

Faith and Philosophy: Journal of the Society of Christian Philosophers

Volume 23 | Issue 4

Article 6

10-1-2006

God and Infinite Hierarchies of Creatable Worlds

Bruce Langtry

Follow this and additional works at: <https://place.asburyseminary.edu/faithandphilosophy>

Recommended Citation

Langtry, Bruce (2006) "God and Infinite Hierarchies of Creatable Worlds," *Faith and Philosophy: Journal of the Society of Christian Philosophers*: Vol. 23 : Iss. 4 , Article 6.

DOI: 10.5840/faithphil200623439

Available at: <https://place.asburyseminary.edu/faithandphilosophy/vol23/iss4/6>

This Article is brought to you for free and open access by the Journals at ePLACE: preserving, learning, and creative exchange. It has been accepted for inclusion in Faith and Philosophy: Journal of the Society of Christian Philosophers by an authorized editor of ePLACE: preserving, learning, and creative exchange.

GOD AND INFINITE HIERARCHIES OF CREATABLE WORLDS

Bruce Langtry

I argue that *For every creatable world there is another one that is better than it does not entail There does not exist an omnipotent, omniscient and perfectly good being.* An agent choosing which world is to be created from an infinite hierarchy of candidates should satisfice. For there are independent reasons for resisting the inferences from *You have better reason to choose A than you have to choose B to You should choose A rather than B* and to *If you were to choose A rather than B then you would be acting in a morally better or more rational way.*

William L. Rowe's recent book *Can God Be Free?* advances what I will call 'the Entailment Thesis':

The proposition *For every creatable world there is another one that is better than it entails the proposition There does not exist an omnipotent, omniscient and perfectly good being.*¹

Section I deals with Rowe's arguments directly in favour of the Entailment Thesis. Section II investigates the content of the thesis. Section III argues against the thesis. Section IV anticipates how Rowe would respond to Section III, and accordingly reinforces my main arguments. Section V raises a final difficulty for defences of the thesis.

I. Rowe's Positive Arguments for the Entailment Thesis

Rowe's main argument is that the thesis follows from the following proposition, 'Principle B,' which he declares to be a necessary truth:

If an omniscient being creates a world when there is a better world it could create, then it would be possible for there to be a being morally better than it.²

Rowe acknowledges that 'we are far from proving' Principle B,³ but offers two arguments in *prima facie* support of it. Here is the first:

For if an omniscient being creates a world when it could have created a better world, then that being has done something less good than it could do (create a better world). But any being that knowingly does



something (all things considered) less good than it could do falls short of being the best possible being.⁴

There is an important ambiguity in the expression 'do something less good than one could do.' Rowe's first premise can be understood in either of the following ways:

- (1a) If an omniscient being creates a world when it could have created a better world, then that being has thereby performed an action that is less good—i.e., less morally good, less rational, or less a manifestation of excellence in acting—than an alternative action that the being could have performed
- (1b) If an omniscient being creates a world when it could have created a better world, then that being has thereby performed an action whose total causal consequences are less good than those of some alternative action that the being could have performed.

(1b) is plausible; but to just about everyone except maximising act-consequentialists, (1a) will surely seem no less doubtful than Principle B itself, and so cannot contribute to a strong argument for the principle. For example, deontologists and rule-consequentialists will be able to envisage cases in which an omniscient being, given a choice of which of two candidate worlds is to be created, has a moral duty to choose the inferior world rather than the better one.

Rowe's second premise can be understood in either of the following ways:

- (2a) Any being that knowingly performs an action that is less good—i.e., less morally good, less rational, or less a manifestation of excellence in acting—than an alternative action that the being could perform falls short of being the best possible being.
- (2b) Any being that knowingly performs an action whose total causal consequences are less good than those of some alternative action that the being could perform falls short of being the best possible being.

Here (2a) is plausible.⁵ But unless we are already maximising act-consequentialists, why should we regard (2b) as more secure than Principle B itself, and so an a good starting point for a persuasive argument for Principle B? As for maximising act-consequentialism, in Section V I will explain why it is far from clear that this doctrine can be used as a premise supporting (2b)—or (3a) or (3b), stated below—and thereby supporting Principle B.

To sum up: (1b) and (2a) look initially plausible, but they do not jointly entail Principle B. It is plain that if we interpret Rowe's argument in such a way that it is free from equivocation with respect to the expression 'do something less good than one could do,' and its premises jointly entail Principle B, then either (1a) or (2b) requires a great deal to be said in its favour before most philosophers will be impressed by Rowe's argument for Principle B, and by the argument based on Principle B for the Entailment Thesis.

Rowe discusses a number of major philosophers from Augustine onwards who have asserted that God always acts in accordance with what is best. But 'act in accordance with what is best' is ambiguous between 'act in the best way' and 'perform the action that has the best consequences.' Once we have noticed the ambiguity, it is hard to see that Rowe has identified, in the philosophers he discusses, any good argument in favour of the claim that God always performs the action with the best consequences.

