

Morality, Religion, and Cosmic Justice

David S. Oderberg, *The University of Reading*

That which the wicked feareth, shall come upon him: to the just their desire shall be given. As a tempest that passeth, so the wicked shall be no more: but the just is as an everlasting foundation (*Proverbs* 10:24–25).

I. Introduction

There is a famous saying, whose origin is uncertain, that no good deed goes unpunished. Although not cited by him, this was no doubt the thought that inspired George Mavrodes's (1986) well-known article "Religion and the Queerness of Morality." In it he argued that although not logically incoherent, a certain sort of world in which moral obligations existed would be "absurd . . . a crazy world" (Mavrodes 1986, 581).

The world he had in mind was what he called "Russellian," after a notorious passage in Bertrand Russell's essay "The Free Man's Worship" (Russell 1903). There, Russell portrayed us as living in a wholly materialistic universe governed by impersonal forces and "accidental collocations of atoms," with no objective purpose or meaning to existence. The greatest of human achievements would ultimately be buried "beneath the debris of a universe in ruins." We could only, concluded Russell, safely build the "soul's habitation" [*sic*] on a "firm foundation of unyielding despair" (Russell 1903, 32). In such a world, according to Mavrodes, the performance of moral obligations does, at least sometimes, lead to no net Russellian benefit, that is, any benefit to the agent that is possible within a Russellian world. Indeed, such actions sometimes confer a net Russellian loss. "We have," he says, "moral obligations whose fulfillment will result in a net loss of good to the one who fulfils them" (Mavrodes 1986, 581).

That such a world would be "queer," "crazy" and "absurd," according to Mavrodes, should only strike those who believe that moral obligations are wholly objective and run deep in the fabric of reality. He does not purport to address the sceptic, the subjectivist, the expressivist, the nihilist or the existentialist. His aim is at those who believe in a rational moral

order that makes special demands on us and requires a “peculiar” (in the non-pejorative sense) sort of adverse judgement of those who do not meet those demands, as well as an equally peculiar commendatory judgement of those who do. For convenience, and at risk of the disapproving stare, let us call such people “serious moralists.” It might be thought that serious moralists ipso facto cannot believe the world is Russellian, but this is not what Mavrodes thinks. He does not elaborate, but we can, if we like, suppose the serious moralist is able somehow to incorporate moral obligations as understood above into his Russellian worldview. Or we can imagine the serious moralist denying he is a Russellian while claiming that but for moral obligations, the world would be Russellian. In other words, he believes in a kind of “Russellian world plus”: the world is as Russell believed, but there are moral obligations too. In either case, on Mavrodes’s argument, such worlds should strike the serious moralist as crazy and absurd given the lack of proportion between the performance of obligation and Russellian benefits. Note that it is hard to see how, on the second view, the serious moralist could appeal to anything beyond Russellian benefits that the extra ontology of moral obligations could bring with it.

There is, in my view, something right about Mavrodes’s diagnosis of the situation of the serious moralist who believes the world is Russellian (or Russellian plus), and of the very idea that a world with moral obligations in it could be Russellian. His cure, namely the truth of religion, in particular of Christianity, is proposed to guarantee that, in the Kantian phrase, happiness ultimately corresponds with virtue, that the Russellian losses in this world are not the last word because the world is not in fact Russellian after all. In this paper, I intend not so much to take off from Mavrodes’s observations but to make some different ones inspired by his line of thinking. Nothing in what follows should be thought of as contradicting what Mavrodes argues, but in an important sense it *bypasses* his argument altogether. I contend there is a way of looking at morality that, while superficially similar to Mavrodes’s, is irrelevant to it. Moreover, I will argue that even if Mavrodes’s qualms about the absurd status of morality in a Russellian world were wholly satisfied, we would still be faced with a fundamental strangeness that suggests, although it does not obviously entail, that the world could not be Russellian.

II. A Brief Statement of the Argument from Cosmic Justice

What I want to focus on is the absurd and crazy status of morality in a world without cosmic justice. The main argument can be put simply:

ACJ:

- (1) We live in a rational moral order.
- (2) If we live in a rational moral order, then there is cosmic justice.
Therefore, there is cosmic justice.

The argument consequential on *ACJ* is one from the existence of cosmic justice to the existence of a religion in which cosmic justice is truly administered:

ARS:

- (1) If there is cosmic justice, then there must be some religion in which cosmic justice is truly administered.
- (2) There is cosmic justice.
Therefore, there must be some religion in which cosmic justice is truly administered.

I begin with some relatively minor clarification of premise 1 of *ARS*. For the moment, take “cosmic justice” as a term of art referring to that activity or set of activities whereby every good person is rewarded and every evil person is punished. I will be more precise about this later, but one point that needs to be emphasised here is that I am not concerned with the idea of an ultimate or final reward/punishment. There is a legitimate question of whether, even for those people who do receive justice on this earth, it is enough justice to dispel concerns about whether the world would be a strange place without a kind of “final reckoning.” My argument puts this to one side, however, concentrating only on the presence or absence of worldly justice. Premise 1 of *ARS*, then, is justified by the observation that, however we understand “cosmic justice” consistently with the general definition just given, there is no cosmic justice in this world if this world is thought of as Russellian. Hence, if there is cosmic justice, it must be administered in non-Russellian terms, that is, it cannot be found in the world if the world is Russellian; so it must be found somewhere else. But the only plausible candidate for that “somewhere else” is within some larger order involving non-Russellian things. And the only likely candidate for such a non-Russellian order is a transcendent one involving a deity and/or other agents who administer cosmic justice. The existence of such an order entails that the world is not Russellian, and the name usually given to such an order is “religious.” There must be, then, some religious system containing cosmic justice, according to the argument. I make no claim as to whether there must be

one or more such religions, or what they are. Nor do I claim any attributes for the agent or agents who administer cosmic justice other than that it or they must have all the properties necessary for being an administrator of such justice. In fact, many of the attributes of the classical monotheistic God can be deduced from the proposition that the agent or agents administer cosmic justice, but that is no part of the argument.¹

III. Benefits and Losses

We need to be more precise about the notion of cosmic justice, but to do so it is desirable first to examine what Mavrodes calls Russellian benefits and losses. These are the only benefits and losses that can accrue to a person in a Russellian world. He never defines them (other than in the way just given, which is only partly illuminating) or even gives examples, so I propose to define a Russellian benefit/loss as follows:

RBL: A Russellian benefit/loss is a benefit/loss that is finite in duration and either:

- (1) wholly bodily in character;
- (2) wholly mental in character; or
- (3) partly bodily and partly mental.

