

EPISTEMIC HUMILITY, ARGUMENTS FROM EVIL, AND MORAL SKEPTICISM

Daniel Howard-Snyder
Western Washington University

Oxford Studies in the Philosophy of Religion, forthcoming

Many arguments from evil at least tacitly rely on something like the following line of thought:

The Inference. On sustained reflection, we don't see how any reason we know of would justify God in permitting all the evil in the world; therefore, there is no reason that would justify God.¹

The conclusion is frequently more nuanced: "it is very likely that there is no such reason" or "more likely than not" or "more likely than it otherwise would be". Some critics reject the premise: we do see how some reason would justify God. These are the theodacists. Others accept the premise but reject the conclusion: the evidence or non-evidential warrant for God's existence is much better than the evidence for no justifying reason. These are the natural theologians and Reformed epistemologists. Some critics, however, insist that even if the premise is true and even if there isn't better evidence or non-evidential warrant for God's existence, we should not infer that there is no justifying reason.² These are the agnostics about the Inference. In this essay I aim to assess an increasingly popular objection to agnosticism.

There are different versions of agnosticism about the Inference. The one I have in mind—henceforth *Agnosticism* with a capital A—affirms at least two theses:³

Agnostic Thesis 1 (AT1). We should be in doubt about whether the goods we know of constitute a representative sample of all the goods there are.

Agnostic Thesis 2 (AT2). We should be in doubt about whether each good we know of is such that the necessary conditions of its realization we know of are all there are.⁴

¹ Three notes in one. (1) The theistic God is the God in question here. At a minimum, this God has enough power and knowledge to prevent evil, and is unsurpassable in moral goodness. (2) By "evil" I mean anything bad, whether pain or suffering or wickedness in deed or character. (3) Arguments from evil that at least tacitly rely on The Inference can be found in, among other places, Rowe 1978, 1979, 1986, 1988, 1991, 1995, 1996, 2001a, 2001b, and 2006, Russell 1989, 1996a, and 1996b, Russell and Wykstra 1988, Snyder 1990, Tooley 1991 and 2008, Schellenberg 2000, Stone 2003. Draper 1989 does too, *contra* Draper 1992, 2004, forthcoming.

² Defenses of agnosticism include Wykstra 1984 and 1996, Alston 1991 and 1996, Christlieb 1992, Sennett 1991, Plantinga 1988, Draper 1992, Bergmann 2001 and 2009, Bergmann and Rea 2003, Howard-Snyder 1996a, Howard-Snyder and Bergmann 2001a and 2001b.

³ Agnosticism with a capital A resembles Wykstra 1984 and 1996, Alston 1991 and 1996, and Bergmann 2001, although there are subtle but important differences between it and these other versions of agnosticism.

⁴ Two notes in one. (1) I have in mind *broadly logical* necessary conditions, as some would put it, or *metaphysically* necessary conditions, as others would say. (2) I have couched The Inference in terms of reasons and AT1 and AT2 in terms of goods. Does it follow that I assume that any justifying reason appeals only to goods? No. All that follows is that I assume that a justifying reason might appeal to goods.

(I will focus on AT1 although I will say a few words about AT2 shortly.) The Agnostic continues: since we should be in doubt about whether the goods we know of constitute a representative sample of all the goods there are, we should be in doubt about whether some good we don't know of figures in a reason that would justify God. But if we should be in doubt about that, then we should be in doubt about whether there is a reason that would justify God. So we should not infer that there is no such reason, even if we don't see how any reason would justify God and even if there is no evidence and non-evidential warrant for God's existence.

The objection to Agnosticism that I aim to assess is the *Moral Skepticism Objection*, or the *Objection*, for short. There are different versions of the Objection. Here's a simple one. Let *Ashley's suffering* name the evil done to twelve-year old Ashley Jones and what she suffered and lost in Stanwood, Washington, September 21, 1997, who while babysitting her neighbor's kids, was raped and bludgeoned to death by an escapee from a local juvenile detention center. Suppose we could have easily intervened to prevent Ashley's suffering without any cost to ourselves. In that case, it would be absurd to suppose that we should be in doubt about whether we should have intervened. *Obviously* we should have intervened. Agnosticism, however, implies otherwise. It tells us that since we should be in doubt about whether the goods we know of constitute a representative sample of all the goods there are, we should be in doubt about whether there is a reason that would justify God's nonintervention. But if that's right, then so is this: since we should be in doubt about whether the goods we know of constitute a representative sample of all the goods there are, we should also be in doubt about whether there is a reason that would justify *our* nonintervention, in which case we should be in doubt about whether *we* should have intervened. So Agnosticism implies that we should be in doubt about whether we should have intervened. But that's absurd. Obviously we should have. So Agnosticism is false.⁵

It will prove useful to have before us the main thrust of this argument. I will call it *the Simple Version of the Objection*:

1. If Agnosticism is true, then we should be in doubt about whether we should have intervened to prevent Ashley's suffering.
2. We should not be in doubt about whether we should have intervened.
3. So, Agnosticism is false.

