfectly reasonable because the *evidential situation* has significantly altered. You now have reasons for taking the simulation solution seriously, and hence accepting that you are yourself currently living in a virtual reality. Since this is only possible if simulation technology is possible, you also have a reason – a powerful one – for concluding that your previous scepticism regarding the latter may well have been misplaced. The fact that the possibility or otherwise of computer-based mentality remains an open issue, in both philosophical and theological traditions, means that revising your position on this issue may well prove to be possible.

The second complicating factor has a quite different source: there is more than one way to create a virtual world.

Neural interfaces: a second route to the virtual

In 2012 Jan Scheuermann, a quadriplegic woman, made headlines for using a neural interface – in the form of two probes on the motor cortex of her brain – to control a prosthetic hand. At the first annual 'Future of War' in February 2015, the director of the US's Defence Advanced Research Agency (DARPA) announced a remarkable advance: in more advanced testing of her interface, Scheuermann had successfully controlled an F-35 in the Joint Strike Fighter simulator by

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

employing thought alone: she just *willed* the plane to move in certain ways, and it did.[9] It's not just the paralysed who are benefiting: neural interfaces are also coming to the aid of people with defective senses. Primitive bionic eyes, which interface with the optic nerve, are already available. Cochlear implants have long made it possible for many deaf people to hear; these work by sending signals along the auditory nerve running from the ears to the brain, mimicking the action of normally functioning 'hair cells' in the cochlea.

It is clear that research on increasingly sophisticated mind–computer interfaces is going to continue, and continue to be very well-funded – in 2017 Facebook announced a project aimed at developing neural interfaces which would allow us to type a sentence just by *thinking* it, and Elon Musk launched 'Neuralink', a brain–computer interface venture whose ultimate goal is to allow us partially to merge with artificial intelligences. Given the successes achieved so far, there is every reason to believe that the neural interfaces available in the future will permit *very* sophisticated interactions between computers and brains.

It could very easily be the case (for example) that a century or so from now, as children grow up, they will do so with devices akin to Banks's neural laces growing deep within their brains, and these devices will allow our descendants to interact with computers in all manner of ways.[10] Forgetting would be a thing of the past – all our memories (or those we selected) would be backed-up on computer for instant recall. There would be no need for mobile phones to keep in touch: provided you are all connected to the same network, you could communicate with your friends directly, brain to brain. By connecting to, and taking full control of the brain's sensory centres these same interfaces would also very likely make full-scale virtual realities possible. If, say, you wanted to experience what it is like to climb Mount Everest, the interface wouldn't just provide you with a 3D visual simulation, it would provide you with fully real-seeming auditory and bodily experiences, based on recordings made by the neural lace of an actual mountaineer. These recordings might even include the mountaineer's own thoughts, memories, and mental imagery – these too might be accurately recorded via neural lace. The entertainments of tomorrow will thus have a degree of immersion that goes way beyond anything available today.

If advanced neural interface technology were readily available teaching and learning methods would also be transformed. Everything about your experience strongly suggests that you are an ordinary inhabitant of twenty-first-century Earth, leading a life that is typical of this period. In reality, you are a child in the early twenty-second century enjoying what then passes for a history lesson: you are learning what it was like for an ordinary person to be alive in the early twenty-first century by spending a few hours *as* a twenty-first-century ordinary person. The environment you perceive, your (apparent) memories, your hopes and beliefs, the language in which you are thinking and speaking: these are all entirely early twenty-first century in character. Needless to say, for the duration of this trip into the past you will have no memory of your twenty-second-

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

century self at all, but afterwards you are able to remember the trip vividly. Virtual reality history lessons like these have been routine for decades; they are made possible by neural implants which make it easy for computers to control people's mental states in fine-grained ways.

Evidently, we have a second potential source of virtual realities to bring into play:

> *Type-2 Simulations*: Advances in technology will make it possible to create completely life-like virtual realities by connecting human brains to suitably programmed machines which have the ability to control every aspect of a subject's conscious mentality. In some Type-2 simulations subjects are free to choose how to act, in others their choices are determined by the machine's program, but they nonetheless *seem* perfectly free.

The fact that Type-2 simulations may well prove possible is of not inconsiderable important to the viability of the simulation solution. Since the experiences of subjects inhabiting these virtual environments are generated by subjects' biological brains – albeit with the aid of neural interfaces – we are no longer being required to accept that computers produce consciousness *all by themselves*. Even if computers cannot themselves generate consciousness, the world you currently find yourself in could be a virtual one. So the simulation solution is secure, irrespective of the precise relationship between computers and consciousness.

Accepting the possibility of Type-2 virtual realities has a further implication. It may well make the simulation solution a live option for those who remain committed to Cartesian-style substance dualism. Contemporary Cartesians accept that our conscious mental lives are heavily dependent on the neural processes in our brains – this is hard to deny given what we have learned since Descartes's day about all the ways changes in brains can impact upon our mental lives.[11] Our perceptual states may be non-physical in nature, but their content is determined by the goings-on in our brains. It is thus far from inconceivable that advanced neural interfaces would allow a computer to override normal neural inputs to the regions of the brain responsible for sensory input, and induce – in effect – a full-scale virtual reality *within the soul*. The resulting experiences would be immaterial states within the immaterial substance, but caused and controlled by the sophisticated VR-device.[12]

A further (and quite different) route to the virtual

Still other forms of simulation are possible, and it would be a mistake to overlook them. Although many simulations are carried out by computer programs these days, the use of *small-scale physical models* has by no means disappeared. Scale models of new aircraft, for example, are tested in wind-tunnels before going into production. Models of this kind are also very relevant to the viability of the simulation solution.