That the benefit/loss is finite in duration captures the Russellian idea that human life (both individually and as a species) will come to an end, not just physically but in every respect. There is no immortality of any kind for human beings, whether bodily or spiritual. Given that condition, all benefits/losses will be only those that accrue to finite members of a finite species. Purely bodily ones are either bodily states (or events, processes, properties) or pertain only to bodily states (or events, processes, properties). Purely mental ones are either mental states (or events, processes, properties) or pertain only to mental states (or events, processes, properties). Most benefits/losses are combinations of these, unsurprisingly: wealth satisfies both mental states such as hopes and desires as well as producing bodily states such as pleasure and physical satisfaction, and so, usually (or at least often), does being loved, experiencing thrills, being

1. For one thing, the agent of cosmic justice would almost certainly have to be all-good, all-knowing, and sufficiently powerful to be able to administer cosmic justice with the scope, depth, and precision it requires. I will say a little more about this later in the paper.

entertained, and having power and influence; one can think of many others. Needless to say, a more detailed taxonomy of benefits and losses would be a large and difficult task. For present purposes, though, what has been outlined is enough to make the point that Russellian benefits/losses pertain only to the goods, both mental and bodily, of the material world. But “material world” needs to be interpreted broadly enough to include various kinds of non-reductionism, as long as they do not entail the existence of any ontology that contradicts the basic Russellian point that the world is fundamentally random and meaningless.

Now there are several important ways in which, for all that he points us in the right direction, Mavrodes’s discussion misses the mark – that is, misses what ought most to bother us about morality in a Russellian world. His principal claim about our world and its arguably Russellian nature is that we have moral obligations whose fulfilment either results in no net Russellian benefit or in a net Russellian loss. (From now on when I speak of losses and benefits I mean “Russellian” unless stated otherwise.) My first worry is that he says nothing about the *non*-fulfilment of obligations. These equally result, much of the time, in a net benefit to the omitters just as performers of obligations suffer a net loss. A small point, perhaps, but focusing exclusively on the lack of benefit in being moral can skew the problem away from what is most important – which, as just implied, is justice.

A more significant issue concerns the very notion of *net* benefit and loss, where the emphasis here is on the “net” rather than on “benefit and loss”; problems with the latter will be raised shortly, although all these worries are bound up together. How, one might wonder, would the accrual of a mere *net* benefit dissipate the queerness that preoccupies Mavrodes? A net benefit is a benefit that outweighs other losses. The idea, then, is supposed to be that if performance of an obligation gives the agent a benefit that outweighs other losses, morality will no longer be queer. But one might as well ask why there should be *any* losses in being moral, if the misalignment of benefits and losses is what is supposed to make moral obligations such strange things. After all, every loss is *some* disincentive to being moral, even if small. So a rational moral order ought to be one in which no one ever loses out by being moral and always loses out by being immoral: it should always be the case that if I am moral, it pays; and if I am immoral, then I pay. Mavrodes could reply that the whole point of a net benefit for a given action is that it makes the action “all worthwhile” or makes the action the rational one to perform all things considered (not just moral matters considered). Yet is that not implicitly to *acknowledge* that

all the losses that come with doing the right thing are *obstacles* to be overcome? True, the action might be worthwhile in the end, but the end is often a long time coming, and it may involve a mere whimper rather than a bang.

Which brings us to a related and perhaps obvious point, that the net benefit (and symmetrical loss for non-performance) might be vanishingly small. As long as the benefit is net, we are asked to suppose, then we are not in a crazy world after all. Yet such tiny benefits are cold comfort. Even that final, sweet thought that “after all, I did the right thing,” for all the satisfaction it might bring the man on his deathbed, might induce no more than a wan smile in the friends and loved ones who know how much this person who did the right thing suffered for it throughout his life. And for the non-performer of obligations, the prospect of a marginal net loss might be easy to brush aside when contemplating the many delights that can accompany a life assiduously dedicated to vice. So perhaps Mavrodes should have stated that the net benefits of morality need to be large enough to remove the queerness. How one is supposed to assess the degree of largeness would still be a mystery. Obligations come, of course, in grades of severity, so we would need a sliding scale of net benefits to place alongside them, a nice job for economists and all those who know the price of everything and the value of nothing, no doubt, but any criterion beyond either the very rough and ready (to put it mildly) or the purely subjective would be elusive. When it comes to morality – I am sure Mavrodes would agree – rough and ready benefit scales are not fit for purpose. As for the subjective, the scales will turn out to be so multifarious as to provide nothing like the objective answer we need to take the sting out of morality’s queerness.

Nor will an emphasis on general versus particular obligations get us any closer to finding a way in which net benefits and losses can remove the craziness identified by Mavrodes. He only speaks about the fulfilment of obligations without clarifying whether he means the particular or the general. I have a general obligation to be kind to my neighbours, and many particular obligations to exercise such kindness in specific circumstances. One might naively suppose that although, as Mavrodes correctly points out, net losses often occur to people who fulfil particular obligations, as long as one complies with the corresponding *general* obligations a net benefit will accrue. Such is not, of course, the case. If I am generally kind to my neighbours they will probably be generally kind to me, all things being equal. Apart from the requirement that all things be equal, it

does not take a leap of the imagination to see that the lives of so many who are generally kind, generally keep their promises, generally refrain from hurting others, and so on are and historically have been snuffed out by the cruelties of nature in the blink of an eye. Moreover, as Mavrodes (1986, 584) points out (speaking of immorality, but the point equally applies to its contrary) – and this is something that applies to particular as much as to general obligations – it is one thing to show that there is a net benefit in being moral and another to show that one is *better off* being moral. Even if there is a net benefit to me in performing my general obligations, for instance, this will not reduce by much the absurdity of morality in a Russellian world if it is also true that the benefit would equally accrue if I were not moral; and as long as there were enough people who would still act kindly to me even if I were unkind to them, such might be the case. Indeed, a combination of this circumstance and whatever Gygean powers of guile and cunning were available to me in a Russellian world might well ensure that the benefit was even greater were I to decline my obligations.