Rowe's discussion of the recent literature contains a second positive argument in direct support of the Entailment Thesis; it also constitutes an argument for Principle B. Here is the major premise:

- (3a) When the agent's motive in acting is to bring about a good state of affairs, the goodness of an action—measured in terms of the quality of its result—is expressive of the agent's goodness.⁶

The context makes it clear that the goodness of an action and of an agent are to be understood as their *moral* goodness. But, so understood, the principle is obviously false. An agent may be motivated by a desire to bring about a good state of affairs, and this desire may be fulfilled, but the moral goodness of the action, and of the agent, may diverge far from the quality of the result. For example, the quality of the result can vary with luck and with the interference of other people, without there being any variation of the moral goodness of either the action or the agent.

Perhaps Rowe overlooks the foregoing points because he is thinking of an omnipotent and omniscient creator, who does not find the outcomes of her providential activity affected by luck or by the interference of other people. But it is hard to see how Rowe can obtain a good argument for the Entailment Thesis by somehow restating (3a) so as to get a premise with a more restricted scope. Consider, for instance:

- (3b) If an omnipotent and omniscient agent's motive in creating a world is to bring about a good state of affairs, then the degree of goodness of the world measures the degree to which the agent is morally good.

Suppose that there exists a being G who is omnipotent, omniscient and moderately good, and consider two worlds V and W such that W is better than V but if G were to create W she would violate an important moral duty recognised by deontologists and rule utilitarians, while if she were to create V she would not violate any moral duty. It is true that even deontologists and rule utilitarians admit that at least some moral duties are defeasible by considerations of utility. But why should they think that given that G, V, and W are as described, we should conclude, as (3b) requires us to conclude, that G's creating W would indicate her possession of a higher degree of moral goodness than would be indicated by her creating V?

I hold that *other things being equal*, in intentionally bringing about the better state of affairs one acts in the better way, all-things-considered. In Section III I will argue that this true principle does not provide the basis for a good argument for the Entailment Thesis.

II. Refining the Entailment Thesis

Inferring the Entailment Principle from Principle B requires the following auxiliary assumption, which I will call *Creator*:

Necessarily, if there exists an omnipotent, omniscient and perfectly good being then it creates some world.

Rowe argues in favour of *Creator* when he says: 'Given that it is better to create a world than not to create at all, God is doing the right thing in choosing to create a world rather than not to create at all.'⁷ Whether this argument is a strong one, and whether *Creator* is plausible, turn out to depend on exactly how the word 'create,' in such contexts as the Entailment Thesis, Principle B and *Creator*, is to be understood.

On p. 41 Rowe says that in creating a world God makes actual the contingent states of affairs contained in that world. He then adds a footnote:

Suppose God creates some living creatures who then create works of art. There being such works of art is a contingent state of affairs not directly created (made actual) by God. But these works of art would not exist had God not created living creatures. So, ultimately, all contingent states of affairs depend for their existence on God's creative activity.

I suspect that Rowe is writing with indeterminism—indeed, with libertarianism—in mind. If Rowe's argument from Principle B to the Entailment Thesis depended on the suppressed premise that necessarily if God exists then he causally determines every contingent state of affairs in whichever world is actual, then Rowe would certainly have realised this, and would have seen a need to highlight and defend such a controversial proposition. On one reading of the footnote, it expresses the view that if God creates a specific world then its constituent states of affairs stand in a variety of causal relations to God, but each one at the very least *depends* on God's activity, and so the world can be said, in a very broad sense, to be 'created' by him.

If this broad sense of 'create' should be applied to the Entailment Thesis and Principle B, then how about 'creatable'? Perhaps the suffix '-able' signifies the mere logical possibility that some person creates the world in question. If so, then here is a candidate logically necessary and sufficient condition of creatability:

A world *W* is creatable if and only if in *W* there is some person on whose activity every contingent state of affairs is at least causally dependent.⁸

(The expression 'every contingent state of affairs' has, in this and similar contexts, some 'basic exceptions' which I will usually leave unexpressed; they include *P's existing* (where *P* is the relevant person) and contingent states of affairs entailed by *P's existing*, if *P's existing* is contingent, and the truth of the true counterfactuals of creaturely freedom, if there are any.)

This proposal will not suit Rowe, because its adoption undermines the Entailment Thesis (and also undermines his arguments for the Entailment Thesis). For suppose that there are two worlds V and W such that (i) W is better than V, (ii) in each of them every contingent state of affairs depends on the activity of some omnipotent and omniscient agent G, but (iii) in W but not in V there occur certain specific states which while dependent on G are nevertheless uncaused by G. If V is actual then we have a situation in which G is omnipotent and omniscient, and 'creates' a world (according to the proposed account) even though there is a better 'creatable' world. But since G could not have caused W to be actual, surely we should not infer that in this case G is not perfectly good. So why should we accept the Entailment Thesis?

Plainly, Rowe needs an alternative account of creating and of creatability which will enable him to retain the spirit, if not the letter, of the Entailment Thesis as formulated in *Can God be Free?*

Let us try to make further progress with such matters by using the familiar expressions 'strongly actualise' and 'weakly actualise' to introduce the additional technical terms 'securely actualise' and 'securely actualisable' as candidate equivalents to the informal words 'create' and 'creatable'.