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

In the *Matrix* movies the vast bulk of the human race is hooked into a communal Type-2 virtual reality, and their computer-supplied experiences do not correspond at all closely to with the real world. In *The Truman Show*, the eponymous hero is an insurance salesman in the small town of Seahaven. The unfortunate Truman does not realize that he is the star of a 24-hour reality TV show, and that the Seahaven he knows is nothing more than a fully realistic TV set, enclosed in a large dome – disguised as a sky – which separates him off from the outside world. He is entirely oblivious to the fact that all the people he has ever known – his family included – are all actors. So far as Truman is concerned, the world he knows *is* the real world; but he's deluded: Truman's world is a deliberately designed and carefully controlled artificial construction that is distinct from the real world. Earlier I defined *virtual worlds* as 'sub-universes within the broader cosmos created not by God (directly), but by the inhabitants – human or otherwise – of technologically advanced societies'. Since Truman-type worlds qualify as virtual realities, at least on this definition, we have a further route to total simulations to consider:

> *Type-3 Simulations*: Advances in technology will make it possible to design and create artificial physical sub-worlds that seem to their inhabitants to have all the appearances – and complexity – of our world. Some Type-3 simulations are autonomous: after their initial creation their designers refrain from interfering at all, irrespective of how things unfold. In other Type-3 simulations the world-designers *do* intervene – more in some worlds, less in others – but never in a way which is evident to their inhabitants.

The *Truman* scenario has some distinctive features. Only Truman is unaware of his predicament and the real nature of his environment; all the other inhabitants of Seahaven are fully aware of the facts. This need not be the case more generally. It seems quite conceivable that the entire population of an artificial world could be in Truman's position, and be entirely unaware that things are not all what they seem. As readers of *The Hitchhiker's Guide to the Galaxy* will recall, in that fiction the entire Earth is an artificial construct built by a technologically advanced species of aliens – the Earth is in fact a large computer running a program – but the entire population is entirely oblivious to this. In both these scenarios there are strong resemblances between the artificial construct and the real world, e.g. both contain human beings. But this need not be the case: some artificial worlds might be very different from the world inhabited by its creators, even if this fact is completely concealed from their inhabitants.

The artificial planets in both the *Truman* and *Hitchhiker* scenarios are full-scale models: the objects they contain are all the same size as their real-world counterparts. Once again, there is no reason to think all Type-3 simulations will be like this: many artificial worlds might instead involve radical miniaturization. Several quite dramatic examples of this can be found in science fiction. In *Superman*, for example, the bottle city Kandor has a population of some six million

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

people.[13] The first *Men in Black* movie features a more dramatic example: an entire galaxy has been shrunk to fit inside an object no larger than a jewel.

Are Type-3 simulations of this sort really possible? I can see no reason to rule them out. The fact that they are impossible with *our* current technologies means very little: beings who have attained a far (far) higher technological level might well be able to create artificial model-worlds which have the same appearance to their inhabitants as the Earth with ease. In the astronomer Kardashev's classificatory scheme, a *type III civilization* has the ability to control energies on the scale of entire galaxies; such a civilisation might well have little trouble creating a full-scale model of a populated planet, along the lines of *The Hitchhiker's Guide to the Galaxy*.[14]

Moreover, as the physicist John Barrow has pointed out, as well as achieving mastery of matter over large scales, a technologically advanced civilization can also be expected to be able to manipulate and control matter at very *small* scales: at atomic levels and sub-atomic levels. It is worth bearing in mind here that discoveries in cosmology and particle physics made in recent decades suggest that the human size-scale (as we might call it) is situated more or less at the centre of the cosmic size-scale. If a typical star is of the order of 10^{10} cm in size, the radii of typical atoms are around 10^{-10} cm; the Milky Way is estimated to be around 10^{21} cm wide, the radius of high energy neutrino is 10^{-21} cm; the approximate diameter of the visible universe is 10^{27} cm, the Planck length – currently thought to be the smallest scale at which physical phenomena can occur – is of the order of 10^{-33}.[15] Given that there is so much room for manoeuvre and manipulation in the realm of the very (very) small, it does not seem implausible to suppose that civilizations possessing very (very) advanced technology will possess the ability to manipulate matter and energy at those scales – an ability which would make it possible for them to achieve feats of miniaturization that to us would seem magical. For this reason virtual worlds akin (in size at least) to Kandor could easily be possible as well.[16]

It is not difficult to envisage some of the reasons why civilizations with these capabilities might choose to deploy them. As with many of their computer-generated counterparts, artificial worlds – whether full-scale or miniature – might well be created for research purposes. Quite how a complex system, particularly a system containing biological entities, will behave over long periods of time may well be impossible to predict, at least for less-than-omniscient beings.[17] The only way for a non-divine world-maker to discover how a given design for a planet or solar system will perform is by building the world (in either computational or material form) and watching and waiting to see how things unfold. Equally, world-creation might be one of the principal means by which members of a highly advanced civilization amuse themselves. It seems perfectly conceivable that many of the artworks of such civilizations could take this form – see Steinhart (2010) for an interesting take on this. Since (non-divine) beings capable of world-creation are likely to have a good deal of spare time at their disposal, they might very well end up creating worlds in vast numbers.