There is also a difficulty specific to general obligations. The fulfilment of a particular obligation at the last moment of one's life can be met with a loss, yet this does not undo any benefits that may have accrued to the agent on performance of all previous particular obligations (and conversely). By contrast, one can only assess whether a net benefit or loss has accrued on fulfilment of a *general* obligation once the time for further fulfilment has passed. Before this time it is always possible that performance – better, part-performance – of a general obligation will be met with a net benefit or loss that might be reversed at a later time, even if it is one second before the agent has died. Such a case would be unusual: imagine the kindly person on her deathbed, a supposed wealthy woman who, moments before expiry, finds herself repudiated and mocked by her loved ones when they discover she is dying penniless. They only ever repaid kindness with kindness on the expectation of a generous legacy. In this world, all her saintly behaviour was for nought and she is made to know it before being tossed coldly into a pauper's grave. A scenario more likely in a penny dreadful than real life, of course, but that is not the point. What matters is that it is logically impossible to guarantee that the performance of general obligations is always accompanied by a net benefit (conversely for non-performance). We can call this the *lateness problem*: it is always too late *for* this life to guarantee that the fulfilment of general obligations has a net benefit. (One could as easily call it the *earliness problem*: it is always too early *in* this life to have such a guarantee.)

Would Mavrodes respond that this absence of a logical guarantee merely heightens the craziness? No, because his view is that there *could* be a world like ours in all respects save for the fact that benefits and losses were correctly aligned with morality, as we will see shortly. Nevertheless, *does* the lateness problem for general obligations heighten the craziness? Yes, if we insist upon benefits and losses as the locus of that craziness. But for reasons already mentioned and for further ones I will provide, benefits and losses as such are the wrong place to look.

As far as general obligations are concerned, then, if the question is “Why be moral?” or better “Why submit to the demands of morality in a crazy world?” (where “crazy” means that benefits and losses are not guaranteed to line up with the performance or non-performance of general duties), it is hard to see how such a question can satisfactorily be answered by appeal to its being *probable*, in practice, that for the agent asking the question life will not be so absurd after all. One might wonder, in the spirit of Pascal, whether the game is worth the candle.

IV. Rewards and Punishments

Having raised a series of worries about net benefits and losses in a Russellian world, I now want to ask whether, even if they could all be allayed, the very concepts of benefit and loss are suitable for reducing or eliminating the absurdity indicated by Mavrodes. Recall his concern that in such a world, the fulfilment of obligations results (often) in a net loss. I added that non-fulfilment (often) results in a net benefit. Mavrodes has this to say: “I think that morality will not be queer in the way I allege, or at least it will not be as queer as I think, if it should be the case that every obligation yields a Russellian benefit to the one who fulfills it” (Mavrodes 1986, 582).

Mavrodes thinks, then, that if benefits (and losses, presumably) can be lined up in the right way with moral obligations, then the queerness will disappear. Here I differ with him fundamentally. First, the benefits that hypothetically accrued would, in a Russellian world, have to be merely accidentally correlated with the fulfilment of obligations. There could be no *law of nature* guaranteeing that my kindness to neighbours was met with a net benefit (whether in each particular case or generally, the lateness problem notwithstanding). It would have to be one, big Russellian cosmic accident, true for all people at all times. Well, what nature giveth, nature can take away, so again the serious moral agent – the one who takes

morality to run as deep as Mavrodes thinks it does – should wonder whether the rigours of morality are worth all the bother given the fragile nature of the benefits to accrue.

Suppose, on the other hand, that we could make sense of there being a law of nature underwriting the correlation between fulfilment and benefit, non-fulfilment and loss. Imagine, to put it as crudely as a Russellian or a Huxleyite might be happy with, that there were a morality gland (caused by the morality gene) that secreted, whenever we acted morally, a hormone producing an indescribably euphoric feeling; and when we declined our obligations, it secreted a noxious chemical that caused us pain and anguish. (The degree of euphoria or severity of pain, let us add, nicely tracks the importance of the various obligations.) Imagine further that biologists of morality have identified the exact causal mechanism behind such correlations and the biological law supporting it. Suppose we all agreed beyond doubt that these good and bad feelings constituted a perfectly adequate benefit and loss in a Russellian world.

In what way could we legitimately think of this far-fetched scenario as one in which the craziness of misaligned benefits and losses was removed? I do not think there is any such way, precisely because the emphasis is still exclusively on benefits and losses rather than on *rewards and punishments*. It is not so much that Mavrodes has missed the point of what the serious moralist should be worried about if the world is Russellian. By raising the misalignment of benefits and losses he homes in on important issues concerning the relation between morality and self-interest, ones that have of course been discussed ad nauseam since the Greeks. More importantly, though, he wanders within the vicinity of the more important target, since rewards and punishments do essentially involve benefits and losses; but they are much more than that. There are no rewards and punishments in the morality gland scenario. We might speak loosely of being rewarded by one's genes for acting morally in the same way that we speak of health as the reward of a good diet and exercise. But this is a mere metaphor. A reward or punishment essentially involves the following:

- (a) It is a benefit/loss.
- (b) It is conferred on an agent by an agent.
- (c) The conferral is because of something the recipient has done.
- (d) The conferrer makes condition (c) known to the recipient.

I am not suggesting (a)–(d) are the complete essence of reward and punishment, but they are at least part of the essence. In particular I omit from (c) anything concerning whether what the recipient has done must

be good in the case of reward and bad in the case of punishment. This is an interesting general question about the concepts at issue but can be left aside for our purposes, since we already have enough conditions to see why the morality gland does not literally reward good behaviour or punish bad behaviour: none of (b)–(d) are satisfied. It might appear that (c) is satisfied: the benefit/loss is sustained by the recipient because of something they have done. There is, however, no *conferral*, and even more importantly the relevant *psychological* element is absent. The “because” in clause (c) above is the “because” of *purpose*, that is, knowledge of the reason why something is done. The morality gland has no purpose for what it does, in which case it is an agent in the attenuated sense of a mere cause. Again, the agent, after feeling the effects of the hormone for long enough, will come to know with a modicum of reflection why they are as they are – but the effects and their causes will not be *made* known to him. Hence, the gene and its hormone certainly are an incentive to good behaviour and a disincentive to the opposite, but incentives are not rewards and disincentives are not punishments.