A person G strongly actualises a state of affairs if and only if G causes and determines it to be actual. G weakly actualises a state of affairs if G does not strongly actualise it and yet performs some action that both causally contributes to its occurrence and is such that, if G were to perform it then the state of affairs would be actual. G *securely actualises* a world if and only if either (a) G strongly or weakly actualises every contingent state of affairs included in the world or else (b) there are counterfactuals of creaturely freedom true in the world and G strongly or weakly actualises every contingent state of affairs in the world except the obtaining of the counterfactuals of creaturely freedom.⁹

Given the foregoing definition of 'securely actualises,' what meaning should be assigned to 'can securely actualise'? If 'can' is treated in this context as expressing logical possibility, then G can securely actualise W if and only if W contains the state of affairs *G's securely actualising W*.

Many philosophers, however, say that even if a person G is omnipotent, there is a sense, important for discussions of providence, in which what G 'can' do is more limited than what it is logically possible for G to do. For example, Molinists declare that *God's* options are limited by contingent truths holding prior to his decisions. Of all the logically possible worlds in which both God exists and there are free creatures, God 'can' securely actualise only some, and which ones they are depend on what counterfactuals of creaturely freedom are true.¹⁰

Given that there are contingent limits to what an omnipotent agent can do, there are competing criteria of the agent's being able to do something. In some contexts, if you succeed in doing something by sheer fluke then this licenses the remark that you *could* do it. In other contexts, people will agree that that you can do something only if you meet certain standards of competence. In yet other contexts, we say that you can do something only if you are not only competent but you know that you have a suitable opportunity — e.g., there is a sense in which, however skilled you are, you cannot play the flute if no flute is within reach, or if a flute is within

reach but you do not know where to find it, or if other people will prevent you from playing if you try. This strong requirement for ascribing ability is obviously relevant to accounts of what omnipotence, omniscience and perfect goodness jointly imply, offered by theories recognising contingent constraints on the plans of an omnipotent being.

Alongside the explanation of 'creates' and creatable' centred on dependence, we now have a second proposal: a person G creates a world if and only if G securely actualises it, and a world is creatable if and only if *some* person can, relative to constraints obtaining in the actual world, securely actualise it.¹¹ Given this explanation, when Rowe says, in the footnote already quoted from his p. 41, that 'all contingent states of affairs depend for their existence on God's creative activity,' the dependence relation he has in mind is that the contingent states of affairs are securely actualised by God.¹²

Furthermore, there is a third proposal that should be taken seriously.

Suppose that there is a world W which contains some contingent state of affairs F (not a counterfactual of creaturely freedom) such that it is logically possible that G securely actualise F, and which also contains F's *not being securely actualised by any being*. It is logically impossible for G to securely actualise W. Suppose in addition, however, that in W it is true that if G were to perform some action (or sequence of actions) A then F would *probably* obtain as a result.¹³ Suppose that in W G does A, and doing A causally contributes to F's obtaining. Even though in W G does not securely actualise W, if in W G securely actualises all contingent states of affairs other than F then there is a sense in which we can sensibly say that in W G *weakly causes* W to be actual. If, in the actual world, G can perform an action of the kind indicated, then G *can* weakly cause W to be actual.

Thus we need to decide between three alternative interpretations of 'creates' and 'can create,' in terms of their suitability for Rowe's adoption in stating and defending the Entailment Thesis, Principle B and related propositions. Here are the three corresponding proposals for necessary and sufficient conditions of *Person G creates world W* and *Person G can create world W*:

- G creates W if and only if W is actual and every state of affairs included in W depends on G's activity. G can create W if and only if in W every state of affairs depends on G's activity.
- G creates W if and only if W is actual and G securely actualises every state of affairs in W. Two alternatives concerning 'creatable': G can create W if and only if in W G securely actualises W. Alternatively: G can create W if and only if in the actual world G possesses the power, skill and opportunity to create W.
- G creates W if and only if W is actual and each state of affairs included in W is either securely actualised or weakly caused to be actual by G. G can create W if and only if in W G creates W. Alternatively: G can create W if and only if in the actual world G possesses the power, skill and opportunity to create W.

Which interpretation of 'creates' is most suitable for Rowe's purposes? I have already indicated a serious difficulty with using the first when affirming and defending the Entailment Thesis. One reason to prefer the

third interpretation to the second is that the third makes the principle Creator—which Rowe needs in order to derive the Entailment Thesis from Principle B—look a lot more plausible than it would look under the second interpretation. For suppose there is some world W which G can weakly cause to be actual but not securely actualise, and which is better than any world that she can securely actualise.¹⁴ Why should G not take steps aimed at weakly causing W to be actual, and succeed in actualising W , and thereby refrain from securely actualising any world?

On the other hand, if Rowe were to adopt the third interpretation then Principle B, on which he relies so heavily, would be far from compelling. Suppose that in the actual world α there exists an omniscient agent G who has either securely actualised α or weakly caused it to be actual, and there is some non-actual world W which is better than α and which G can weakly cause to be actual but cannot securely actualise. (Doubts about whether this supposition is logically possible would obviously provide other reasons for rejecting the third interpretation.) It follows that an omniscient being G has created a world (namely, α) when there is a better world (namely, W) that G could have created. But why should we conclude that it is possible that there to be a being morally better than G ? There is an obvious possible explanation, consistent with moral perfection, for G 's creating α (by securely actualising it or weakly causing it to be actual), rather than creating W (by weakly causing it to be actual instead): although W is better than α , the relevant probabilities are such that G 's taking the steps she in fact took had a higher *expected value* than her taking the steps which would probably but not certainly have resulted in W (and there were no overriding reasons for reaching a contrary decision).