Here's the plan of the paper. In section 1, I clarify Agnosticism. In sections 2 and 3, I sketch a criticism of the Simple Version of the Objection. Absent qualification, it goes like this. Our assessment of the Simple Version of the Objection should reflect the epistemic implications of our moral theories or principles. There are two types of moral theory and principle: (i) those that posit right- and wrong-making features of an act that should leave us in doubt about its moral status and (ii) those that posit right- and wrong-making features of an act that should not leave us in doubt about its moral status. If we endorse an instance of the first type, then prior to our assessment of the Simple Version of the Objection we should already be in doubt about whether we should have intervened to prevent Ashley's suffering; in which case we should think that premise (2) is false. However, if we endorse an instance of the second type, then prior to our assessment of the Simple Version of the Objection we should deny that Agnosticism implies that

⁵ Variations on this theme are condoned in Snyder 1990, Russell 1996a and 1996b, Almeida and Oppy 2003 and 2005, Tooley 2004, Pereboom 2005, Jordan 2006, Zagzebski 2007, Piper 2007, Gale 2007, Maitzen 2007. Criticism appears in Alston 1996, Howard-Snyder 1996a, Bergmann 2001 and forthcoming, Trakakis and Nagasawa 2004, Morriston 2004, Bergmann and Rea 2005, Schnall 2007, Trakakis 2007, and Anderson 2005 and unpublished.

we should be in doubt about whether we should have intervened; in which case we should think that premise (1) is false. Either way, the epistemic implications of our moral theories or principles imply that the Simple Version of the Objection is unsound. In section 4, I consider whether my criticism is at home with atheism, on the one hand, and theism, on the other hand. In section 5, I turn to two real live proponents of the Objection and I show how my criticism of the Simple Version of the Objection applies to their version. Some people have responded to my argument with something like the following speech: “In the course of your overall argument, you consider moral theories and principles that are tailor-made to make the argument succeed. However, these theories and principles are massively implausible. If you had considered *my* favored theory and principles, which, of course, are highly plausible, your argument would have been an obvious failure, a failure you could have easily avoided if you had possessed the foresight to consult me.” So be forewarned: you may well respond in a similar fashion. If so, be patient. I reply in section 6.

1. Some preliminaries about Agnosticism

Note, firstly, that Agnosticism is *not* a kind of theism. It is perfectly compatible with atheism. Thus, to call it “skeptical theism,” as many people do, is to evince ineptitude at naming things.

Second, it is important to be clear about what the Agnostic means and does not mean when she says *we don't see* how any reason we know of would justify God in permitting all the evil in the world. She does not mean to comment on our visual capacities. Rather, she means that we don't understand or comprehend how any reason we know of would justify God.

Furthermore, she does not mean that we don't see how any reason would justify God in permitting *any* of the evil in the world, nor does she mean that we don't see how any reason would *partially* justify God in permitting all of the evil in the world. She means that we don't see how any reason would *fully* justify God in permitting *all* of the evil in the world. Ashley's suffering is a case in point.

Sometimes the Agnostic will say “We don't see how any *good* we know of justifies God in permitting all the evil in the world”. This is shorthand. What she means is, “We don't see how any *reason we know of that appeals to a good* justifies God in permitting all the evil in the world”. I will frequently revert to the Agnostic's shorthand way of speaking.

Third, as I noted at the outset, the Inference is frequently nuanced with probabilistic vocabulary. Suppose you hear someone assert, “I don't see how any reason would justify God in permitting Ashley's suffering, so there *probably* is no such reason”. They might simply be expressing their degree of confidence in the proposition that there is no such reason or their tendency to give high odds on it if they were to bet on its truth, in light of their inability to see how any reason we know of would justify God in permitting Ashley's suffering. That's fine. But if that's *all* that they are doing—if probability talk in their mouth in this context is *merely* talk about their degree of confidence or betting tendencies—then the Agnostic is not interested in what they have to say. She isn't interested in their psychology. She's interested in the truth. Only when some more objective, extra-mental understanding of probability talk is invoked or assumed will she care to listen. (Invoked or assumed, for example, by way of the Principal Principle, according to which, roughly, subjective probabilities should track objective probabilities.⁶)

⁶ See Lewis 1980. By the way, the most well-known proponent of the evidential argument from evil, William Rowe, agrees with me on this point. He tends to think of the matter in terms of propensities.

Fourth, it is important to be clear about what the Agnostic means and does not mean when she says that we should be in doubt about whether some good *we don't know of* would justify God in permitting all the evil in the world, including Ashley's suffering.