Reward and punishment are an essential part of justice: they are the administration or doing of justice as between agents who themselves are not only in relations of justice and injustice but in other moral relations as well. In the human legal system, unjust acts are punished and the act of punishment is itself a just act. The “adverse judgment,” as Mavrodes (1986, 580) puts it, of others who fail in their obligations is also a kind of justice: it is right and fitting that such a judgement should be made. Similarly, both law and human convention reward good behaviour: the legal system usually only rewards good behaviour fairly indirectly and in specific circumstances, say by using government contracts to reward efficiency, but human convention regularly rewards, whether by acts or favourable judgements, the behaviour of people that involves mere cases of ordinary justice such as returning lost property or resisting opportunities to cheat. We should by now have enough considerations before us to see that the morality gland does not dispense justice.

Suppose, though, that the morality gene had been synthesised by the proverbial mad scientist who released it into the population, whereupon it entered the genotype of every individual; and suppose his reason were precisely to reward the good and punish the wicked. Would this not be a genuine administration of justice sufficient for a Russellian world? Leave aside the obvious problem of previous generations who lacked the gene and the difficulty of supposing, any more than in the case where the gene occurred naturally, that such a world would have anywhere enough

similarity to our own for us to make reasonable assumptions about what it would be like. The mad geneticist would still have to make known to each agent, on each occasion, why they were being rewarded or punished. It is not enough that, through sufficient stimuli, they might *come* to know through reflection what was going on and why; the scientist would have to *inform* them of this. We all come to know through sufficient stimuli that on many occasions where we do bad things we get hurt even if not by any other person: when we despoil the environment it often comes back to damage us, but we are not punished by it. The scientist would have to be a mind-reader, capable of action at a distance, powerful beyond all human ken and wise enough to make King Solomon look like a fool. He might not have to be literally omniscient, yet his knowledge would have to be so extensive as not to make much difference, including knowledge of vast swathes of the past, perhaps also of the future and even knowledge of how things would have been in other possible situations (what theists call “middle knowledge”). This mad genetic engineer starts to look something like the God of classical theism, not a possibility the devout Russellian should wish to contemplate.

In any case, the point of such speculation is whether any *contingent* features could be added to the actual world, on the assumption it is Russellian, to dispel its absurdity. Mavrodes seems to think they could, but I beg to differ. Even if benefits and losses could somehow be tweaked to align with moral and immoral behaviour, and even if the correlation were underwritten by a law of nature, the world would still be every bit as queer as it should seem to the serious moralist, and not because of any further adjustments that could be made. Tweaking the world to the level of the mad scientist and his morality gene would be either to abandon the Russellian outlook in favour of some sort of theism or else to posit a world that was at worst metaphysically impossible, at best so remote from our own as not to be worth wondering about. Even if the latter is the truth, the fact is that our world is just not like that. For us, the mad geneticist is many worlds away.

V. Why Be Bothered about Cosmic Justice?

So what is it about rewards and punishments that makes the world so strange? For a start, it is nothing to do with whether reward and punishment should be understood in terms of retribution and recognition rather than deterrence and incentive. The strangeness cuts across that old debate:

many wicked acts go without retribution and/or deterrence. The vicious person who dies in his vice has no chance of suffering retribution or being deterred, nor does the saint who dies in virtue have any chance of being recognised in her life, let alone incentivised. (Of course there is posthumous recognition, but that is not itself a reward.)

Consider the human legal system and moral convention. No one would hesitate to call a society strange that possessed a legal system in which no evil deeds were punished (on the assumption that evil deeds were done). Perhaps its statute books were full of laws against this and that, yet not one was ever enforced. Such an arrangement would barely deserve the term “society,” but even if it could function an outside observer would call it, literally, crazy. Maybe we should deny that such a community of individuals even possessed a legal system, but that merely exacerbates the sense of craziness: what *are* they up to, we would wonder; and why all those so-called laws?

Similarly, imagine that no good acts, no obligations fulfilled were ever praised, recognised, or rewarded with benefits of some kind or other. Not even a community of the grimmest, dourest, black-clothed duty doers would recognise such a society as sane. Could people who refrain from rewards for moral behaviour really take morality seriously? Yes, it might be said: the evidence they take it seriously is precisely that they do not stoop to reward it. Note that this is *not* the same as saying that they take it seriously because they do not want anyone to think that morality *pays*. To repeat, the issue is not one of benefit and loss but of reward and punishment. The question I am raising is *not* whether I should be moral because it is in my interests to be so. Nor is it whether I should be moral because it is in everyone’s interests, or society’s interests as a whole. The question is whether, in a society without rewards for good behaviour and punishments for bad behaviour, the people really take morality seriously. To reply that they do not deign to sully themselves with rewards and punishments reminds one of the husband who, when asked by his wife if he loves her, says “Of course I love you.” “Then why don’t you ever show it?” His answer: “I do not need to show it. I do not need to shower you with affection, give you gifts, tell you how much I love you, as though I have to make a point of it. If I do it on some days I have to do it on all days in case you think I don’t love you on the days I do nothing. But if I do it on all days, it’s as though I have to keep reminding you that I love you, as if you or I should be in doubt. Eventually the displays of affection would become trivial and meaningless. Of course I love you. After all, I married you, didn’t I?” And so on ad nauseam.

Rewards for good behaviour are, as we might say, just like punishment for bad behaviour – fitting and proper acts of recognition. They manifest a society's attitude to morality, making plain that its members are, after all, serious moralists. So far, no problem: for there is no society, and never has been, in which people are *not* rewarded for good behaviour and punished for wicked deeds. But we have not yet reached our baseline for assessing the sanity of the system. For it is not merely that some people are rewarded and others punished. We want *everyone* to be rewarded and *everyone* to be punished for the appropriate behaviour, *all* of the time. It might not look like it, but the plain fact is that every "thank you," every "how could you?" every smile, and every headshake of disapproval is a reward or a punishment, however minor and seemingly trivial in the scheme of things. And when we cannot reward or punish, it is not because we do not *want* to, unless we are thoughtless, negligent, or just plain mean – all of these attitudes that themselves, in a sane society, will, when perceived, raise the very same shakes of the head and looks of annoyance that first-order bad behaviour generally elicits.