It might be objected that talk of expected value is out of place when we are discussing decision-making of omniscient agents. Not so—at least, not unless it is logically necessary that if an omniscient agent exists then either determinism or Molinism is true. Suppose that God is omniscient and knows, timelessly or in advance, that he will weakly cause world α to be actual. There is an obvious circularity in the suggestion that God could rationally use this item of knowledge as part of his reason for resolving to weakly cause α to be actual.¹⁵ Hence even given that God is omniscient, his reasons for action will in some cases need to involve consideration of probabilities—unless either Molinism or determinism is true, in which case the second account of creating is to be preferred.¹⁶

I conclude that Rowe faces severe difficulties whichever of the three accounts he adopts. Nevertheless the general drift of *Can God Be Free?* is plain enough: the second account, creating-as-securely-actualising, is the closest to what Rowe has in mind. Nowhere does Rowe discuss, or even acknowledge, the possibility that God's actions sometimes fail to achieve God's aims. Accordingly, in the rest of this paper I consider the Entailment Thesis under the second interpretation.¹⁷

With the expressions 'securely actualise' and 'can securely actualise' in hand, I can introduce some more terminology that will prove useful later—especially in Section V.

A *hierarchy of α -securely actualisable worlds* is a set S of worlds such that each member of S is commensurable with—i.e., either better

than, inferior to, or equal in value to—each other member, and in the actual world α there is some omnipotent and omniscient person G such that G can, relative to α , securely actualise any member of S, and there is no world outside S that is both commensurable with each member of S and which G can, relative to α , securely actualise.

(From now on I will use 'actualisable hierarchy' to abbreviate 'hierarchy of α -securely actualisable worlds', and, ignoring the possibility of G's weakly causing a world to be actual, I will use 'actualise' to abbreviate 'securely actualise'.)

A world W is a *prime member* of an actualisable hierarchy H if and only W is a member of H and there is no better member.

A world is *prime* if and only if it is a prime member of some actualisable hierarchy.

There are three alternatives concerning what actualisable hierarchies there are:

- there is none that has a prime member
- they all have at least one prime member
- some of them have at least one prime member and others do not

The first of these alternatives is equivalent to *There are no prime worlds*.

III. An Attack on the Entailment Thesis

I am now construing the Entailment Thesis as saying that the proposition *There are no prime worlds* entails the proposition *There does not exist an omnipotent, omniscient and perfectly good being*.

If there are no prime worlds, and an omniscient being actualises a world, then she actualises a world when she could have actualised a better one. Rowe thinks that in this case she performs an action that is less good than a relevant alternative action that she can perform, and therefore falls short of being the best possible being.

I deny Rowe's claim. It is logically necessary if there is an omnipotent, omniscient and perfectly good being then she cannot act in a morally better way, or in a way that is better all-things-considered, or more rationally, than she in fact acts. If there are no prime worlds, then what follows is that in every actualisable hierarchy there are infinitely many pairs of worlds V and W such that both W is better than V and an omnipotent and omniscient being can actualise W, but her actualising W need not be a morally better action, or better all-things-considered, or more rational, than her actualising V.

Other things being equal, in intentionally bringing about the better state of affairs one acts in the better way, all-things-considered. Nevertheless there are certain situations in which the inference from *The agent could have selected a better world to be actualised* to *The agent could have acted in a better way* is defeated. One defeater consists of the conjunction of the following three conditions:

- for every world that could have been selected there is a better one that could have been selected, and
- the world that was in fact selected is good enough relative to the foregoing circumstances, and
- failure to select any world would have led to an outcome that was far inferior to each of the worlds that is good enough.

Alongside issues concerning God's choices if there are no prime worlds, there is puzzle concerning rational, self-interested choice. Suppose that a superior being offers to prolong your life for a finite number of good days, with the number of days to be chosen by you, without restriction on how many days you may pick. Rational decision theory can reasonably be required to advise you about what you should do in the foregoing situation.

It is tempting to argue: For any number of days, N , that you might obtain, you would be better off if you obtained $N+1$ instead. So for any N , you have better reason to select $N+1$ instead. So you would be acting more rationally if you selected $N+1$ than if you selected N . So you should not select N . This conclusion, however, is inconsistent with the obvious truth that you should select *some* number, rather than walk away from the offer.

You should satisfice—that is, choose a number that will secure an outcome that is good enough.¹⁸ Indeed, you are rationally required to do so. Satisficing will lead to a better outcome for you than failing to satisfice. Nevertheless satisficing involves choosing an outcome despite the fact that there is an alternative available outcome given which you would be better off.

It might be objected: When a policy is rationally required, and there are various *rationally permissible* ways of implementing it, some of these ways can be more rational than others. Even though it is more rational to satisfice than to fail to satisfice, and therefore there are some numbers that it is rationally permissible for you to choose, for each such number N it remains more rational for you to satisfice by choosing $N+1$ than to satisfice by choosing N .