(a) I might know of something, in one sense, but not in another. I know of String Theory, in the sense that I know that it attempts to unite quantum mechanics and the theory of General Relativity, the most popular versions posit one-dimensional oscillating lines and eleven spatial dimensions, and so on. Anyone can know of String Theory in this sense by simply consulting an encyclopedia. But I don't know of String Theory in another sense, in the sense that would require me to have a substantive understanding of the mathematics of quantum mechanics, the theory of General Relativity, multi-dimensionality, and the like. Which sense does the Agnostic mean when she speaks of *goods we know of* and *goods we don't know of* (if there are any)? She means the second sense. When she speaks of goods we know of, she means goods we comprehend and understand in at least somewhat of a substantive way; and when she speaks of goods we don't know of, she means goods we don't comprehend or understand (if there are any), not even in a somewhat substantive way.

(b) When the Agnostic says that we should be *in doubt* about whether some good we don't know of would justify God in permitting all the evil in the world she does not mean that we should *doubt that* there is such a good. To be in doubt about something is not to doubt that it is so. To doubt that something is so is to be (at least) more inclined to think it is false rather than true; to be in doubt about something is to be of two minds about it, ambivalent, undecided. I am in doubt about whether there is sentient extra-terrestrial life, whether the United States will be a world power in a thousand years, and whether the number of Douglas firs in Lake Padden State Park is odd. But I am not in the least bit inclined to think these things are false. Rather, given what information I have at my disposal, I don't know what to think about them. According to the Agnostic, the same goes for the Inference. Given the information she has at her disposal, she is in doubt about whether there is a reason that would justify God even though she can't see how any reason she knows of would do the trick. She thinks she is in no position to make such a judgment. She is in the dark. She confesses ignorance on the matter.

(c) When the Agnostic says that we *should* be in doubt about whether some good we don't know of would justify God in permitting all the evil in the world she means either that we have a duty to be in doubt about it, that it is wrong not to be in doubt, that we're irresponsible if we are not in doubt, or, alternatively, she means that it is fitting for us to be in doubt about it, being in doubt is the appropriate state of mind.⁷ Speaking for myself, although I do not reject the first way, I tend to think in terms of the second way. I tend to think of the Agnostic as saying that it is proper for us to be in doubt about whether some good we don't know of would justify God, proper in the sense that being in doubt about the matter exhibits a humility that befits the range of our cognitive powers whereas not being in doubt exhibits excessive self-confidence. (Hence the first phrase in the title of this paper.) Perhaps there is an important connection between these two senses. For example, perhaps we have an epistemic duty to possess and exhibit an appropriate humility in our cognitive doings. But whether or not there is any such connection between epistemic duty and epistemic humility will be none of my concern here. I merely register the fact that I think in terms of the latter and not the former, and what I have to say in the sequel should be so understood.

⁷ The duty or fittingness in question is epistemic, as opposed to moral, legal, prudential, etc. Epistemic evaluation is evaluation vis-à-vis the goal of gaining truth or some other epistemically desirable states.

(d) When the Agnostic speaks of a *good*, she means to refer to an abstract state of affairs which, if it were to obtain, would be good. She does not mean to refer to a concrete object or event. Goods are abstracta not concreta.⁸ Thus, when she says that we should be in doubt about whether some unknown *good* would justify God, she means that we should be in doubt about whether some unknown *abstract state of affairs the obtaining of which would be good* would justify God.

(e) The Agnostic assumes that states of affairs are necessary beings. They exist at every possible world. So every good state of affairs exists. However, not every good state of affairs obtains. Some good states of affairs do not obtain. For example, if no one has ever been free with respect to being the sort of person that they are, then the state of affairs of our sometimes being free with respect to the sorts of persons we are *exists* but it does *not obtain*.

An interesting question arises: can a reason that justifies God's permission of suffering appeal to a good state of affairs that never obtains? The received view is that it could not: a reason that justifies God's permission of suffering can appeal to a good state of affairs only if that state of affairs will obtain if it does not do so already.⁹ The Agnostic may well disagree, however. Here's why. Suppose I'm driving down a country road and I see a motorist stranded with a flat tire. She lacks the resources to fix it. I have the resources and time to help. If I don't help, she will suffer the usual frustration of being stranded. Now here's a good state of affairs: I freely help her change the tire. Suppose God permitted her to suffer so that *that* state of affairs would obtain; in that case, his reason appeals to *my freely helping her change the tire*. Suppose that reason justifies his permission of her suffering. Then a reason that justifies God's permission of suffering can appeal to a good state of affairs that does not obtain; for that state affairs never obtains since, as it turns out, I drive on by. This case exhibits how a reason that justifies God's permission of suffering can appeal to a good state of affairs that never obtains.

Two objections routinely arise. I can address them here only briefly.

