

THE METHOD OF PERFECT BEING THEOLOGY

Jeff Speaks

Perfect being theology is the attempt to decide questions about the nature of God by employing the Anselmian formula that God is the greatest possible being. One form of perfect being theology—recently defended by Brian Leftow in *God and Necessity*—holds that we can decide between incompatible claims that God is F and that God is not F by asking which claim would confer more greatness on God, and then using the formula that God is the greatest possible being to rule out the one which confers less greatness on God. This paper argues that this form of argument, while intuitively quite plausible, does not work.

Perfect being theology is the attempt to decide questions about the nature of God by employing the Anselmian formula that God is the greatest possible being, which we can state as follows:

[GPB] $\Box \forall x (x \text{ is God iff } x \text{ is the greatest possible being})$

How might this formula help us to decide questions about the divine attributes? A plausible answer to this question, which is employed in many discussions of the divine attributes, is that we can decide between rival incompatible hypotheses about God by asking which hypothesis would make God greater. Given that God is the greatest possible being, it seems, we can at least rule out the hypothesis which would accord God the lower level of greatness.

One finds this line of thought, in various guises, all over contemporary philosophical theology. Murray and Rea, I think, state the orthodox view when they say that the Anselmian formula “provide[s] us with a rule or recipe for developing a more specific conception of God.”¹ The aim of this paper is to raise some questions about the argumentative strategy informally sketched above, and hence about the idea that perfect being

¹Michael Murray and Michael Rea, *Introduction to the Philosophy of Religion* (Cambridge: Cambridge University Press, 2008), 8.



theology provides anything like a satisfactory recipe for answering questions about the divine attributes. In doing so I won't challenge the truth of [GPB]—I think that it is true—but only the idea that it can play the envisaged role in philosophical theology.

The paper has three parts. In the first, I examine the version of this argumentative strategy which has recently been defended at length by Brian Leftow.² Leftow, I think, succeeds in making explicit the method which has been implicit in many recent discussions of the divine attributes. I'll argue that this method fails. In the second section, I'll ask whether Leftow's method can be fixed if we rely, not on the claim that God is the greatest possible being, but on the claim that God is the greatest conceivable being. I'll argue that it cannot. In the last section, I turn to the question of how Leftow's method might be modified by the perfect being theologian to avoid some of these problems. I'll argue that there is, unfortunately, no easy fix.

I. Leftow's Version of Perfect Being Theology

Leftow explains his method in the following passage:

Nothing could be a better G . . . than God in fact is. God can be F. God would be a better G were he F than were he not F . . . Suppose now for reductio that God is not F. Then God is not as good a G as He could be. So if God is not F, it is false that nothing could be a better G than God in fact is. But this is true. So, prima facie, God is F. (9–10)

This, it seems, gives us the wanted recipe for deciding between the hypotheses that God is F and that God is not F—we choose the hypothesis which makes God greater.

There are two complications here worth mentioning before them setting to the side. First, Leftow emphasizes that the conclusion of this chain of reasoning is only prima facie true because it might turn out that the claim that God is F is inconsistent either with the deliverances of revelation or with the claim that God is G, where we have arrived at the claim that God is G via the same style of reasoning used in support of the claim that God is F. Since none of my worries about the style of reasoning under discussion turn on cases of conflicting prima facie claims about the divine attributes, I'll set this sort of worry to the side, and assume that we're dealing with a conflict-free example.

Second, we might distinguish different senses in which a property is good to have, or better to have than lack. Some properties are good for me relative to a kind of which I am a member; for instance, I'm a member of the kind basketball player, and, relative to that kind, it is better for me to shoot baskets every day than not to—that is, I would be a better basketball player were I to shoot baskets every day, than were I not to. But this fact, even given that I am a member of that kind, doesn't entail that this prop-

²Brian Leftow, *God and Necessity* (Oxford: Oxford University Press, 2012). All further references are to this book.

erty is one which it is better for me to have than lack, full stop. In fact, I don't think that it is.