Decision theory provides demands (including prohibitions), recommendations, and assessments (including comparisons) of rationality. These must fit together harmoniously. Specifically, our decision theory:

- (C1) should not both recommend, of each of several candidate actions, that it be not performed, and also recommend that there be one of these actions that is performed, and
- (C2) should recommend that an action be not performed if the theory identifies some rival action (or refraining from action) as more rational, and
- (C3) should recommend (and, indeed, demand) that an action be performed if the theory identifies the action as rationally required.

The objector violates at least one of these conditions of theoretical coherence. Her theory evidently identifies, for each available choice of the number of days, at least one rival choice as more rational. According to condition (C2), her theory should therefore recommend, for each available

choice of number, that you not make it. But her theory admits that satisficing, i.e., making one of the available choices, is rationally required, and so according to (C3) it should recommend that you make one of them. But these two recommendations jointly violate (C1).

You are rationally required to select some finite number of days. Our decision theory should affirm this. It follows, via (C3), that our decision theory should recommend that you select some finite number of days. No doubt it will add some supplementary recommendations—e.g., that the number selected should be greater than any number of days you have a reasonable expectation of living even if you decline the superior being's offer. Let us call any number that fulfils this and similar clauses 'admissible'. According to (C1), if our theory should recommend that you select some admissible number of days, then it should not recommend of each admissible number, that you not select that number; in which case, we can infer, via (C2) that our theory should not, for every admissible number, identify at least one other number (e.g., the number's successor) as one which it would have been more rational to select.

I conclude that there is at least one admissible number N such that N it is not more rational to select $N+1$ instead, even though, *ex hypothesi*, you will be better off if you select $N+1$ than if you select N .

The foregoing problem has an ethical variant. Suppose that you are a child's guardian, and have a moral obligation to act in her interests. A superior being offers to prolong the child's life for a finite number of good days, with the number of days to be chosen by you. What should you do? An argument similar to the one above leads to the conclusion that, morally speaking, you should sacrifice for the sake of the child.

We should not respond to the problem of how an omnipotent and omniscient being should choose when there are no prime worlds by concluding that *There are no prime worlds* entails *There is no omnipotent, omniscient and perfectly good being*. Drawing that conclusion would provide no way out of the underlying ethical problem, which is raised equally by an who is limited in knowledge and power, and who may or may not be morally perfect, but is faced with an infinite hierarchy of better and better states of affairs that she can decide will be brought about. What moral demands and recommendations does our well-developed moral theory issue to such an agent? What should it say about when her selecting a better state of affairs would be acting in a morally better way than her selecting a less good state of affairs?

It should recommend that she sacrifice, i.e., that she select some good state of affairs even though she could select a better one. Therefore it should not also recommend, of each available good state of affairs, that she not select that one. Therefore it should not declare that, whichever state of affairs she selects, there is at least one alternative member of the hierarchy such that selecting it would be a morally better action.

Some paragraphs ago, I claimed that when an agent is choosing a world to be actualised from an infinite hierarchy of better and better available worlds, the inference to the conclusion that whatever she did she could have acted in a better way is defeated. An objection was stated: that however good the world chosen, the agent would always have chosen better by choosing a better candidate. We are now in a position to identify where the

objection goes wrong. We should resist the inference from 'If the agent had chosen W instead of V then she would have chosen a better world' to 'If the agent had chosen W instead of V then she would have chosen better.'

IV. Standards for Satisficing

In *Can God Be Free?* Rowe responds to an earlier version of the line of reasoning advanced in Section III.¹⁹ He focuses on my original claim, modified above, that in the scenario in which there are no prime worlds the conclusion that God could have acted in a morally better way is defeated by the fact that some worlds are good enough. He applies pressure to the idea that some worlds are good enough.²⁰

Rowe observes that in my view whether or not a particular world is good enough for God to actualise is not an intrinsic feature of the world, but is relative to the range of options God faces. He complains that I seem, wrongly, to think that it is *obvious* just what worlds are acceptable for actualisation, relative to the choice situation in which for every world that could have been selected there is a better one. So I should now say more about this non-obvious matter.

Let us try to get some clues by returning to rational, self-interested choice in response to the superior being's offer to prolong your life. Plainly if you were to choose the number 1 you would be choosing irrationally. You should satisfice by picking a number that is high enough. But when is a number high enough, relative to this choice situation?

The superior being's offer assures you of the opportunity to live a finite, long and good life. In the absence of any further information about your future, what fuzzy minimum length of life will be required to enable you to live a life which, taken as a complete whole, instantiates your values, your ideals, to a high degree? It depends on what these values and ideals are, both before and after you have received the superior being's offer; and it also depends on what you count as a high degree. Let us suppose that your conception of a good life involves a wide range of goods functioning together in a harmonious and well-rounded way. If you die tomorrow then you know that your life will have a regrettable lack of balance: not only will major current aims and hopes be frustrated, and lots of loose ends be left lying around, but, more significantly, some of your most heart-felt values applicable to an individual human life will be unrealised. Perhaps, for example, you think it important that a person does all that she can to overcome estrangement from family and old friends, but you have not yet done so. Presumably there will be some range such that any number of days within the range will fulfill the following condition: although it would be better for you if you lived even longer, a life that falls within the range is not in consequence flawed, deficient or disappointing, relative to your deepest values and ideals for the shape and content of your life. In that case, if you choose a number that is considerably above the bottom of the range—about a thousand times those near the bottom, say—then surely you have satisficed wisely.²¹

Why 'surely'? Because decision theory should not yield the consequence that a wise choice is impossible in the situation supposed. Given this assumption, some such criterion as the one I have sketched is required.