One important qualification needs to be made, however.² It might be objected that we have many reasons not to reward and punish in some situations, reasons deriving from forgiveness, pardon, or even the possibility of undesirable consequences of doing justice that may make it better, in the circumstances, to let a misdeed pass or a commendable act go unacknowledged. (A mild rebuke might make an otherwise good but sensitive child overreact; praise might go to a good but impressionable person's head.) So how can we even want to do justice in all cases (where we either have authority to do so or it is otherwise within our remit, to speak deliberately vaguely)? It is not possible here to discuss the important notions of forgiveness and pardon in depth. All that needs to be said, however, is that on either of two broad views of forgiveness and pardon, the present argument still applies. Suppose forgiveness and pardon, or the requirements of some particular circumstances, involve doing away with punishment altogether in some cases. Then the argument for cosmic justice applies only to those, all too common, cases where an agent is impenitent, lacking remorse, in no way merits forgiveness or pardon and where there are no other circumstances that call for justice not to be done. Alternatively, one might think that forgiveness and pardon, or the requirements of some particular circumstances, only ever serve to *moderate* a punishment, not to do away with

2. Thanks to Tim Chappell for pressing this point.

it. They might even make an appropriate punishment extremely mild without expunging it. If this is the case – and it is the view I favour – then the present argument still applies to all cases of wrongdoing – forgiveness, pardon and specific circumstances notwithstanding.

With this qualification in mind, it is still correct to say that we want everyone to get their rewards and their punishments and we try to realise this as much as we can (again, where we either have authority to do so or it is otherwise within our remit). The legal system has the intention of punishing whatever miscreance is within its authority, although it fails for all-too-human reasons. And so we fail as well in the application of our conventions because none of us can know everything or be everywhere and everywhen. The failure, more strikingly, is often a result of the fact that evildoers are often just too good at getting away with it and saints too humble and self-effacing even to think of publicising their actions. Overlying all of this is plain circumstance, bad luck and cruel nature. Think of the barbaric dictators with the blood of millions on their hands, who died outside the grip of justice – Stalin and Pol Pot are prime examples, and at the time of writing Kim Jong-il is soon to be another. Adolf Hitler is an interesting case. In one sense he escaped justice – never captured, never tried, let alone punished or executed. Yet he died a miserable suicide in his bunker surrounded by the ruins of his fanciful thousand-year Reich, fully aware of his failure. Was this justice, however? Could such an end really have made up for all the carnage and murder?

History, too, multiplies examples of heroic virtue that have never so much as seen the light of day. We usually find out about it after the person has died – the sick and dying they cared for, the charities they set up or ran, the money given to good causes, and the wayward individuals they reformed. Making sure their left hands did not know what their right hands were doing, barely a whisper of their activities attracted attention. Where was their reward or recognition? That they no doubt would have shunned it had they found it is irrelevant. Again, the question is not whether morality pays, so their rejection of a proposed reward disproves a point I am not trying to make. It is whether we should regard it as fitting and proper that they *do* receive their reward, however lacking in self-interest their motivation. I can put the point even more starkly. Suppose I were one of these saints (most certainly a mere supposition) and that I did receive my reward. (If it only came at the end of my life, then imagine I had been vouchsafed it from the beginning.) But suppose also that no other saint received *their* reward. If

I were interested in whether I should be moral because morality pays, I would be happy with my lot. And if the question were “Should anyone be moral if morality pays?” then the answer would be “It depends on who anyone is.” My concern, by contrast, cuts right across the question of morality and self-interest. No matter how great my reward, I should properly be outraged that other equally deserving people had not received theirs.³

In the case of punishment, were I to get my just deserts for a life of vice and no one else – or only some others – got theirs, then although the only outrage I would feel, because of the warp of my conscience, was the misplaced indignation that others had “got away with it” but not me, the *fitting and proper* reaction on the part of a sane and serious moralist should be outrage that some individuals received their punishment but not others. The concern with justice, then, is the concern with how *all* people are treated, not excluding oneself – but only inasmuch as one is just another person. The concern with *cosmic* justice, further, is the discomfort of the serious moralist with the possibility that justice is not done to all people at all times in respect of their actions. Such a possibility, I submit, makes a Russellian world a real source of fundamental unease about the way of things, since if the world is Russellian the serious moralist’s discomfort is wholly justified.

VI. Cosmic Justice and the Rational Order

Return now to

ACJ:

- (1) We live in a rational moral order.
- (2) If we live in a rational moral order, then there is cosmic justice.
Therefore, there is cosmic justice.

Clearly, a lot of work is being done by the concept of a rational moral order. Sometimes writers on the nature and justification of morality speak

3. True, there is an asymmetry here: people tend to be more concerned that wrongdoers receive their just punishments than that doers of good receive their rewards, but I take this to reflect an innate egocentricity in most people that tends to fade when great acts of goodness are compared with great acts of evil: ‘great acts of evil’. In more ordinary cases, egocentricity means that if I am punished for wrongdoing but others are not even if their acts were comparable, my first reaction is likely to be “Why me?” But if I am rewarded and others are not for comparably good acts, I am disinclined to think “Why them?” Rather, my focus will be on the contentment I experience from my own reward.

of a “moral enterprise” (e.g. Baier 1995), understanding moral order in terms of a collective social project or undertaking. The terminology is unfortunate: morality no more requires entrepreneurs than project managers. Morality is a *way of life*, or better, a way of *being* that makes comprehensive demands on every fibre of a person’s body and every recess of his or her heart and mind. Often it works against everything a person wants, loves or believes. Even when it harmonises with all of these it still regulates, in an all-encompassing fashion, every way in which a person carries out a plan, acts on a decision, treats or reacts to another, and so on. Morality has, somewhat infelicitous to say, a vice-like grip on human life. As such, we should expect that grip to be softened by the balm of rationality and explicability.⁴

For all the risks, I cannot help but compare morality to science. Science is, fair to say, more of an enterprise than morality: you can choose whether to be a scientist, that is, engage in scientific practice, but you cannot choose whether to be moral in the sense of engaging in moral practice. You can choose to be immoral, of course – you can even choose to be a nihilist, the ultimate immorality – but you do not thereby choose to refrain from engaging in moral practice. That notwithstanding, the person who chooses to be a scientist makes some very obvious assumptions: that the world is fundamentally rational and explicable; that there is a certain uniformity in the behaviour of nature; that even if human minds are too limited ever to understand it all, still the workings of nature are at least graspable by some mind or other, perhaps greater than ours only by