Hence we might ask whether Leftow's argument employs the kind-relative or the non-kind-relative sense of "good." I think that it is difficult to answer this question in a satisfactory way. If we choose the kind-relative notion, then some choices of a kind seem to give the wrong results, whereas others (e.g., the kind deity) seem to presuppose a knowledge of the divine nature that we might have wanted perfect being theology to provide rather than presuppose. If we choose the non-kind relative notion, then we face the problem that different properties are good for me than for, say, a parrot—which makes it seem as though we have to specify an object, or class of objects, in order to get results about which properties are better to have than lack even in the non-kind-relative sense. But again it looks like some classes of objects will yield the wrong results, whereas others (e.g., the choice of the singleton set containing God) will assume the sort of knowledge of God's properties that we want our method to deliver.

While I think that this is a genuine problem for the proponent of perfect being theology, here I will set it to the side. None of the problems to follow depend on any assumptions about the sense of goodness (or betterness) employed in perfect being theology.

Though my focus is on the form of argument characteristic of perfect being theology rather than any particular application of that form of argument, it will be useful to have a specific argument on which to focus. Here I'll focus on Leftow's argument that God is responsible for the "truth-explainers" for all necessary truths, which he labels "NEC":

NEC. (P)(\Box P is true \rightarrow God is, contains, has, has attributes that have (etc.) or produces all its truth-explainers) (115)

Leftow explains in some depth both what he means by a "truth-explainer," and the reasons behind the "is, contains, has, etc." qualification—for our purposes, we can ignore these details and turn directly to his argument in favor of NEC, which looks to be an application of the argument schema sketched above:

It would be an awesome thing to be unconstrained even by modal truths and facts of modal status. . . . A being whose power is not externally limited even by these would seem more powerful than one whose power was subject to such a constraint. (NEC) secures this awesome lack of constraint. So if (NEC) is viable, theists have 'perfect-being' reason to accept it. (134)

At first glance, this looks like a perfectly acceptable form of argument: we consider two incompatible claims about God; we have as a premise that God would be greater if the first claim were true than if the second were; we employ the assumption that God is the greatest possible being; we derive the result that the first claim is true, and the second false.

But let's think in more detail about how this argument is supposed to run. A reasonable first pass at the argument might be:

1. NEC $\square \rightarrow$ God has greatness of level X
 2. \neg NEC $\square \rightarrow$ God has greatness of level Y
 3. $X > Y$
 4. Y is not the greatest possible level of greatness. (1,2,3)
 5. God has the greatest possible level of greatness. (GPB)
 6. \neg (God has greatness of level Y). (4,5)
-
- C. NEC (2,6)

But this argument faces a problem: the step from 1,2 and 3 to 4 is invalid. For this inference to be valid, we'd have to assume that NEC is possibly true. Otherwise, we'd have no way of knowing that X is a possible level of greatness at all, and hence no way of knowing that Y is not the greatest possible level of greatness.

And, looking at the above quotes from Leftow, it might seem that it was his intention to include the possibility of NEC as a premise; that, presumably, is why his sketch of the perfect being theologian's form of argument includes the assumption that "God can be F," and his description of the argument for NEC includes the assumption that "NEC is viable." So suppose we add the premise

0. \diamond NEC

Then 4 follows from 0,1,2 and 3; problem solved.

But the problem is solved only at the cost of creating a new one. The problem is that NEC and its negation are both necessary if possible, and hence that the additional premise 0 is—given that possible necessity is equivalent to necessity—itself equivalent to the conclusion to be argued for. The claim that NEC is necessary if possible is independently plausible; but we can also argue for it by reductio. For consider the claim that NEC is, while true, not necessary. Then there is some world w in which NEC is false, and hence in which God does not explain the modal truths. But then by Leftow's lights God in w has less than the greatest possible level of greatness, which is inconsistent with the plausible thought that God is not just the greatest possible being, but also necessarily the greatest possible being. (More on the generality of this form of argument below.)