Objection 1. In this case, the reason that justifies God appeals to a good state of affairs that does obtain, namely *my being free with respect to helping her*, which obtains whether or not I help. I have two points to make by way of reply. First, I doubt that my being free with respect to helping her is a state of affairs that is good *in itself*. If it is, it has very little value, certainly not enough to ground a reason that justifies God in permitting the motorist's suffering. At any rate, it is not nearly as good in itself as my freely helping her, not least because the latter but not the former entails her being helped and my choosing and doing good freely. If, however, my being free with respect to helping her is an *instrumental* good, say, by virtue of being necessary for the obtaining of some other state of affairs that is good in itself, then it isn't clear whether there is an objection here at all. For it is quite natural to suppose that in this case the state of affairs that is good in itself and requires my being free with respect to helping her is my freely helping her, which in the case at hand never obtains. Second, the objection simply changes the case: "in this case, the reason that justifies God appeals to...*my being free with respect to helping her*." Well, no it doesn't. I said it appeals to my freely helping her, *not* my being free with respect to helping her. Of course, the objector will retort: "Well, in *that* case, the reason you attribute to God is *not* a justifying reason since it appeals to your freely helping her and *that* state of affairs never obtains." I have nothing to say in reply except, look very carefully at that last sentence and think of human analogues.

⁸ Cp. Bergmann and Rea 2005, 242.

⁹ E.g. Rowe 1996, 264: "...that good state of affairs must become actual [obtain] at some point in the future, if it is not already actual [does not already obtain]".

Objection 2. God is a perfect knower, so God knows that I will not freely help her, and so the reason that justifies God cannot appeal to my freely helping her. By way of reply, I question the first inference. For, in my view, the extent of a perfect knower's knowledge is adequately captured as follows: x is a perfect knower at time t only if x knows at t each true proposition that can be known at t. Add that at no time t can a perfect knower know at t any proposition that describes the future free acts of human beings (relative to t), and it follows that God can be a perfect knower even if God does not know (prior to my choice) that I will not freely help her. If I'm wrong about the extent of a perfect knower's knowledge, then it is much more difficult for me to see how a reason that justifies God's permission of suffering can appeal to a good state of affairs that never obtains. If I'm right however, then it is much easier.¹⁰

Fifth, we need to understand what the Agnostic means and does not mean when she says we should be in doubt about whether the goods we know of constitute a *representative sample* of all the goods there are (which is AT1).

(a) In general, a sample can be representative of a population with respect to one feature but not another. For example, the employees at Microsoft are representative of the human population with respect to planet of origin but not annual income, place of residence, or nationality, among other things. When the Agnostic says we should be in doubt about whether the goods we know of constitute a representative sample of all the goods there are she means that we should be in doubt about whether the goods we know of are representative of all the goods there are with respect to *being apt for justifying God's permission of all the evil in the world*.¹¹ In what follows, I will typically leave the qualification made in this paragraph tacit.

(b) In general, a sample, S, is representative of a population, P, with respect to feature F, if and only if the frequency of members in S that are F is almost the same as the frequency of members in P that are F. Thus, when the Agnostic says that we should be in doubt about whether the goods we know of constitute a representative sample of all the goods there are with respect to being apt for justifying God's permission of all the evil in the world she means that we should be in doubt about whether the frequency of members of the goods we know of that are apt for justifying God's permission of all the evil in the world is almost the same as the frequency of members of the total population of goods that are apt for justifying God's permission of all the evil in the world.

Sixth, what reason do we have to think that we should be in doubt about whether the goods we know of constitute a representative sample of all the goods there are? Excellent question! The first thing to note about it is that it presupposes that we *need* good evidence to be in doubt about the matter for it to be the case that we should be in doubt. That's arguably false, however. To be in doubt about something is the stance from which we need good evidence to move, to believing it or believing its denial. We don't need good evidence to be in doubt for it to be the case that we should be in doubt. So, absent good evidence to believe that the goods we

¹⁰ Thanks to an anonymous referee for discussion of these matters.

¹¹ Furthermore, although the Agnostic says that we should be in doubt about whether the goods we know of are representative of all the goods there are with respect to being apt for justifying God's permission of all the evil in the world, she does *not* say that we should also be in doubt about whether the goods we know of are representative of all the goods there are with respect to being apt for justifying our permission of those evils. Maybe we should be in doubt about that too; maybe not. It is not part of the Agnostic's position *qua* Agnostic to take a stand on the matter. Cp. Bergmann and Rea 2005, 242, and Almeida and Oppy 2005, 84-85.

know of are representative of all the goods there are (or its denial), we should be in doubt about the matter.¹²

Even if we don't need good evidence to be in doubt about something for it to be the case that we should be in doubt about it, we might nevertheless have good evidence to be in doubt about it. In this connection, the Agnostic argues that evidence to think that the goods we know of are representative of the total population is bad evidence, and that more general considerations in favor of the Inference fail. Moreover, she argues, several considerations properly induce doubt about whether there are God-justifying goods outside our ken.¹³ Since my aim in this essay is to assess a specific objection to Agnosticism and not to assess the general case for it, I will say no more about the latter.