4. It might be thought that I am endorsing a view of morality that is sharply criticised by Bernard Williams (1985, Ch. 10). Williams calls morality a “peculiar institution,” but he has a number of objections to “the morality system.” He regards it as too abstract, too detached and excessively/exclusively focused on the notion of obligation, committed to the idea that moral obligations trump all other “ethical” considerations, and even seems not to think that morality involves a system at all. It is hard to know what his main worry is, but it must be said that the “morality system” he criticises is not one that I endorse. I explicitly acknowledge the numerous concepts apart from obligation that figure in moral thinking; the focus on obligation here is merely for ease of exposition. He is right to criticise much moral philosophy for its detachment from real life, but if this is supposed to mean that there can be no system of abstract principles at the heart of morality then we should differ. It is hard to see how obligations could *not* override other considerations, but this does not mean that morality is merely a system of obligations that overlies the rest of human life. On the contrary, the view I am defending sees all of human life as permeated by morality, so that all of life is moral in that sense. If Williams thinks that the very fact of there being no escape from morality is problematic – that there is something wrong with the idea of an entirely moralised world – then again we should differ. In short, the picture of morality I sketch is not the “morality system” deprecated by Williams, although it contains notions he would probably find questionable.

degree or else greater in kind. How could science even be practised if some such assumptions were not at its foundation?⁵ And if so, how much more should this be the case for morality, which is not a limited, chosen enterprise? (Here, the disanalogy only strengthens the point.) A rational moral order, then, is one in which, at least, there is a reason for the grip of morality on our lives. Morality as a comprehensive system – of rules, principles, guidelines, requirements, permissions, virtues, vices and so on – should be fundamentally rational and explicable. There should be a reason *why*, for example, the cruel tyrant died unpunished. The sort of reason required in this context is not “Because no one succeeded in bringing him to trial, despite all their best attempts,” any more than the question “Why are neutrinos like that?” has the mere answer “Well, that’s what our best experiments have revealed.” Neutrinos are like that, assumes the physicist, because there is something about the fundamental structure of reality, with all its interconnections, that has as a feature a certain nature and structure for neutrinos. In the moral case, if we make similar assumptions, the thought should be that the tyrant might have escaped unpunished at one level, but at a deeper level – the level of the rational moral order – there is a *reason why* this happened that does not offend against our profound sense of justice. The reason encapsulated in the thought that a greater punishment awaits the tyrant than any that could have been administered by human hands is one kind of reason that should satisfy that very sense.

Scientists assume a certain uniformity in nature: things, overall, are never going to be radically different from the way they have been; every departure can be explained, at a deeper level, in terms of another uniformity. (Think of mass extinctions, environmental catastrophes, the death of stars or the possible heat death of the universe itself.) The laws of nature

5. These assumptions should not be infused with too much detail. I do not propose them as propositions that rule out of court, *ab initio*, various philosophies of science that can be labelled instrumentalist, anti-realist, pluralist or particularistic and certainly not any that are anti-reductionist. It may well be that the more we flesh out these assumptions, the more we will see that at bottom they are incompatible with certain theories of scientific methodology and/or ontology. But stated simply as here, I take them to be held by all scientists and by those philosophers of science who are not anarchical in their view of what science does. It is a confusion to suppose that wholly general explanation must be reductionistic. Moreover, one can be a Fodorian about the special sciences and still assume and aspire to a comprehensive and rational explanation of the connections between them. As for anti-realism, even the instrumentalist, for example, regards the world as rational and explicable if only inasmuch as there are regular, uniform correlations between cause and effect, input and output, and so on. An anti-realist about so-called “theoretical entities” still believes that there are rational connections between the phenomena we want to “save.” This is enough commitment to comprehensibility to make my analogy.

are assumed to apply at all times and places. Moreover, scientists look for comprehensive explanations since these are, by their nature, more satisfying than partial or piecemeal ones. The more overarching an explanation – for instance, unifying electricity and magnetism or, so it is perhaps hoped in vain, general relativity and quantum mechanics – the more satisfying it is or is presumed to be because it unifies otherwise disparate phenomena. Scientists do not assume that at some level of generality explanation just *cuts out*; on the contrary, to put it crudely, they think it gets *better*. In the moral case, the analogous thought is that the coexistence of both justice and injustice in this world embodied in the ubiquitous misalignment of rewards and punishments with the behaviour to which they are proper is explained at a more general and comprehensive level. The partial satisfaction inherent in piecemeal and unpredictable justice in this world is embraced by a more satisfying and highly predictable justice that, if not present in this world, must exist in some other.

Further, scientists assume the world to be fundamentally comprehensible, even if not by limited minds such as ours. The moral analogue is the idea that even if *we* do not fully (or remotely, for that matter) understand why so many evil people “get away with it” and so many good people do not receive their due, it must be understandable by *someone*, whether by us in some better state of cognitive capacity or by some other mind much greater than ours. This of itself does not, of course, suggest that *we will* understand it or that there is some other mind that *does*. But if we combine this idea of epistemic accessibility with the previous thoughts that the moral system is both rational and uniform, and add to them the basic concern that cosmic justice is both fitting and proper, we do not prove that there must be such a mind and/or that we will, eventually, be in a position to understand how that justice works ourselves, but we do suggest strongly that one or both are true.⁶

6. Clearly the kind of rational moral order outlined here is substantial in its metaphysical commitments. To many moral theorists, the term “rational moral order” will have a more minimalist connotation. They will think that all one needs for the moral order to be rational is a system of internally coherent normative requirements of a special kind, with sufficient scope and applicability to count as a system that *could* be instantiated in a society of human beings (or rational agents more broadly conceived). Such a conception neither says nor implies anything about which if any of the requirements will be met by agents, including whether, say, justice will be done in cases where it ought to be done. Nothing in my more metaphysical conception of a rational moral order is inconsistent with the minimal conception; rather, it goes beyond it. Moreover, it would require argument to show that there is a more minimal conception similar to the one just mentioned that is both plausible and *rules out* anything more metaphysically laden. On the other side, the metaphysical

VII. Cosmic Justice and the Point of Morality

I have claimed that morality is neither a project nor an enterprise but a way of being with no opt-out. The interesting question here is not so much “Why be moral?” but “What’s the point of morality?” The latter question is one about the rationality of the system. It will not do to reply that we have no choice in the matter. On the contrary, it is the very fact that we have no choice that makes the question more urgent. If there is no point, then we are entitled to ask why the human situation is not fundamentally, inescapably, meaningless and hopeless. Constructing meanings for ourselves along the way and decorating them with man-made hopes is no answer to the question since it misses the point. If the human situation is fundamentally without hope or meaning then no amount of artificial scaffolding can hide that fact from us, at least not forever.