Given that NEC is necessary if possible, it looks like the addition of premise 0, while needed to make the argument valid, also obviates the need for the argument. If we knew that NEC was possible, we wouldn't have to appeal to anything about the greatness of God to establish its truth.

One might say, in reply to this problem, that only possible levels of greatness can be compared, so that if premise 3 really is true, then X and Y must be possible, and the move to premise 4 valid. Hence, the reply continues, there is no need to add the problematic premise 0.

But this objection, if correct, simply undercuts Leftow's method in a different way—for if we cannot compare possible with impossible levels of greatness, then we can't use Leftow's method to decide between distinct

hypotheses about God each of which is necessary if possible, for one will always be ascribing an impossible property to God.

The problem here is a very general one for the perfect being theologian making use of a method like Leftow's, for very often hypotheses about God are necessary if possible. Whenever such hypotheses are at issue, the problem just discussed will arise: a "perfect being" argument in favor of one over the other will be valid only if it includes the premise that the hypothesis to be argued for is possibly true; and in such cases that premise will be trivially equivalent to the conclusion.

Below I'll consider some responses to this problem. But rather than trying to rescue Leftow's style of reasoning in its full generality, we might simply narrow its scope, and concede that it can be used only to decide between hypotheses about God which are each possible. This would of course greatly restrict the range of applicability of Leftow's style of argument; but this would still be better than nothing.

It is true that when we are considering a pair of possible claims about God, none of the problems discussed above arise; but another one does, which seems to me just as decisive. Suppose, as many do, that God is not just the greatest possible being, but essentially the greatest possible being, and now consider a pair of claims P , $\neg P$ about God which we stipulate to each be possibly true. For Leftow's sort of "perfect being" argument to work, it must be the case that either

Were P the case, God would be greater than if $\neg P$ were the case.

or the reverse. (Otherwise, the claim that God is the greatest possible being would give us no grounds for choosing between the two.) So suppose that this is true. Given that $\neg P$ is possible, it follows that there is some world w in which God is less good than God in fact is. But then in w God would have less than the greatest possible level of greatness, which contradicts our assumption that God is essentially the greatest possible being, and falsifies [GPB].³

The prospects for Leftow's form of argument, then, look somewhat dim. In cases where the relevant hypotheses are necessary if possible, the relevant argument will be valid only if it contains a premise which is trivially equivalent to the conclusion. In cases where the hypotheses and its negation are each possible, the argument will deliver a result only if we deny that God is necessarily the greatest possible being—a move which hardly seems to be in the spirit of perfect being theology.

³Indeed, this line of reasoning also brings out a further problem for the appeal to conceivability in this context. For many proponents of the idea that God is the greatest conceivable being will also wish to hold that God is the greatest conceivable being in every conceivable world. But then for any pair of claims P and $\neg P$ which meets the above condition, there will be no pair of conceivable worlds, one of which is such that P is true in it while the other is such that $\neg P$ is true in it. For in that case there would be, contra the supposition that God is the greatest conceivable being in every conceivable world, at least one conceivable world in which God is not the greatest conceivable being.

II. Possibility vs. Conceivability

There is a very natural response to this problem, which many readers will have been wanting to make for a while. This is to say that we shouldn't, in the end, rely on the claim that God is the greatest possible being at all, and instead should focus on the claim that God is the greatest conceivable being:

[GCB] $\Box \forall x$ (x is the greatest conceivable being iff x is God)

Isn't this, after all, closer to the original Anselmian claim that God is a being than which no greater can be thought?

Then, with [GCB] in hand, we might replace premises 0 and 5 with the following:⁴

0*. NEC is conceivable.

5*. God has the greatest conceivable level of greatness. (GCB)

The resulting argument will, it seems, be valid without the assumption that both NEC and its negation are possible.