So what counts as a world that is good enough, relative to the task of a person who is to decide which world will be actualised, out of an infinite range of better and better worlds? As I have already pointed out, we need not assume that the chooser is perfectly good. The issue concerns what our moral theory tells her about which of the available worlds are such that she would be choosing in a morally better way if she selected one of the others, and which of the available worlds are not. If there are worlds in the latter group, then they are the ones that are said, in the present context, to be *good enough*.

What answer our moral theory gives will depend on its assumptions about value. Suppose, for instance, that it identifies the basic value states in any world as all consisting in a person's feelings of pleasure or displeasure during a short unit interval of time, measures their value simply by the intensity of the feelings, and declares that the value of a world is the sum of the values of the states. In that case there seems to be no principled way of identifying, amongst the worlds that have positive net value, some that are good enough. But a theory of the foregoing kind is very implausible.

So suppose instead that our moral theory embodies the assumption that worlds are to be ranked chiefly by the extent to which they include many large communities of people characterised by justice, a rich culture, and fellowship between individuals each of whose lives embodies a wide range of goods functioning well together. A theory of this kind is likely to support the judgment that there are worlds such that, whether or not there are even better worlds, they are not flawed, deficient or disappointing relative to the values that underlie the ranking of worlds. (From now, for brevity's sake, on I'll just say '(non-)disappointing'; the other adjectives should be regarded as tacitly present.) Such judgments fit in with a much wider array of judgments of quality—e.g., we can regard a meal cooked by a friend as very good, without inquiring into what a 5-star restaurant chef would have succeeded in doing with similar raw meat and vegetables; we can rate the first performance of a new piano sonata as wonderful, without listening to other performances of it by pianists of higher repute.

How good does a world have to be to be good enough, relative to the choice situation specified? Here is a sufficient condition: when an agent is selecting a world to be actualised from an infinite hierarchy of better and better candidates, a world is good enough if it is non-disappointing in the light of the values that underlie the ranking of worlds, and moreover is abundantly better than those worlds that only just barely escape the accusation that they are disappointing.²²

Can God Be Free? yields grounds for thinking that Rowe would object to the last few paragraphs, along the following lines.²³ Let us compare two hypothetical agents Alice and Beryl, each of whom is faced with the task of selecting a world to be actualised from an infinite hierarchy of better and better candidates, and each of whom adopts the following policy:

I will satisfice, and moreover I will choose in a great-hearted manner, and so select a world that is abundantly better than any world that is just barely non-disappointing.

The difference between Beryl and Alice is that Beryl rules out of further consideration worlds that Alice is quite prepared to regard as living options, because Beryl has a much higher standard than Alice of what is required by great-heartedness, and of what should count as being abundantly better than the minimally and near-minimally non-disappointing.

Surely, the objector declares in support of Rowe, our moral theory should deliver the verdict that in virtue of Beryl's higher standards and the accompanying choice-disposition, Beryl has a better moral character than Alice, and furthermore, given that Alice's choice of a world is classified as morally inferior by what we ourselves describe as the *higher* moral standards of Beryl, we should regard Alice's choice of a world as morally inferior to various alternative choices. The example can be reiterated, since for every possible agent there will be another who has much higher standards pertaining great-heartedness, and of what should count as being abundantly-better-than-barely non-disappointing. Hence for every possible agent faced with the task of selecting a world to be actualised from an infinite hierarchy of better and better candidates, there is another possible agent who faced with a similar task who makes a morally better choice and who has a better moral character. Hence the account given in Section III is mistaken.

Here is some parallel reasoning concerning self-interested choice in the longevity example. Suppose that you have chosen a number of days. Someone now alleges that your standards of what counts as a non-disappointing life, and the dispositions which have led you to hold those standards, have led to a less rational choice of a number of days than alternative standards and dispositions that you might have possessed, and reflect a less rational practical intellect. Indeed, it is alleged, not only has this in fact happened: it had to happen, because whatever specific standards you might have held, there would have been others which it would have been more rational to hold. Hence, even if relative to those standards you actually hold your choice of number of days was at least as rational as any other, your choice was not, *all-thing-considered*, as rational as a choice of this or that larger number would have been.

It is plain that there is an error somewhere here. As I argued back in Section III, the view that for every number there is another number which it is more rational for you to choose should be recognised (in the light of plausible general principles of decision theory) to be inconsistent with the manifest truth that, having received the superior being's offer, practical reason requires you to choose some number. Hence the last sentence of the preceding paragraph, and so the reasoning that led up to it, should be rejected.

But where, specifically, does the reasoning go wrong? Well for one thing, it is far from clear that there is indeed an infinite series of standards of a non-disappointing life that it is more and more rational for you to hold, as the objector in the longevity example assumes. Similarly, in the case of choice between worlds, it is far from clear that there is an infinite series of standards of a minimally good enough world, that are not only higher and higher in the sense of more and more demanding but are also such that, other things being equal, a person who holds a higher standard has a morally better character than a person who adopts a lower standard.