Except this: If the goods we know of constitute a random sample of the total population of goods, then our sample is generated by a process that gives every member of the total population of goods an equal chance of being selected into our sample.¹⁴ But if our sample is generated by a process that gives every member of the total population of goods an equal chance of being selected into our sample *but we lack the concepts needed to comprehend or understand every member of the total population of goods*, then we might have selected from the total population of goods a member that we in fact lack the concepts to comprehend or understand, in which case our sample might have included a good we know of but lack the concepts to understand or comprehend. But that's impossible. It is impossible that we know of a good but lack the concepts to comprehend or understand it.¹⁵ So if our sample is generated by a process that gives every member of the total population of goods an equal chance of being selected into our sample, then we possess the concepts needed to comprehend or understand *every* member of the total population of goods. Thus, if the goods we know of constitute a random sample of the total population of goods, then

- (i) Each member of the total population of goods is such that we possess the concepts needed to comprehend or understand it.

The Agnostic bids us to reflect on the fact that (i) is like other propositions in an epistemically relevant respect, for example

- (ii) Each member of the total population of empirically adequate physical theories is such that we possess the concepts needed to comprehend or understand it,

¹² Two notes in one. (1) If you think that you need good evidence to be in doubt about something in order for it to be the case that you should be in doubt about it, then let its introspectively seeming to you that you lack good evidence to believe it or its denial be the good evidence in question, and adjust the text accordingly. (2) If you think that on occasion it's alright to move from being in doubt about something to believing it or its denial even though you don't have good evidence to do so, then do one of two things: either let "good evidence" include experiences of various sorts or else grant that the case at hand—moving from doubt about whether the goods we know of constitute a representative sample of all the goods there are to believing it or its denial—is not one of those occasions. Thanks to Hud Hudson and an anonymous referee observations that led to these two points.

¹³ See, e.g., Wyktra 1984 and 1996, Alston 1991 and 1996, Howard-Snyder 1996a, Bergmann 2001.

¹⁴ ...or else there is a subclass, C, of the total population of goods, P, such that C is not S and C is representative of P with respect to being a would-be God-justifier, from which S is generated by a process that gives every member of C an equal chance of being selected into S. See Hawthorne 2004, note 15. This way for a sample to be random need not concern us here.

¹⁵ Recall the brief discussion of what it is to "know of a good" on page xxx above.

and

(iii) Each member of the total population of ontologies of what we call “physical objects” is such that we have the concepts needed to comprehend or understand it.

Even if, unbeknownst to us, (ii) and (iii) are both true, we should be in doubt about them. It would be an extraordinary stroke of good epistemic luck if our evolutionary history to this point left us with every concept needed to comprehend and understand every physical theory and ontology of physical objects. Similarly, says the Agnostic, and for the same reason, even if, unbeknownst to us, it turned out that (i) is true, we should be in doubt about it. So we have some good evidence to think that we should be in doubt about whether the goods we know of constitute a random sample of the total population of goods, and so we have some good evidence to think that we should be in doubt about whether the goods we know of constitute a representative sample of all the goods there are, which is exactly what AT1 says.¹⁶

Finally, a word about AT2, the thesis that we should be in doubt about whether each good we know of is such that the conditions of its realization we know of are all there are. As has been pointed out on occasion, we can *know of* a good without *seeing how* it would justify God in permitting horrific evil. This can happen in at least two ways. First, we might know of a good but fail to fully appreciate its goodness. Second, we might know of a good but fail to know of all the conditions of its realization.¹⁷ AT2 is about the second way.

What goods might be such that we know of them but fail to know of all the conditions of their realization? Union with God is one candidate. It’s hard to say whether or not created persons must be permitted to undergo horrific suffering in order to enter into the deepest union with God. To be sure, we have some idea of what it would require by way of understanding what union between human persons requires. But is our understanding of what union with God requires *complete*? Suppose we are in the following frame of mind. No aspect of God’s nature that we know of is such that we think that by virtue of it, God cannot permit horrific suffering; moreover, for all we can tell, there are aspects of God’s nature that we don’t know of in virtue of which a created person can enter into the deepest union with God only if she is permitted to undergo horrific suffering. If we are in that frame of mind, then our understanding of what union with God required would be not only be incomplete, it would be—much more importantly—*incomplete* in such a way that we should be in doubt about whether we know of all the conditions of its realization. Are we in that frame of mind? I think I am. The Agnostic thinks you should be too. If she’s right, then we should be in doubt about whether union with God is such that the conditions of its realization we know of are all there are. Thus, we should think AT2 is true.¹⁸

So much for preliminaries. I now turn to my main task.

2. Taking consequences very, very seriously

Consider the following theory:

¹⁶ Thanks to Joshua Spencer for discussing this argument with me.

¹⁷ See Alston 1996, 315-16, 323-25, Howard-Snyder 1996a, 308 n13, and Bergmann 2001, xxx.

¹⁸ According to J.L. Schellenberg, *there is some* aspect of God’s nature we know of such that we *should* think that in virtue of it God cannot permit horrific suffering. That aspect is God’s perfect love: perfect love is maximally empathetic, maximal empathy is maximally opposed to horrific suffering, and maximal opposition to horrific suffering cannot permit it. See Schellenberg 2007, chapter 11. For criticism, see Howard-Snyder unpublished.