The natural next question is how we should understand "conceivable." And on this topic, one thing is immediately clear: to avoid the problem with the original formulation of the argument, we need some sense of "conceivable" on which some propositions are conceivable but not possible. For were we to choose some sense of "conceivable" which did not meet this constraint, 0* would entail 0. But then 0*, like 0, would trivially entail the conclusion by itself, without any help from (GPB) or (GCB), and the reformulation of the argument in terms of conceivability would fail to solve the problem it was introduced to solve.

This constraint on the relevant sense of "conceivable" suggests that we start with a very weak sort of conceivability, quite distant from possibility, which we might define as follows:

P is conceivable₁ iff $\neg P$ is not a formal logical truth.⁵

On this interpretation, 0* should look quite plausible even to those not antecedently inclined to accept NEC; NEC's negation is simply not a formal logical truth, so it is conceivable₁. This is a good thing. But for just the same reason, on this interpretation, 5* looks quite implausible.

⁴We'd also have to replace "possible" with "conceivable" in premise 4. Since this is irrelevant to the argument to follow, I ignore this complication in the main text.

⁵Here for simplicity I'm being a bit sloppy about the distinction between propositions and sentences. For presumably the things that are conceivable or not are propositions, whereas the things that are formal logical truths are sentences. The defender of the view being sketched in the text, if in possession of a sufficiently fine-grained view of propositions, might resolve the difficulty by saying that, in an extended sense, a proposition is a formal logical truth if it is expressed by at least one sentence which is a formal logical truth. I don't think that for present purposes much hangs on the details here.

Consider, for example, Descartes's question of whether God "was free to make it not true that all the radii of the circle are equal."⁶ It is surely conceivable₁ that God be free to do this, since the negation of that claim is not a formal logical truth. And it is at least plausible that, all things equal, it would be better for God to have this ability than to lack it. But it would surely be a bit quick to conclude on that basis that God's existence requires God to be able to make the radii of a circle unequal. Conceivability₁ is such a weak notion that the requirement that God be the greatest conceivable₁ being is implausibly strong.

It is natural to think that we can get round this sort of problem by strengthening the relevant notion of conceivability; along these lines, one might try

P is conceivable₂ iff \neg P is not a formal logical or mathematical truth, and is not a formal logical consequence of a formal logical or mathematical truth.

This might appear promising, since it looks reasonable to say that both NEC and its negation will be conceivable₂, and one might make a case that the claim

God can make the radii of the circle unequal.

is not conceivable₂, on the grounds that it entails that it is possible that the radii of the circle be unequal, and the negation of this claim is a mathematical truth. Could conceivability₂ be the middle ground between conceivability₁ and possibility that we want?

Unfortunately not. For one thing, it does not even quite handle the case of the radii of the circle, since the negation of the claim that it is possible that the radii of the circle be unequal is not a mathematical truth. And the entailment in question is not a matter of formal logical consequence, since it relies on the (presumably synthetic) claim that ability entails possibility. And we can't simply solve this second problem by replacing "a formal logical consequence" in the definition of conceivability₂ with "an entailment," since—given that every necessary falsehood entails everything—conceivability₂ would then entail possibility, making the detour through conceivability pointless.

And even setting these problems to the side, conceivable₂ is powerless to handle variants of the radii/circle case, like—to use another of Descartes's examples⁷—

God can make a mountain without a valley.

which we should presumably want to come out as not, in the relevant sense, conceivable, but which is clearly conceivable₂.

⁶Letter to Mersenne, 27 May 1630.

⁷Letter to Arnauld, 29 July 1648.

Hence it seems that we need a much stronger sense of conceivability; perhaps something like

P is conceivable₃ iff $\neg P$ is not knowable a priori.