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

With every discovery that we have made about the universe we inhabit (most recently the detection of planets in other solar systems) the Fermi paradox has grown every more baffling. In a universe that is so old and so big, with so many potential sites for life, why do we see no trace of intelligent life elsewhere? The solution, in part at least, may be that many intelligent civilizations have elected to devote their energies to investigating – and expanding into – the realm of the very small, rather than engaging in cosmic-scale engineering projects or interstellar travel.[18]

Is the solution self-undermining?

If the hypothesis that we inhabit a virtual world had turned out to be wholly absurd, the simulation solution would be untenable. But inevitably speculative as they were, the considerations advanced over the course of the last few sections suggest this hypothesis cannot simply be dismissed. Virtual realities of many different types and scales are frequently to be found in science fiction, but it would be a mistake to think they will forever be confined to fiction.

However, the simulation solution is not yet out of the woods; it must confront a further challenge, of a very different sort.

> Even if virtual realities really are possible there's nonetheless a problem with what you're proposing. We are attempting to reconcile the existence of natural evils with the existence of God. If we do this by accepting that our reality – or what hitherto we have taken to be such – is no more than a virtual world produced with the aid of some sort of advanced technology, aren't we also removing any grounds we might have had for supposing that God exists in the first place? If this world is a virtual one, it's really nothing but a form of fiction or fantasy. In which case, we no longer have any reason to believe the claims of any religion are true. If the simulation solution renders belief in God untenable, it can scarcely be considered to be a satisfactory solution to the problem of natural evil. For the latter is precisely the problem of understanding how *God* could have created a world containing so much pointless pain and suffering.

The problem is potentially fatal. If the simulation solution really does undermine the reasons one might have for believing the world *was* created by an omniscient, omnipotent, and omnibenevolent God, it would, in effect, be self-undermining. But once again, when examined more closely, this problem turns out to be less serious than it initially appears. To a surprisingly large extent, the most important reasons one might have for believing in God do not require us also to believe we are living in the ground-floor of reality, rather than a higher-level simulation of one kind or another.

We can begin by considering what might initially strike one as the most problematic case: religious scriptures. Is accepting the simulation solution compatible with taking any founding religious text – whether it be the Old Testament, the New Testament, the Koran, the Avesta, or the Book of Mormon – as veridical? Since (to take just one example) many Christians would find a solution to the problem

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

of evil less than satisfactory if it obliged them to regard Jesus Christ as being a fictional character on a par with Sherlock Holmes or Harry Potter, the simulation solution might seem to be in serious trouble. Does it not oblige us to do precisely this?

To assume we are under any such obligation would simply be a mistake. Simulations *can* be misleading – there is certainly no guarantee that they reflect the characteristics of ground-level reality – but equally, they do not *have* to be misleading: simulations which *don't* misrepresent reality to any significant degree are perfectly possible too. Consequently, theists who find the simulation solution very appealing and who also find themselves firmly convinced that the Bible (say) is true in all essentials, have a way of rendering their beliefs coherent. They need simply take themselves to be living in a world which is (a) simulated, but (b) closely resembles the real world in the relevant respects. The relevant inter-world resemblances in this context – to a first approximation – will very probably include broadly similar laws of nature, social histories, and technological trends and developments. They will certainly be such that the Bible, and many of the events portrayed in it, are not confined to the virtual, but rather have counterparts in the ground-level reality as well. The only parts of scriptures which *cannot* accurately represent base-level reality are those involving natural evils of the sort which will not exist in the non-virtual world – assuming the simulation solution is correct.

In this connection it is worth noting that a position broadly along these lines has long existed in the Christian tradition, or at least in one of its prominent offshoots (deemed heretical by some). Many variants of Gnosticism flourished in the first and second centuries, but most responded to the problem of evil in the same manner: by holding that our world wasn't created by God but by some lesser being(s). To focus on just one example, in *Against Heresies* (Book 1, chapter 24) Irenaeus relates that in the second century Basilides held that our world is not in fact the direct creation of the all-powerful true God. Rather, a group of secondary deities emanated from God, and these were jointly responsible for creating the first (and highest) heaven. A further order of still more lowly sub-deities emanated from the first group, and they created their own second-order heaven – and so down through 365 levels. Since, according to Basilides, our world is the very lowest of these heavens it is scarcely surprising that it is so imperfect. Seeing our plight, the true God took pity on us and sent Christ down into our world to 'bestow deliverance on them that believe in him, from the power of those who made the world. . . . those who know the truth of this history will know themselves free of the power of the princes who built this world' (Irenaeus (1885), 1, 23). As Borges observes, what Basilides supplies us with – amongst other things – 'is a quiet resolution of the problem of evil by means of a hypothetical insertion of a gradual series of divinities between . . . God and reality' (Borges (1999), 66).