Objective Maximizing Act Consequentialism (OMAC). An agent's act is permissible solely in virtue of the fact that its total consequences are no overall worse than those of any option open to him; otherwise, it is impermissible.

There are different concepts of consequence that might be plugged into OMAC and the resulting versions of OMAC will have different implications. Although I'm not defending OMAC here, as will be apparent shortly I will assert that OMAC has certain epistemic implications. Whether I'm right or not will depend on what concept of consequence is deployed. So I have to say something.

Without trying to be precise, I have in mind a version of OMAC that most of my self-identifying maximizing consequentialist friends affirm. They say that what counts as a consequence of an act is *any* future event or fact causally downstream from the act. Some of them like a counterfactual condition on causation according to which A caused B only if B would not have occurred if A hadn't. In that case, we can think of a chain of counterfactuals of the form *If A had not occurred then B would not have occurred* linking the act and the future event or fact. If you are happy with this concept of causal consequence, go with it. If not, go with whatever "link" you like *provided that* it has the implication that my maximizing friends want, namely, that what you do right now will have causal ramifications until the end of time and *all* of them are morally relevant.¹⁹

OMAC implies that the total consequences of intervening to prevent suffering and the total consequences of nonintervention make the difference as to whether we should intervene. If nonintervention in Ashley's case has overall better total consequences than intervention, then we should not intervene—even if the foreseeable consequences of intervention are vastly better than those of nonintervention. Now: since we are in no position to say what the unforeseeable consequences of intervention and nonintervention contain and the unforeseeable consequences swamp the foreseeable ones, we should be in doubt about whether the total consequences of intervention are overall worse than those of nonintervention. In that case, given OMAC, we should be in doubt about whether we should intervene to prevent Ashley's suffering. Thus, prior to assessing the Simple Version of the Objection, we should already be inclined to deny premise (2)—*if we endorse OMAC*.

In effect, I have just summarized the first step of the well-known Epistemic Objection to OMAC. Later steps connect that step with the denial of OMAC. I am not taking the later steps. I am only taking the first.²⁰ Let me explain why I take the first step. What I have to say is not original.

The unforeseeable consequences of an act and its alternatives swamp the foreseeable consequences. Thus, what we can foresee is a minute fraction of the total consequences. Moreover, we are in the dark about what the unforeseeable consequences of an act and its alternatives contain. In an important article, James Lenman underscores how deeply darkness envelops us on this score by pointing out how much of our behavior has massive and inscrutable

¹⁹ Cp. Mason 2004, 317: "consequentialism demands that we make decisions that have as their justification *the whole future*" (emphasis added). Unless otherwise indicated explicitly or by context, in what follows all talk of consequences should be understood along these lines including talk of consequences in contexts other than OMAC. Thanks to an anonymous referee for pressing me on this matter.

²⁰ Later steps include, e.g., the claim that an ethical theory must be practical or action-guiding, or better, the claim that ought implies can. On the better step, see Howard-Snyder 1997.

causal ramifications.²¹ Killing and engendering, and refraining from killing and engendering, ramify in massive ways because they are directly identity-affecting actions. They directly “make a difference to the identities of future persons [that is, a difference to what people there will be] and these differences are apt to amplify exponentially down the generations.”²² Much of our other behavior is indirectly identity-affecting, as, for example, when a word harshly spoken, or eating raw garlic, or introducing your girlfriend to your best friend Ray makes a difference to who sleeps with whom tonight, or tomorrow morning, or next month. To illustrate the main point here, imagine Richard, a first century BC bandit in southern Germany who, while raiding a small village, spares the life of a pregnant woman, Angie.²³ Angie, it turns out, is the great-great-...[add 97 ‘great-’s]...great-grandmother of Adolph Hitler. By permitting Angie to live, Richard played a role in the occurrence of the Holocaust. Moreover, anyone who refrained from killing any of the intermediate ancestors of Hitler before they engendered the relevant child, or assisted in introducing the parents of each generation, or refrained from introducing them to others, and so on, played a role as well. Which one of these people throughout the generations had an inkling that their behavior would contribute to such a horror?

Another source of massive causal ramification is causal systems that “are extremely sensitive to very small and localized variations or changes in their initial conditions”.²⁴ Such sensitivity will underscore the skeptical implications of OMAC if such systems occur

in even a small number of domains that have a significant influence on the human world. One such domain is perhaps the weather: differences in the weather make extremely widespread differences to the behavior of huge numbers of people. Such differences affect, for example, people’s moods, the plans they make for any given day, and the way these plans evolve as the day goes on. For any significant difference in weather over a large populated area, some of these effects are certain to be identity-affecting.²⁵

Another such domain is financial markets:

[T]hese are influenced by countless, often quite intrinsically insignificant, human actions, and probably—directly or indirectly—by a very high percentage of intrinsically more significant ones. And the effect of market movements on human life is again enormous and certainly often identity-affecting.²⁶

As it was with Angie and Richard, so it is with Ashley and us. We are in the dark about the unforeseeable consequences of intervention and nonintervention; moreover, the foreseeable consequences are but a drop in the ocean of the total consequences, and all but that drop is inscrutable to us. So, *if we endorse OMAC*, then, when we turn to assess the Simple Version of the Objection, we should already be in doubt about whether we should prevent Ashley’s suffering; that is, we should already be strongly inclined to deny premise (2).²⁷

²¹ Lenman 2000. I can’t recommend this article strongly enough.