If it is a priori knowable that the radii of a circle must be equal, and knowable a priori that there can be no mountain without a valley (and knowable a priori that ability entails possibility), then something like this appears to handle all of the cases discussed so far.⁸

It handles them, though, by making conceivability quite close to possibility. And one might reasonably wonder—especially in this domain—whether we’ve now made conceivability so close to possibility that the gap between the two can no longer do the work that we want it to. For there is a gap between possibility and conceivability₃ only in the case of necessary a posteriori propositions; but it does not seem especially plausible that claims of the sort we are interested in—like NEC and its negation—are examples of the necessary a posteriori. Indeed, one might even think that the use of conceivability₃ in the present context is a bit self-undermining—for surely Leftow’s argument in favor of NEC is an a priori argument.

But there is another problem, which is that we can simply modify the Cartesian examples above to make trouble for the version of the argument which employs conceivability₃. For suppose that “Hesperus is Phosphorus” and “Water is H₂O” are examples of the necessary a posteriori, and consider the claims

God can make Hesperus distinct from Phosphorus.

God can make water distinct from H₂O.

We presumably want these claims to not come out as, in the relevant sense, conceivable, for just the same reason sketched above: all things being equal, each of these claims would make God greater than would its negation. We don’t want this fact to lead us to conclude that God can bring about impossibilities. But each of these claims is pretty clearly conceivable₃; after all, if it is not knowable a priori that Hesperus is Phosphorus, then it is not knowable a priori that it is impossible for Hesperus to be distinct from Phosphorus, and so not knowable a priori that God lacks the ability to make Hesperus distinct from Phosphorus—which is enough to make the above claims conceivable₃. So if we think that God cannot bring about impossibilities, it looks like we have to deny that God is the greatest conceivable₃ being.

One might respond by denying that claims like these really are examples of the necessary a posteriori—but then it is hard to see how one could reasonably maintain that the claims we are interested in—like NEC and

⁸Presuming the plausible if non-trivial claim that a priori consequences of a priori claims are themselves a priori.

its negation—are. And if they aren't, then for the claims we are interested in, conceivability₃ collapses into possibility.

I have, obviously, not considered every conceivable definition of conceivability, and it may be that one would fare better than those that I have considered. But I hope that the problems into which the above definitions run show the conflicting pressures on any such definition, and make plausible my conclusion that no other definition is likely to do the work that the proposed modification of Leftow's argument would require.

III. How Should the Method of Perfect Being Theology be Revised?

The problems with Leftow's version of the method of perfect being theology result from the fact that, plausibly, there is no property F such that (i) it is possible for God to be both F and not-F, and (ii) it is better for God to be F than not-F (or vice versa). But this problem suggests a solution. Perhaps to decide whether or not to include F among the divine attributes, we should ask, not whether it is better for God to be F than not-F, but whether, in general, it is better for an arbitrary thing to be F than not-F.

Unfortunately, it is harder than it looks to make this thought precise without collapsing back into a version of perfect being theology which shares all the problems of Leftow's. A natural starting point is to say that F is among the divine attributes iff every thing which is F is greater than every thing which is not-F; that is,

$$\forall x \forall y ((Fx \ \& \ \neg Fy) \rightarrow x > y)$$

But this lets every property on the list which is either uninstantiated or universally instantiated, since either situation will make the conditional trivially true by falsifying its antecedent on any assignment of values to the variables. Slightly better is the view that F is a divine attribute iff it satisfies the schema

$$\Box \forall x \forall y ((Fx \ \& \ \neg Fy) \rightarrow x > y)$$

But this faces a closely related problem, which is that it makes any impossible property—as well as any necessarily universally instantiated property—a divine attribute.

This problem might seem easily solvable; we can just restrict our attention to properties which are such that both they and their negation are possibly instantiated by some things and not others. We can do this by letting the divine attributes be those properties which satisfy the following schema:

$$\Diamond \exists x Fx \ \& \ \Diamond \exists x \neg Fx \ \& \ \Box \forall x \forall y ((Fx \ \& \ \neg Fy) \rightarrow x > y)$$

This looks like progress.