Far more recently, and operating in a very different tradition, Bostrom has suggested that our world might not merely be virtual, it could be generated by purely computational activity on a computer that is itself part of a virtual world: 'virtual

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

machines can be stacked: it is possible to simulate one machine simulating another machine, and so on, in arbitrarily many steps of iteration' (Bostrom (2003), 253). While there are clear parallels between Bostrom's hierarchy of nested virtual worlds and the nested hierarchy of virtual heavens in Basilides's baroque cosmology, there are also significant differences. Bostrom's virtual worlds are not generated by supernatural beings or forces; they are the products of ordinary physical machines, operating in accord with natural physical laws. Moreover, since these machines were originally created by ordinary humans rather than any kind of divine being, the scenario is fully compatible with orthodox monotheism.

Arguments for the existence of God

How would embracing the simulation solution impact upon theological arguments for the existence of God? Needless to say, there are a great many such arguments, many of which can be formulated in different ways, and there is much disagreement among theists as to their merits, so a comprehensive survey of this complex terrain cannot be attempted here. But a brief overview will illustrate that the force of these arguments is largely independent of supposing that we inhabit base-level reality.

Consider first the ontological argument. In the eleventh century Anselm argued along these lines. Whatever other properties God possesses, we know that *no being greater than God* can be conceived. Now suppose God doesn't exist. We can easily conceive of another being, one which is such that *no greater can be conceived, and which also exists*. But this is absurd: such a being would be greater than God, and nothing can be greater than God, who we have already agreed is the greatest conceivable being. Since the attempt to deny God's existence leads to absurdity, we must conclude that God exists. Centuries later Descartes and Leibniz offered their own alternative versions of the ontological argument, and in the twentieth century logicians and philosophers such as Gödel, Malcolm, and Plantinga have gone on to offer further variants. There is no need to enter this complex debate here. Based as it is on *a priori* metaphysical reasoning that holds good (if it does) in all logically possible worlds, the validity and persuasiveness of the ontological argument are entirely independent of whether or not we are living in a virtual world.

The same goes for the cosmological arguments. In response to the question 'Why is there something at all, rather than just nothing?' some are prepared to accept the answer 'No reason at all, the universe exists, and that's all there is to it.' But there are others who find this unacceptable, and plenty who find it most plausible to conclude that the chain of explanation must terminate in a being which by its nature couldn't fail to exist, i.e. God. If you find one or more versions of the cosmological argument compelling, there is no reason at all why accepting the simulation solution should lead you to revise your position. Accepting this

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

solution does entail that God was not *directly* responsible for either the existence or the character of our world – the simulation creators are to blame for that. But God still enters the picture as the direct creator of the base-level of reality.

Some find the existence of morality is powerful evidence for God. There are acts – the murder of a child, the torturing of a blameless adult – that are just obviously wrong. Moreover, in the eyes of many, such acts are *objectively* wrong. Since God is the only possible source of objective morality, we have no option but to conclude that God does exist. Or so one version of the argument runs. If you find this argument persuasive, accepting the simulation solution gives you little reason to reject it. Morality is not confined to the ground-floor level of reality; inflicting agonizing pain on innocent people is wrong, irrespective of whether the people concerned are constituted of real or simulated flesh and blood – a point Banks's fictional cyber-hells illustrate very well. Consequently, if you believe morality is both real and objective, and you are firmly of the view that objective morality must be grounded in a divine being such as God, then you are justified in concluding that God exists on this basis even if you also believe that you inhabit a virtual world.

It is difficult to see that most 'pragmatic' arguments are undermined when translated into virtual environments. If you currently think Pascal's wager provides you with grounds for believing in God, coming to believe you inhabit a virtual world does not alter the situation at all: you still have as much to gain by believing in God (or lose by not doing so).

How do arguments from design fare? In accepting the simulation solution we are of course accepting that our own world *does* have a designer – albeit one who is less than divine.[19] We also know that irrespective of whether we are in a simulation of Type-1, Type-2, or Type-3, the natural laws of the underlying reality have to be such as to make possible the kinds of advanced technologies required for simulations of entire worlds. Is this form of nomological organization in itself evidence of divine workmanship? At least some of those who were previously inclined to accept a version of the design argument might well think so. Moreover, anyone who has suffered back problems, or wondered why human childbirth is so painful, will have been vividly aware of the design deficiencies in the human body. Since God was not directly responsible for the design of our world, these deficiencies – along with all other design-related natural evils – can no longer be laid at his door.

Is adopting the simulation solution compatible with taking any form of *religious experience* as providing evidential support for God's existence? As far as I can see, it is – though given the variety of different forms of experience falling under this heading, one must tread carefully here. One of the relevant forms of experience takes the form of a quasi-perceptual encounter with God. In other forms of religious experience God is not himself present *in* the experience, but the experience has special qualities which nonetheless in some way *points to* the existence of the divine. Neither case poses any difficulty for the simulation solution. For in a

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

theistic context, even if we assume that the vast bulk of our experience is being caused by the simulation devices sustaining or controlling our consciousness, these devices are not the only possible causes of our experience. Since God exists (or so we are assuming), if he so chooses he can easily reveal his existence to us in our experience: either *directly*, by presenting some aspect of himself to the perceiving subject, or *indirectly*, e.g. by ensuring that some (if not all) are able to discern the hallmark of the divine in their worldly experience.