²² Lenman 2000, 346.

²³ Lenman 2000, 344-46.

²⁴ Lenman 2000, 347.

²⁵ Lenman 2000, 348.

²⁶ Lenman 2000, 348.

²⁷ This argument only assumes that causation is transitive in the sorts of causal series in question.

OMAC posits right- and wrong-making features of an act which, given our cognitive situation and a sensible view about what is and is not of value, should leave us in doubt about its moral status. But perhaps appearances are deceiving. Perhaps there is a way friends of OMAC can avoid this skeptical implication. Let's look into the matter briefly.

A popular reply used to be that the consequences of our acts "approximate rapidly to zero like the furthest ripples on a pond after a stone has been dropped in it".²⁸ Or, as a referee crisply put it: "Consequences fizzle fast". To be sure, there are some concepts of a consequence according to which this is true. But, as I intimated above, those concepts are none of my concern here. I am concerned with a version of OMAC according to which, as I said, "what you do right now will have causal ramifications until the end of time and *all* of them are morally relevant". We all know objective maximizing consequentialists who take this line. And my point is simply this: no one privy to the facts to which Lenman calls our attention can retain the view that consequences *in that sense* fizzle fast.²⁹

An appeal to expected value might be a more promising strategy. Suppose act *A* is an alternative action open to me. There are many possible outcomes, $O_1, O_2, O_3, \dots, O_n$, each of which might obtain, for all I can tell, if I were to perform *A*. Each outcome has a value: $V(O_i)$. Moreover, for each outcome, O_i , there is a conditional probability of its obtaining given that I perform *A*: $Pr(O_i/A)$. The expected value of *A* is the sum, for all of these outcomes, of all of the products determined by $V(O_i) \times Pr(O_i/A)$. Expected value can be put to use as follows. Although the foreseeable consequences of intervention and nonintervention in Ashley's case are a vanishingly small proportion of their total consequences, and although we are ignorant of their unforeseeable consequences, it does not follow from OMAC that we should be in doubt about whether we should intervene. For, despite our vast ignorance, we should not be in doubt about whether the expected value of intervention is greater than the expected value of nonintervention: we should think it is greater. Thus, we should not be in doubt about what we should do: we should intervene.

But *why* should we not be in doubt about whether the expected value of intervention is greater than the expected value of nonintervention? To answer that question, we need to answer two others. First, what general procedure should we follow to determine whether the expected value of an act is greater than the expected value of available alternative acts? Second, if we follow that procedure in Ashley's case, will it leave us in a position where we should not be in doubt about intervening?

In an important article, Fred Feldman contends that the nature of expected value itself recommends the following general procedure:³⁰

1. List all of the alternative actions available to us.
2. List all of the possible outcomes of the first alternative.
3. For each outcome of the first alternative, specify its value.
4. For each outcome of the first alternative, specify its probability on that alternative, given the information available to us.
5. For each outcome of the first alternative, multiply its value by its probability on that alternative.
6. Sum these products. This sum is the expected value of the first alternative.

²⁸ Smart and Williams 1973, 33. See also Moore 1903, 153.

²⁹ Cp. Lenman 2000, 350-51.

³⁰ Feldman 2006; it's a must-read.

7. Repeat steps 2-6 for each of the other alternatives.

Let's apply this procedure to Ashley's case.

Step 1 tells us to list the alternative actions available to us. What are they? At first blush, there are exactly two options: intervention and nonintervention. But that's a gross oversimplification. The fact is that there are a thousand (tens of thousands? millions? more?) ways in which we can intervene and many more ways in which we can fail to intervene. Each of them must be placed on our list. To the extent that we are in doubt about whether our list is complete, we should be in doubt about the results we arrive at.

Suppose we somehow identify a few of the most salient alternative actions. Step 2 tells us to list the possible outcomes of the first act on our list. Suppose the first act is firing a warning shot in the air to scare away the perpetrator. Recall that an outcome of an act is a total way the world might go if the act were performed. And note that the "might" in question is epistemic. We need to ask: how many total ways might the world go if we were to fire a warning shot, relative to the information at our disposal? There are millions of such ways, perhaps many, many more. We need to list each of them. To the extent that we should be in doubt about whether our list is complete, we should be in doubt about our results.

Suppose we somehow identify several of the most salient outcomes, say, a thousand of them. (To the extent that we lack a principled way to do this, more grounds for doubt arise.) The next two steps tell us that we need to assign numbers to those outcomes. Step 3 tells us to assign a number that represents the true value of each outcome. Step 4 tells us to assign a number that represents the probability of each outcome, given our firing a warning shot. (For each alternative act, the sum of the probabilities assigned to each outcome must equal exactly 1.) We haven't the foggiest idea what numbers to assign. We are awash in a sea of doubt.