It is not just the inescapability of morality that generates the question of its point but the massive demands it makes on us, demands that always and everywhere trump every other. Could it be that such a system of absolute and unconditional regulation of every human being has, at bottom, no point? No serious moralist should think that any amount of appeals to self-interest, the interests of others, the interests of society, or benefits and losses, can furnish that point. All of these incentives can be present and yet, if cosmic justice is not done, morality will still be futile and pointless. Moreover, cosmic justice, on this view, has to be *exceptionless*. It is not adequate that *enough* justice be done for me to have good grounds to believe that morality has a point; and this not because I might be an exception. Rather, if *anyone* is an exception, then morality loses its point. The outrage one should feel at the lack of justice involving another

conception of a rational moral order I am proposing is more plausible than the minimal conception inasmuch as it concerns the way in which the internally coherent and applicable system of rules *is* instantiated or lived. Perhaps it is better to say that the defender of the more minimal conception is talking at cross-purposes with the defender of the more metaphysical conception. The former thinks of morality as a system of rules and requirements. The latter, while not denying this, is not interested in it as such, but in the way that system is instantiated or lived by the agents to whom it applies. Similarly, when one analyses a human legal system one might be interested only in the abstract framework of laws. One might, however, be interested in the framework of laws as realised in the society to which it applies: one might want to know how the laws are enforced, the way in which sanctions are applied, what the consequences of certain laws have been for that society, what the incentives and disincentives are for obedience and disobedience, respectively, and so on. So when I talk of the rational moral order, I am concerned with the way in which an abstract system of rules and requirements has its concrete manifestation in human life. (Thanks to Brad Hooker for discussion of this issue.)

person, even if justice be done to oneself, is, if anything, purer and nobler than if justice were done to all *but* oneself.

It is important to appreciate what I am *not* claiming and what I mean by the “point” of morality. Cosmic justice does not compete directly with other objectives that morality might have. One might think the objective of morality was human flourishing, the doing of what is right because it is right, the maximisation of utility or something else. Still, as long as the doing of justice was either a *component* of such an objective or a necessary correlate of it, then my claim is that the absence of cosmic justice would make morality pointless or futile by thwarting or undermining its objective. The objective of morality, then, is what morality aims at. The *point* of morality is the meaningfulness of pursuing that objective. Sisyphus had an objective – endlessly to roll a stone up a mountain. Pursuing the objective, however, was pointless. It is in this sense that I claim the pointlessness of morality without cosmic justice. On the other hand, one might not think morality *aims* at anything – the idea being regarded as confused, comparable perhaps with the claim that there was an objective to being rational or seeking the truth. Even so, morality considered as without objective would still be pointless in the sense of meaningless without cosmic justice, or so the present argument goes.

VIII. Some Objections

A number of objections to the line of argument I have been presenting suggest themselves. The first is familiar: one person’s *modus ponens* is another’s *modus tollens*. The critic can deny the conclusion of either *ACJ* or *ARS* and so infer that premise 1 of *ACJ* is false: there is no rational moral order. Well, the assertion that there is no cosmic justice might after all be no more an assertion – the declaration that there is a phenomenon that the above line of argument denies on grounds that are supposed to grip anyone who takes morality seriously in the way suggested. More charitably, it might point us towards the tsunamis of ink that have been spilled in the debate over whether the world is Russellian after all. Perhaps there are arguments to this conclusion that are strong enough to overturn the argument from cosmic justice. To be sure, my argument is not intended to knock down the Russellian’s impressive pile; for that, I would need to introduce the anti-Russellian’s own more persuasive (in my view) machinery. This is a battle for other times and occasions. For now, all I claim is that no Russellian should fail to take the argument from cosmic

justice into account in his reckoning – if he takes morality in the way that I, following Mavrodes, have outlined.

Another way of taking the denial that there is cosmic justice is as an assertion of the contradictory of the conclusion of *ACJ* rather than the contrary: it is not that there is cosmic *injustice*, only that there is neither cosmic justice nor injustice. The criticism implicit in this point is that *ACJ* only has bite if we think there is cosmic injustice in a Russellian world. But for there to be injustice there has to be a *failure* of justice; and a failure implies the existence of one or more agents who have themselves failed to administer justice. Note that the critic is appealing here to the *second* order of justice, not the first. At the first order, there is obvious injustice: people wrong each other, violate each other's rights and the like. The denial of cosmic justice, on this interpretation, is the claim that there is no failure, at the second level, to make good these first-level injustices. There is, in other words, no failure in the *administration* of justice. The defender of *ACJ* will reply, of course, that there are obvious failures in the administration of justice: a person is assaulted and his or her assailant goes to court and walks away scot-free because of a corrupt judge or a negligent prosecutor. Note that there is an issue of culpability here. What if the defendant walks free because of a quirk in the legal system that is no one's fault? Where is the injustice? I will return to that in a moment, but it is enough for now that the defender of *ACJ*, so it seems, can point to at least one culpable failure.

The problem for the defender, though, is that the critic can dig in his heels at this point and insist that such cases are not second order after all, not in the sense required by *ACJ*, although they do pertain to the administration of justice. The critic might say that these failures are just more cases of first-order injustice: the victim, in our example, was treated badly once by the assailant and again by the judge. How does this generate a case of cosmic injustice? What the defender needs, so the critic might insist, is to point out a failure to administer even the administration of justice; and for every individual or institution he points to – an appeal court, the highest court in the land, the ministry of justice and the head of state – as failing culpably to treat the victim in the way he deserves, the critic will insist that these are all just more first-order injustices. Who, he will ask, has actually *failed* to make good all of these wrongs? And so where is the injustice on the cosmic scale?

Three points of reply are in order. First, I am not sure that much hangs on whether the critic asserts the contradictory of the conclusion of *ACJ* rather than the contrary. To assert merely a lack or absence of

cosmic justice rather than actual injustice is still to identify a flaw in the moral order which is equally worrying: justice cuts out at the cosmic level not because any person is being cosmically unjust and so should have his or her injustice remedied by another who is cosmically just, but because there is a complete absence of the administration of justice at the cosmic level.