Bostrom turbocharged

Before moving on there is a further point to note. So far as Bostrom's simulation argument is concerned, when estimating how likely it is that you're living in a simulation there are three factors that are relevant: (i) how likely you think it is that humankind will survive long into the future and continue to develop sophisticated technology, and (ii) the quantity of simulations you think our descendants are likely to produce, and (iii) their character – how many simulations will reflect ordinary twenty-first century life, as opposed to twenty-second or twenty-third. These are the *only* considerations which matter. In the theological context the situation is very different.

Let's assume you are firmly convinced that God exists and worried (very) about the problem of natural evil. In the absence of any better solution to the latter, if you think simulations are possible at all, you will conclude that it's *almost certain* that you live in a simulation simply because of the level of natural evil our world contains. You will have powerful reasons for concluding that this reality is virtual even if you would otherwise think it is fairly *un*likely our descendants will have the ability or desire to produce world-simulations of the twenty-first century in vast quantities. The fact that this world is so imperfect – combined with your religious beliefs – makes it very likely that you are living in a virtual world irrespective of most other considerations. In effect, the simulation argument is *turbocharged* when transplanted into a theological context.

A ban on world-making?

To bring matters to a close, and also link up with where we started from, I want to consider one last objection:

> Even if a benevolent and all-powerful God *could* create a world where the laws of nature are such that sophisticated world-simulation technology is possible, there is no way that such a God *would*. The consequences for the misuse of the technology are simply so great. These misuses range from making it possible to create cyber-hells of the sort Banks envisaged, along with systematically deceiving whole populations as to their real conditions. If we solve the problem of natural evil in the way you propose, we have to accept that we are all locked in something akin to a hallucination from which we cannot escape. A benevolent God would have good reasons for *not* allowing such things to happen.

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

This objection does not strike me as very plausible – though it does introduce issues which merit more discussion than I can devote to them here.

As recent history illustrates all too well, the laws of nature in our world – assuming for the moment that it is real rather than virtual – already permit technologies which can cause enormous quantities of human suffering.[20] It is not obvious that the horrors that would be possible in virtual realities will be significantly different in level or kind from those that are possible in the real world – particularly if we bear in mind the new horrors that future technological developments will make possible. It is undeniably true that if God created a world in which Type-1 virtual worlds are possible, then in doing so he made it possible for far greater numbers of conscious subjects to exist than would otherwise be the case. But the inhabitants of the vast majority of these worlds may well lead enjoyable and fulfilling lives that are very much worth living. Banks's cyber-hells may not be the norm among virtual realities (and hopefully aren't).

As for the deception problem, for the reasons outlined earlier, anyone whose religious beliefs derive largely from religious scriptures and who also finds the simulation argument an attractive solution to the problem of natural evil should believe that we aren't being greatly deceived about the basic character of our world – we are living in something akin to a work of ordinary fiction at worst, not a total fantasy. The fact that there are different types of simulation is also very relevant here. A short-lived solitary Type-2 simulation is very different from a Type-1 total simulation or a large-scale Type-3 simulation. The latter will typically involve millions of subjects interacting with one another in a common (virtual) environment; the solitary simulation involves just *one* subject, and the mere appearances of other people. There's certainly a case for saying that subjects in solitary simulations of this sort *are* victims of a deception. But even so, if subjects enter such simulations voluntarily – for purposes of entertainment or education, as might well be the case – there is no obvious harm in them.

Subjects in a Type-3 simulation possess physical bodies and inhabit a physical world, albeit a world (or sub-world) that has been designed and created by non-divine beings, and which may not resemble the real universe in the ways its inhabitants assume. The subjects inhabiting long-lasting communal Type-1 simulations are in a very different sort of situation, existing as they do in software-sustained simulations. However, here are plenty of reasons for saying these people are not deceived if they believe they live in a *world*, even if their world is not physical in nature in the way its inhabitants might suppose. First of all, by virtue of being *communal* such subjects are not alone: there are other subjects within the same virtual world with whom they can interact. Second, the virtual environment such subjects find themselves in is filled with publically observable virtual objects. The table across the room may not be a solid material body – in Type-1 simulations it will consist of nothing more than a collection dispositions to produce certain patterns of experience in the simulation's

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

inhabitants – but *that there is a table across the room* is something other subjects will agree on, if they happen to look in that direction.[21]

It's true, of course, that since a (virtual) world of this kind is dependent upon the continued functioning of a VR-device it is not a metaphysically basic entity. But since theists will anyway hold that God is the metaphysically fundamental entity, it is hard to see why they would allow this fact to disqualify a virtual world from being a *world*.[22]

Further issues

Taking the simulation solution seriously gives rise to a number of issues worthy of further investigation. The solution entails that ground-floor reality contains *far less* natural evil than our world, but this needn't mean that it contains *no* natural evils at all – perhaps some are necessary for 'soul-building' purposes. If so, what sort and how many? There is also the issue of how it is possible for a world to be similar to ours in many respects but still possess far less by way of natural evil. Could the laws of nature be precisely the same, and the desired effect be produced by judicious selection of initial conditions? Or will the laws have to be slightly different? Or will occasional (subtle) divine interventions be required? Assuming it is virtual, might our world – with all its failings – be the product of human world-makers who are conducting experiments with a view to answering this very question?