Instead of developing this point further, I will move straight on to the question: relative to the premise that there is indeed such an infinite series of standards, what recommendations do our decision theory and value theory jointly deliver about which standards you should have hold?

They should at least recommend that choosers adopt (or re-affirm) *some* standards of what counts as a non-disappointing life and of what counts as a good enough world, and cultivate *some* relevant dispositions, rather than, say, choosing at random or in accordance with whatever dispositions the choosers happen to find themselves with. If the objector were to suggest the contrary, then she would not be entitled also to make her various moral claims about Alice, Beryl and other agents.

(C1) and (C2), in Section III, make it plain that our decision theory and value theory should not recommend of each available standard that you not hold it, and of each available disposition that you not cultivate it. Hence, given they should not, for every available standard, say that there is at least one alternative which it would be more rational, or better-all-things-considered, to hold.

I can now safely re-affirm my own conclusions about satisficing by divine and non-divine choosers.

V. Hierarchies with Prime Members and Hierarchies without Them

Although my main criticisms of the Entailment Thesis are now complete, it is worth drawing attention to one other consideration that Rowe should find awkward. It becomes apparent in the context of the supposition, which Rowe would find it hard to rule out, that an omnipotent and omniscient agent G is faced with some actualisable hierarchies containing prime members, and some actualisable hierarchies without them.

Let V be a prime member of some actualisable hierarchy, and let W be a good member of some hierarchy that lacks prime worlds. Surely G will be indifferent between actualising V and instead actualising W. After all, V is incommensurable with W.

It might be objected that if G actualises W then she acts in such a way that she has better reason for actualising some other world instead, whereas if she actualises V then this is not the case. Now I agree that in general if some world W^* is better than W then this constitutes a reason for actualising W^* instead of W. But in the context of G's specific overall choice situation, W^* is better than W does not entail G has better reasons, all-things-considered, for actualising W^* ; my grounds for saying this are that it follows from the truth, for which I argued in Section III, that in this context W^* is better than W does not entail *It would be more rational, all-things-considered, to actualise W*. Hence the first premise of the objection is erroneous. G should indeed be indifferent between actualising V and actualising W.

An especially interesting point emerges from the previous paragraph. Rowe is committed to saying that G's being morally perfect is compatible with G's actualising V but is incompatible with G's actualising W. Hence Rowe is committed to the view that G should rank actualising V higher amongst her options than actualising W, even though V and W themselves are incommensurable in value. Plainly, this view is incompatible with maximising act-consequentialism.

Hence given the assumption that some actualisable hierarchies contain prime members and some do not, Rowe cannot invoke the truth of maximising act-consequentialism in any argument he might appeal to in direct support of Principle B, or in support of (2b) or (3a).²⁴ Moreover the view that an agent may rank one candidate action higher than another even though their respective total causal consequences are incommensurable is incompatible with both Rowe's (3a), which constitutes the major premise of Rowe's second argument in direct support of the Entailment Thesis (and also for Principle B). I do not see how Rowe can rule out its being true that some actualisable hierarchies contain prime members and some do not.²⁵

University of Melbourne

NOTES

1. William L. Rowe, *Can God Be Free?* Oxford University Press 2004, p. 92. Rowe uses 'world' to mean 'possible world,' and on pp. 40–41 he proposes that a *world* be thought of not as a whole consisting of created objects such as trees, rivers and minds, but instead as 'a maximal state of affairs *W* such that for any state of affairs either it or its negation is included in *W*.' I concur. 'Creation' therefore pertains not only to *bringing into existence* things like stars, rivers and humans, but also to *providential control* of what happens to things.

2. *Ibid.*, p. 97.

3. *Ibid.*, p. 121

4. *Ibid.*, p. 89, 120f. Consider the proposition that even if for every creatable world there is one that surpasses it in goodness, an omnipotent, omniscient agent's failure to create an unsurpassed creatable world constitutes a fault that shows that the agent herself is not unsurpassably good. Remarks on pp. 104ff indicate that Rowe is not asserting this proposition. Instead, he is saying that an omnipotent, omniscient agent's creating a world when there is a better one that she could create shows that she is not unsurpassably good.

5. Construed uncharitably, (2a) implies that if there is a best possible being then each of her actions—including trivial ones like brushing one's teeth—has an equal degree of moral goodness to each other action. I assume that a more charitable interpretation can be devised. Similarly with respect to (2b).

6. *Ibid.*, p. 100. I am not quoting Rowe; rather, the proposition sums up what he says over several sentences.

7. *Ibid.*, p. 111; cf. p. 92

8. This formulation preserves the idea that being created is being (at least) dependent, and also relies on the point that if a world's fulfilling a specified condition is itself included in the world then it is logically possible that the world fulfils the condition.

9. Clauses (a) and (b) tacitly involve the basic exceptions which I mentioned earlier. I am using the dummy name 'G' instead of the word 'God' in order to avoid assuming that necessarily if a being securely actualises a world then that being is God. That assumption would require support by a lot of metaphysical argument which should be bypassed here.