A second reply is that we do not need a failure to administer justice in order to see that some situations are unjust. If the murderer escapes justice through no one's fault, this state of affairs is still unjust in the sense that we feel rightly that it is fitting and proper for it to be remedied. It is tempting to assimilate such states of affairs to what is now known as moral luck (Nagel 1979; Williams 1981), only in this case it is not an agent who is assessed for something involving matters beyond her control, but a state of affairs is assessed although it involves matters beyond any agent's control. There is perhaps something to this comparison, but we should be careful. In classic cases of moral luck, it is not always plausible to think it fitting and proper that the luck should influence our moral assessment of the relevant action (for example genes or spatio-temporal location). In the cases we are considering, I submit that it is always fitting and proper that the bad luck involved in an evildoer's escaping justice or a saint's missing her reward should influence our judgement that the situation is inherently unjust and demands some sort of remedy. More generally, some situations simply are unjust although no one has been unjust in causing or failing to remedy the situation. As long as a person's rights have not been respected, we are justified in calling the situation unjust even if no one has failed to respect their rights.

Connected with this point is the third reply that we can demonstrate the injustice of a situation on a cosmic scale by appeal to the injustice of its constituents. If evildoer A escapes justice, saint B misses her reward, and so on across every case throughout time (although this is not strictly necessary), then these all add up to cosmic-level injustice. Is there a fallacy of composition lurking here? No, since the fallacy of composition must be handled with care; it all depends on what we are talking about. If the parts of an object are all and exclusively wooden, so is the whole, but if they are all square it does not follow that the whole is square. All the fallacy means is that one cannot infer the property of a whole, *whatever the property*, from the same property of its parts. In the case of states of affairs, the same applies. If the constituents are all desirable, for instance, it does not mean that the complex state of affairs constituted by them is desirable. On the other hand, if all the constituents obtain at noon, then so does the

complex whole. It would take us too far afield to work out a criterion for distinguishing those properties that produce the fallacy from those that do not, but it seems clear enough that if all the constituents of a complex state of affairs are unjust, then so is the whole. Therefore, we have a reason for thinking that there is cosmic injustice, or at least a lack of cosmic justice, in a Russellian world. In short, the critic does not fare any better in denying the consequences for justice if the world is Russellian than he does by insisting that the world *is* Russellian and that we are therefore stuck with admitted cosmic injustice or a lack of justice.

Another objection is that too much emphasis is being placed on whether the world is Russellian. For the world might not be Russellian and yet still unjust, for instance if a supernatural evil demon were running the show, or an “evil god” as in the Evil God Challenge (Law 2010). This would be another way of challenging the conclusion of *ACJ* but without relying on the assumption that the world is Russellian. True, but there is no assumption lurking behind *ACJ* that if there is no cosmic justice the world must be Russellian; it is the converse, namely that if the world is Russellian there is no cosmic justice. But there is cosmic justice since we live in a rational moral order. So the world cannot be Russellian. Premise 1 of *ARS*, however, implies that cosmic justice rules out not just Russellian worlds but all worlds in which cosmic justice is not done. There must be some religion in which it is done, and since the existence of an Evil God would entail the truth of a religion in which cosmic justice is not done, the Evil God Hypothesis cannot be true either. Or if there were an evil god, there would have to be a broader system of cosmic justice, a kind of super-cosmic justice, which made good all the defects caused by an evil god. And for every iteration of such hypotheses at higher levels, there would always have to be, at some highest level, a system of cosmic justice.

IX. Conclusion

Premise 1 of *ARS* rules out any religion in which cosmic justice is not done. It excludes religions with petty, spiteful, whimsical, capricious or arbitrary divine beings. It rules in, for the time being, religions in which, at least by their official teachings, there is a comprehensive system of rewards and punishments for good and bad behaviour, respectively. On further inspection, what looks just might not be just, and what one religion counts as good and bad might not accord with

another or with reason and reflective judgement. The Argument from Cosmic Justice is not supposed to provide a method for working out which religion is true. All it does is rule out obvious non-starters and propose that there must be some economy or dispensation (to use theological language) in which justice is comprehensively done. For all the argument shows, the religion in which this is the case has not even been made known yet to human beings. There are independent reasons for thinking this fanciful as well as for thinking that we already have very good candidates for religious systems that do teach the existence of cosmic justice and so are aligned on this score with the metaphysical truth. The classical monotheisms all have excellent claims in this regard. They cannot all, however, adequately explain every case of an innocent and sinless man's being subjected to the indignity of public humiliation, torture, and execution. To this extent, Christianity has something more informative to say than its rivals; but on this specific point the Argument from Cosmic Justice must remain silent.⁷

Department of Philosophy
University of Reading
 Reading RG6 6AA, UK
 d.s.oderberg@reading.ac.uk

References

- Baier, K. (1995). *The Rational and the Moral Order*. Peru, IL: Open Court.
- Law, S. (2010). "The Evil-God Challenge." *Religious Studies* 46: 353–373.
- Mavrodes, G. (1986). "Religion and the Queerness of Morality." In L. Pojman and M. Rea (eds.), *Philosophy of Religion: An Anthology*, 5th edition. Belmont, CA: Wadsworth, pp. 578–586. (Originally in R. Audi and W. Wainwright (eds.), *Rationality, Religious Belief, and Moral Commitment: New Essays in the Philosophy of Religion*. Ithaca, NY: Cornell University Press, 1986.)
- Nagel, T. (1979). "Moral Luck." In T. Nagel (ed.), *Mortal Questions*. Cambridge: Cambridge University Press, pp. 24–38.

7. I am grateful to Tim Chappell and Craig Thompson for comments on a draft of this paper, as well as to staff and students at the University of Hertfordshire and at the University of Reading Work in Progress seminar.

- Russell, B. (1903). "The Free Man's Worship." In L. I. Greenspan and S. Andersson (eds.), *Russell on Religion: Selections from the Writings of Bertrand Russell*. London: Routledge, pp. 31–38.
- Williams, B. (1981). "Moral Luck." In B. Williams (ed.), *Moral Luck*. Cambridge: Cambridge University Press, pp. 20–39.
- . (1985). *Ethics and the Limits of Philosophy*. London: Fontana Press.