Irrespective of its ultimate fate, one thing the simulation solution does reveal is that the seriousness of the problem of evil is context dependent. Put yourself in the shoes of a theist living in technologically primitive times. The idea that there could be technologies that are capable of creating virtual worlds that are inhabited by sentient beings is unlikely to cross your mind at all, and if it does you aren't likely to take it at all seriously. Since the simulation solution is unavailable, the threat to theism posed by natural evils is very serious – you may even be tempted by supernatural explanations. The situation is entirely different if far higher levels of technology are prevalent. Let's suppose you find yourself living in a period of human (or post-human) history when simulation technology that is capable of generating completely real-seeming virtual realities is as readily available as television is today. Religious believers in *that* sort of world who also find themselves surrounded by natural evils will find it very natural to conclude – with little further ado – that their world must be a virtual one. This will be by far the most obvious and plausible explanation for their world's being as it is – by the most obvious way of reconciling their religious convictions with their empirical experience.

Intriguingly – but also instructively – the situation we find ourselves in lies between these two extremes. Simulation technology of the required potency is undeniably on the horizon, but it has not yet arrived. Consequently, the simulation solution will not seem as obvious or compelling for people living today as it will be

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

for those who find themselves living in technologically more advanced times. Expect this to change if and when simulation technology *does* become real.[23]

References

Ball, J. A. (1973) 'The zoo hypothesis', *Icarus*, **19**, 347–349.

Banks, I. M. (2010) *Surface Detail* (London: Orbit).

Barrow, J. D. (1998) *Impossibility: The Limits of Science and the Science of Limits* (Vintage: London).

Baxter, S. (2000) 'The planetarium hypothesis: a resolution of the Fermi paradox', *Journal of the British Interplanetary Society*, **54**, 2010–2016.

Borges, J. L. (1999) *Selected Fictions*, E. Wienberger (ed.) (New York: Viking).

Bostrom, N. (2003) 'Are you living in a computer simulation?', *Philosophical Quarterly*, **53**, 243–255.

Chalmers, D. (1996) *The Conscious Mind* (Oxford: Oxford University Press).

Chalmers, D. (2005) 'The matrix as metaphysics', in C. Grau (ed.) *Philosophers Explore the Matrix* (Oxford: Oxford University Press), 132–176.

Dainton, B. (2002) 'Innocence lost', <http://www.simulation-argument.com/dainton.pdf>.

Dainton, B. (2012) 'On singularities and simulations', *Journal of Consciousness Studies*, **19**, 42–85.

Domingo, P. (2015) *The Master Algorithm: How the Quest for the Ultimate Learning Machines will Remake Our World* (Philadelphia PA: Perseus).

Feser, E. (2006) *Philosophy of Mind* (Oxford: Oneworld).

Feser, E. (2014) *Scholastic Metaphysics: A Contemporary Introduction* (Neunkirchen-Seelscheid: Editiones Scholasticae).

Foster, J. (1991) *The Immaterial Self* (London: Routledge).

Irenaeus, St (1885) *Against Heresies*, A. Roberts & W. Rambaut (tr.), in *Ante-Nicene Fathers* Vol.1, A. Roberts, J. Donaldson, & A. Cleveland Coze (eds) (Buffalo NY: Christian Literature Co.).

Goldberg, N. (2004) 'Is Aristotle's philosophy of mind functionalist?', *Philosophical Writings*, **26**, 47–65.

Johnson, D. (2011) 'Natural evil and the simulation hypothesis', *Philo*, **14**.

Kardashev, N. S. (1964) 'Transmission of information by extraterrestrial civilizations', *Soviet Astronomy*, **8**, 217–221.

Lewis, C. S. (1944) *The Problem of Pain* (New York: Macmillan).

Lloyd, S. (2000) 'Ultimate physical limits to computation', *Nature* **406**, 1047–1054, <http://arxiv.org/pdf/quant-ph/9908043.pdf>.

Madden, J. (2013) *Mind, Matter and Nature: A Thomistic Proposal for the Philosophy of Mind* (Washington DC: Catholic University of America Press).

Maitzen, S. (2006) 'Divine hiddenness and the demographics of theism', *Religious Studies*, **42**, 177–191.

Meixner, U. (2017) 'Substance dualism', in Y. Nagasawa & B. Matheson (eds) *The Palgrave Book of the Afterlife* (London: Palgrave), 277–296.

Merricks, T. (2001) 'How to live forever without saving your soul: physicalism and immortality', in K. Corcoran (ed.) *Soul, Body, and Survival: Essays on the Metaphysics of Human Persons*, (Ithaca NY: Cornell University Press), 183–200.

Nussbaum, M. C. (1978) *Aristotle's De Motu Animalium* (Princeton NJ: Princeton University Press).

Perrine, T. & Wykstra, S. J. (2017) 'Skeptical theism', in C. Meister & P. K. Moser (eds) *The Cambridge Companion to the Problem of Evil* (Cambridge: Cambridge University Press), 85–107.

Plantinga, A. (1974) *God, Freedom and Evil* (Grand Rapids MI: Eerdmans).

Plantinga, A. (2006) 'Against materialism', *Faith and Philosophy*, **23**, 3–32.