10. Molinism is a theory of divine knowledge and providence, and their relation to human freewill, explained in Luis de Molina, *On Divine Foreknowledge*, ed. A. J. Freddoso, Cornell University Press 1988

11. Sharpening this account would require elaborate investigation of the identity and role of background circumstances relative to which a specific 'can' statement is true. But a more precise explanation is not needed here.

12. If so, and if Rowe is a libertarian, then Rowe holds that if God exists then Molinism is true.

13. Let action A have two components: strongly actualising the laws of nature, some of which are probabilistic rather than deterministic, and strongly actualising all other relevant states of affairs occurring before F. It is in virtue of the probabilistic laws that F can be said to be *physically probable* given A.

14. This supposition is consistent with there being an infinite series of better and better worlds that G can securely actualise. After all, 2 is greater than any member of the series $1/2, 2/3, 3/4 \dots$, even though the series ascends forever.

15. William Hasker makes essentially this point in his *Providence, Evil and the Openness of God*, Routledge 2004, pp. 104f. Let me expand the argument. If God is timelessly or at all times omniscient, then there is no place for non-metaphorical talk of his deliberating about whether to causally contribute to the occurrence of some event, as if this is somehow an open question. But we can still speak of God's timeless or everlastingly past reasons for action. These reasons for action cannot include the premise that God will in fact causally contribute to the event. So what can they include? Perhaps counterfactual conditionals, concerning what would happen if God were to act in such-and-such ways. But unless either Molinism or determinism is true, the only relevant true counterfactual conditionals are probabilistic ones. One way of employing probability estimates is as inputs to calculations of expected value.

16. Rowe, as shown on pp. 125, 135 of his book, believes that worlds have quantitative values. Even if he is mistaken on this point, it may still in some cases be possible to compare worlds' values in quantitative terms, and thereby to provide for a sense in which the expected value of weakly causing one specific world to be actual can be estimated as greater than the expected values for other worlds. This is most apparent in cases where two worlds differ only in some very simple, measurable, value-relevant respect.

17. Given the second interpretation, we should not believe that Creator is trivially true, on the grounds that to refrain from bringing about anything is in this context to bring about something. Consider, for example, the following proposition: *If God were to refrain from bringing into existence any physical universe, then a physical universe might well come into existence anyway*. If this proposition were true, then God's refraining from bringing into existence any physical objects, finite minds etc., together with any positive actions he did perform, would not amount to his securely actualising some world. It might be objected that the italicised proposition is logically impossible. Even so, it is not a *trivial* truth that it is logically impossible.

18. 'Satisfice' is a technical term from decision theory and economics. I say more about satisficing in Section 4.

19. The earlier version was contained my article 'God and the Best,' *Faith and Philosophy* 13 (1996), pp. 311–28. Rowe's discussion of the article is on pp. 121–27 of his book.

20. Rowe, *op. cit.*, pp. 124–26

21. I have, of course, been assuming that there are no time intervals within which the value of your life decreases as its length increases.

22. On his p. 125, Rowe considers a being who creates a world W100 from an infinite series of increasingly better worlds. Such a being, Rowe says, could instead have created a being a billion times better than W100. Surely there might be a being whose goodness simply would not let him create W100 when there was this superior alternative option. If so, then the being who creates

W100 is not absolutely perfect. The last few paragraphs form the basis for a reply to Rowe. Furthermore, it is at best very risky to assume that if an agent faces an infinite hierarchy of better and better actualisable worlds then in selecting one world then she declines worlds a billion times better than the one she chooses. If there is an infinite hierarchy of better and better worlds, it may nevertheless be that these worlds cannot sensibly be regarded as themselves possessing quantitative values, even fuzzy ones, and if the worlds do possess quantitative values, these values may converge on some finite limit. (Rowe ignores these epistemic possibilities when he says, on his p. 135: 'In the [infinite] series of increasingly better worlds . . . each world logically could be better than a world below it by the *same degree or amount*.'.) In the former case, no world will be a billion times better than any other; in latter case, if the value of the selected world is close to the limit then it may be that no world is a billion times better than it. (Of course in either case some worlds may be a billion places above V in the ranking.)

23. Cf. Rowe's response to Daniel and Frances Howard-Snyder's views on his pp. 92–98, and his principle III* and assumption A3, which he states on his pp. 115, 119 respectively and employs in his discussion of W. J. Wainwright's paper 'Jonathan Edwards, William Rowe and the Necessity of Creation'—Rowe's and Wainwright's pieces are both to be found in Jeff Jordan and Daniel Howard-Snyder, eds., *Faith, Freedom and Rationality*, Rowman & Littlefield 1996.

24. I do not assert that Principle B entails act-consequentialism. Nor do I assert that Rowe himself is a maximising act-consequentialist. After all, he recognises the existence of supererogatory acts (*Can God be Free?* p. 82.) Of course it might be argued that my longevity example provides a strong objection all varieties of consequentialism that rank one action as better than another if and only if its causal consequences are better overall. I need not press this suggestion here, but rest with the point that Rowe will have to work hard before he is in a position to use any such variety of consequentialism to lend plausibility to Principle B.

25. I thank anonymous referees, and the editor of *Faith and Philosophy*, for their helpful comments on an earlier version of this paper.