But we've been ignoring so far a problem which, at first, looks like a mere detail. The problem is that the above schema does not require that the values of x and y that we compare for greatness be the same with respect to all properties other than F. And this is obviously something that

we should require. For suppose that we are wondering whether omnipotence is among the divine attributes. One might argue, using the above schema, that it is not, by pointing out that a being which is perfectly good and omniscient but not omnipotent is greater than a being which is evil and lacks knowledge but is omnipotent. This is hardly the sort of argument which should rule out omnipotence as a divine attribute!

So we need to build in the requirement that the beings we compare for greatness be the same with respect to every property other than the one to be evaluated. A first pass at this would be:

$$\diamond\exists xFx \ \& \ \diamond\exists x\neg Fx \ \& \ \square\forall x\forall y ((Fx \ \& \ \neg Fy \ \& \ [\forall G:G\neq F](Gx \ \text{iff} \ Gy)) \rightarrow x>y)$$

In plain English, this says: F is a divine attribute if it satisfies the following conditions: possibly, something has it; possibly, something lacks it; and, necessarily, any two things which are alike except that one is F and one is not are such that the first is greater than the second.⁹

But this leads to a problem which I don't see how to solve. The problem is that there are properties which are such that, for anything which possibly instantiates that property, it is better to have that property than to lack it—but still, for all that, not plausible candidates to be divine attributes. Consider, for example, the property of having a well-functioning nervous system. It is plausible that if we compare any two beings alike except for the fact that one has this property and one lacks it, the former will be greater than the latter. This is because having a well-functioning nervous system entails the possession of other properties, like having a nervous system, having a body, etc. Hence if we compare any two beings alike except that one has a well-functioning nervous systems, we will always be comparing a pair of beings with nervous systems, with bodies, and in general alike except for the fact that one's nervous system is, and the others is not, functioning properly. Surely in such a case the former will always be greater than the latter. Hence, according to the test above, having a well-functioning nervous system will come out as one of the divine attributes. But, for all that, we don't think that God has a nervous system, whether well-functioning or not.

It might seem at first as though this problem has an easy solution. The problem, you might think, comes from the fact that the quantifiers in our schema are completely unrestricted. That is what allows us to test for whether a property is a divine attribute by comparing beings—like embodied beings with nervous systems—whose properties are, intuitively, irrelevant to the attributes of God.

But how should the quantifier be restricted? A restriction to, for example, persons won't do the job, since it won't avoid the problem about nervous systems just discussed. We could solve the problem by restricting

⁹Note that this is plausible only as a sufficient condition for a property being a divine attribute. Existence, for example, will fail this test, since it is not the case that $\diamond\exists x\neg(x \text{ exists})$. But presumably we want existence, along with necessary existence, to be a divine attribute.

the quantifiers to range only over deities. But of course once we do that, we've come full circle, since we'll have arrived by a circuitous route at a view of the method of perfect being theology which is extensionally equivalent to Leftow's. With the quantifiers thus restricted, no property which God has essentially will satisfy the first two conjuncts of the schema above. And, if God is essentially the greatest possible being—as [GPB] holds—no property which God possibly has and possibly lacks could make true the last conjunct, since that would entail that God is greater in some worlds than in others. This, of course, is just the problem faced by Leftow in a new guise.

IV. Conclusion

None of this casts any doubt on the truth of [GPB]. But it does cast some doubt on the idea that perfect being theology is a recipe for discerning the divine attributes. For surely if it is such a method, we should be able to give a clear statement of what that method is. But that is just what we have been trying, and failing, to do. To be sure, the failure of the versions of the method of perfect being theology considered above don't amount to anything like an impossibility proof. But I think that they do show that the method of perfect being theology is much more problematic than is usually thought.¹⁰

University of Notre Dame

¹⁰Thanks for helpful discussion to Sam Newlands, Mike Rea, Meghan Sullivan, the participants in the Philosophy of Religion Workshop at Notre Dame, and an audience at the Baylor-Georgetown-Notre Dame philosophy of religion conference in 2013.