Plantinga, A. (2017) 'Supralapsarianism, or "O Felix Culpa" ', in M. L. Peterson (ed.) *The Problem of Evil*, 2nd edn (Notre Dame IN: University of Notre Dame University Press, 2017), 363–389.

Putnam, H. (1975) 'Philosophy and our mental life', in *Philosophical Papers*, II (Cambridge: Cambridge University Press), 291–303.

Rudder Baker, L. (1995) 'Need a Christian be a mind-body dualist?', *Faith and Philosophy*, **12**, 489–504.

Rudder Baker, L. (2007) 'Persons and the metaphysics of resurrection', *Religious Studies*, **43**, 333–348.

Rudder Baker, L. (2011) 'Christian materialism in a scientific age', *International Journal for Philosophy of Religion*, **69**, 1–12.

Schellenberg, J. L. (1993) *Divine Hiddenness and Human Reason* (Ithaca NY: Cornell University Press).

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

Steinhart, E. (2010) 'Theological implications of the simulation argument', *Ars Disputandi*, **10**, 23–37.
Tononi, G. (2008) 'Consciousness as integrated information: a provisional manifesto', *Biological Bulletin*, **2015**, 216–242.
Tooley, M. (2015) 'The problem of evil', <http://plato.stanford.edu/entries/evil/>.
van Inwagen, P. (1995) 'Dualism and materialism', *Faith and Philosophy*, **12**, 475–488.
Webb, S. (2015) *If the Universe is Teeming with Aliens . . . Where is Everybody? Seventy-Five Solutions to the Fermi Paradox and the Problem of Extraterrestrial Life* (Dordrecht: Springer).
Wolfram, S. (1984) 'Computer software in science and mathematics', *Scientific American*, **254**, 188–203.
Wolfram, S. (2002) *A New Kind of Science* (Champaign IL: Wolfram Media).
Zimmerman, D. (2004) 'Christians should affirm mind–body dualism', in M. Peterson & R. Van Arragon (eds) *Contemporary Debates in Philosophy of Religion* (Oxford: Blackwell), 315–326.

Notes

1. To take just one example, a 2016 report by the International Agency for Research on Cancer estimates that there are currently around 80,000 deaths per year caused by childhood cancers. Although cancer is prevalent in many species of animals, in some it is rare, e.g. elephants, bowhead whales and naked mole rats. Some of the underlying genetic causes for these striking divergences have been discovered, e.g. whereas humans have just one copy of the tumour-suppressing TP53 gene, elephants have twenty. Was an all-powerful God incapable of arranging matters so that human genomes are equipped with the same cancer-protective powers as the genomes of elephants or bowhead whales?
2. When a seven-year-old Japanese girl who survived the 2011 Tohoku tsunami asked Pope Benedict XVI why so many children suffer in natural disasters the Pope admitted: 'I also have the same questions: why is it this way? Why do you have to suffer so much while others live in ease? And we do not have the answers, but we know that Jesus suffered as you do . . .'.
3. In *God, Freedom and Evil* Plantinga puts forward the Satanic explanation as merely a logical possibility, but in later writings he evidently takes it rather more seriously: 'The thought that much evil is due to Satan and his cohorts is of course entirely consistent with God's being omnipotent, omniscient and perfectly good; furthermore it isn't nearly as improbable with respect to "what we now know" as most philosophers tend to assume' (Plantinga (2017), 377)
4. When asked whether he was ever tempted to see the presence of a divine creator behind the marvels of nature, the naturalist David Attenborough replied:

> When I expose people to such marvels, such majesty and such details of animal life, I also have in mind at the same time, the horror of that small boy in the Congo region of Africa, whose eye is at that same moment slowly penetrated by a river worm from behind, where his brain is, after it has wormed its way all through his body from his feet upwards. I have to keep both glory and horror in mind. When I do that, I do not see a benign creator in all this. (BBC TV interview)

   While there is no denying that the advance of science has led to the discovery of many wonders, it has also led to the discovery of entirely new threats and horrors.
5. The core thesis common to the various forms of so-called 'sceptical theism' has been put thus: 'there is a divine–human gap such that we humans should, for many evils in our world, not expect to grasp the divine purposes and reckonings behind God's allowing these evils' (Perrine & Wykstra (2017), 86). Amongst its other merits the simulation solution does not require us to hold that God is bafflingly inscrutable, or that we are cognitively closed to divine reasoning or to the real nature of good and evil. It also provides a solution to the various problems associated with 'divine hiddenness' that have recently been much discussed (see Schellenberg (1993) and Maitzen (2006)). If we accept the simulation solution we can suppose that God is only hidden in a problematic fashion in *our* (virtual) world, thanks to the design choices made by the less-than-divine creators of our world.
6. According to Seth Lloyd's calculation – in his 'Ultimate Physical Limits to Computation' (2000) – a 1 kg lump of matter, fashioned into the ultimate laptop, would be capable of performing 10^{40} operations per second at ordinary temperatures. Today's brains (and supercomputers) are less than glacial in comparison.

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

7. Even if the learning algorithms devised by (human) computer scientists fall short of their providing machines with a genuine all-purpose intelligence, there is a rival approach which could still succeed. Scientists engaged on the 'brain emulation' project aim to create computer simulations of human brains that are directly drawn *from* human brains with the aid of advanced neurosurgical and scanning techniques. Considerable resources are already being diverted into such ventures. In the US, reverse engineering the brain is one of the National Academy of Engineering's 'Grand Challenges' for the twenty-first century. In 2013 the EU invested 1 billion euros in its Human Brain Project (HBP) – its progress can be monitored here: <https://www.humanbrainproject.eu/en/>.
8. A lack of unanimity on the mind–body relationship is nothing new among Christian theologians. As long ago as the second century we find materialists and dualists locking horns. Some early writers (e.g. Irenaeus and Tertullian) adopted a wholly materialist metaphysic of human persons, and argued that resurrection involves the reassembly of the pieces of matter which originally constituted our bodies before our deaths, whereas others (e.g. Origen) favoured the dualist view that we consist of material bodies and immaterial souls.
9. <http://www.washingtonpost.com/news/speaking-of-science/wp/2015/03/03/a-paralyzed-woman-flew-a-f-35-fighter-jet-in-a-simulator-using-only-her-mind/>.
10. In 2015 a team of researchers announced the first 'injectable' neural interface, in the guise of a tiny nano-scale *mesh*: after being squeezed through the syringe's needle, this unfolds itself and merges with the neural tissues of the brain, potentially allowing us to connect electronic devices with individual neurons. For further details see 'Syringe-injectable electronics', Liu et al., <http://www.nature.com/nnano/journal/v10/n7/full/nnano.2015.115.html>.
11. See, for example, Meixner (2017).
12. A Cartesian might well insist that high-level cognitive abilities and decision-making reside solely in the immaterial soul in a way that makes them invulnerable to computational control. Even if this were the case, it would still be possible for the totality of a subject's sensory and perceptual experience to be machine-controlled, and this is all that's needed for Type-2 virtual realities. For more on Type-2s see Dainton (2012).
13. In the *Superman* tales of the 1950s and 1960s, Kandor was the capital city of Krypton, until it was captured and radically shrunk by the evil villain Brainiac. After recapturing it, Superman kept it on a ledge – in a glass jar – in his Fortress of Solitude.
14. The scheme is set out in Kardashev (1964). The fact that Earth-bound astronomers see no signs at all of cosmic-scale engineering in our universe should not be taken as evidence that it is impossible, or does not exist. It may simply be that the Earth *is itself* the product of a large-scale engineering project, but its creators have chosen to provide us with a night sky (or the appearance of one) that bears no trace of stellar manipulation.
15. Barrow (1998, 133) distinguishes a total of seven levels of control over increasingly small entities. Civilizations at level *Type I-minus* are 'capable of manipulating objects over the scale of themselves'; those at *Type IV-minus* are 'capable of manipulating individual atoms, creating nanotechnologies on the atomic scale and creating complex forms of artificial life'; those at *Type VI-minus* are able to manipulate 'the most elementary particles of matter (quarks and leptons) to create organized complexity among populations of elementary particles', whereas those with *Type Ω-minus* capabilities (the highest level) are able to manipulate the basic structures of space and time.
16. And if you think that the natural laws in our universe are such that a human-type consciousness cannot be generated by physical systems significantly smaller than a human brain there is a point to bear in mind: the laws in the underlying non-virtual reality might be different in this regard.
17. For more on computational irreducibility see Wolfram (1984), *Idem* (2002).
18. Alternative solutions to the Fermi paradox include the idea that the Earth is a Type-1 or Type-2 simulation, and unbeknownst to us the real universe is teeming with unconcealed alien civilizations – see Baxter (2000) and Ball (1973) respectively. The simulation solutions to the Fermi paradox and the problem of evil are similar in that both derive from a conflict between how our universe seems to be, and how it should be given certain hypotheses (e.g. life appeared almost as soon as it could on Earth and a sizeable proportion of species possessing technology avoid destroying themselves, the universe was created by a maximally benevolent and all-powerful God). There is, however, a significant difference: alternative solutions to the problem of evil that are remotely as plausible as the simulation solution are (very) thin on the ground, but

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press

it is otherwise with the Fermi paradox – by no means all of the seventy five solutions to the latter surveyed by Webb (2015) have much plausibility, but many do.

19. The simulation solution thus casts the First Commandment in a fresh light.
20. 6 August 2015, BBC report from Hiroshima survivor: 'It was hell on Earth'.
21. Moreover, the table in your room exists when no one is looking at it because the relevant information – the information required to supply appropriate sensory experience – remains safely stored in the VR-device memory, ready to be triggered if and when someone enters the room. A broadly Berkeleian view could thus be true of our level of reality, even if not ground-floor reality.
22. A point David Chalmers forcefully makes in his 'The matrix as metaphysics' (2005).
23. There is a further complication to note in this connection: anyone who adopts the solution now or in the near future should also believe that powerful simulation technology has in fact *already* been developed – albeit in that portion of reality which is simulating our world. My thanks to Nick Bostrom, Stephen Clark, Richard Gaskin, Simon Hailwood, Daniel Hill, Nicholas Nathan, Howard Robinson, Attila Tanyi, Peter Tyler, and Tom Winfield; I am also grateful for helpful comments from audiences at York, CEU, The Arctic University of Norway, and anonymous referees for this journal.

https://doi.org/10.1017/S0034412518000392 Published online by Cambridge University Press