Suppose we somehow assign the correct numbers. To arrive at the expected value of our first alternative action, we must multiply value and probability one thousand times, once for each outcome (step 5). Then we must add the products (step 6). By the time we finish this last step, Ashley's fate will have been long decided. Of course, even if, by some miracle, we arrive at this point in a second or two, we must now repeat the procedure for each of the salient alternative actions available to us, of which there are many (step 7). Our work has just begun.

An ethicist from one of the top three Philosophy departments on the Leiter Report declared with utmost sincerity before an audience of sixty philosophers at the 2004 Bellingham Summer Philosophy Conference that she "just intuitively" which act has the highest expected value without doing the calculations. While such declarations make for memorable theatre, it is difficult to take them seriously. At any rate, I suspect that most of us lack the powers of intuition she claimed to possess.

So: should we intervene to prevent Ashley's suffering, given OMAC? If the appeal to expected value is all we have to go on, we might as well toss a coin.

Perhaps friends of OMAC can avoid these worries by appealing to the Principle of Indifference which, for our purposes, we can put like this:

Indifference. If we have no evidence favoring any of n mutually exclusive possibilities, we should assign each a probability of $1/n$.

Indifference might be put to use in Ashley's case as follows. It is virtually certain that she will be saved if we intervene and it is virtually certain that she will not be saved if we do not intervene.

Those are the foreseeable consequences of intervention and nonintervention. We are in the dark about the unforeseeable consequences of intervention and nonintervention, however. For example, we have no evidence to suppose that the unforeseeable consequences of nonintervention will not be much better than the unforeseeable consequences of intervention. This fact drives the Agnostic's worry. The corrective is to remember that, by the same token, we have no evidence to suppose that the unforeseeable consequences of intervention will not be much better than the unforeseeable consequences of nonintervention. Thus, since we have no evidence favoring one of these two mutually exclusive possibilities, Indifference tells us to assign each a probability of $\frac{1}{2}$, in which case they cancel each other out. So we are left with the foreseeable consequences. On that score, there is no doubt about what we should do—we should intervene.³¹

I have three concerns about the appeal to Indifference here.

First, we have no good reason to believe it. Our grounds for believing it are either a priori or empirical. I haven't the space to consider all a priori grounds that have been offered. Here's the most recent attempt I know of:

Let's say that possibilities n_1 and n_2 are *evidentially symmetrical* for you if and only if you have no more evidence to think that n_1 is the case than you have to think that n_2 is the case, or vice versa. Now, when two possibilities are evidentially symmetrical for you, you should assign a probability to them that adequately reflects your evidence for them. Thus,

Evidential Symmetry. If n_1 and n_2 are evidentially symmetrical for you, then you should assign exactly the same probability to n_1 that you assign to n_2 .

An obvious corollary of this principle is this:

Indifference. If you have no evidence favoring any number n of mutually exclusive possibilities, then you should assign each a probability of $1/n$.³²

What should we make of this argument?

I have two objections. First, it is *not* obvious that Indifference is a corollary of Evidential Symmetry. Indeed, it is obvious that Indifference is *not* a corollary of Evidential Symmetry. For even if you should assign exactly the same probability to n_1 that you assign to n_2 when your evidence for them is symmetrical, it does not follow that you should assign $1/n$. You might well assign a vague or indeterminate probability, perhaps even the interval $[0,1]$ in some circumstances, to each of them. In fact, this way of representing the probability of possibilities under complete ignorance is a much more accurate representation of that cognitive condition than is assigning a sharp probability to each of them. Second, suppose n_1 and n_2 are evidentially symmetrical for you. Does it follow that you should assign each of them a probability? Of course not. You should assign no probability at all. A fortiori, you should not assign exactly the same probability, contrary to Evidential Symmetry.³³

Independent a priori grounds to believe Indifference have been notoriously difficult to find, which has led some of its friends to say that it is a part of the concept of rationality itself.³⁴

³¹ See Keynes 1913, 309-10, Norcross 1990, Kagan 1998, 64ff.

³² I have gleaned this argument from White forthcoming (ms). Whether he'd endorse it, I don't know.

³³ The idea isn't new, of course. See, e.g., van Fraassen 1990 and Levi 1974. This is not the place to assess the objections to vague or indeterminate probability put forward in White forthcoming.

³⁴ A recent example: Mason 2004, 319.

This is a desperate move. It implies that “Dennis is no less rational for denying Indifference” expresses a *conceptual* impossibility. It does not. Moreover, it implies that its foes have a deficient grasp of the concept of rationality. They do not. Other friends of Indifference say it *just seems* true to them. When faced with autobiographical declarations like this, about all one can say is that it *doesn’t* just seem true to me.

Perhaps we can do better with empirical grounds for Indifference, at least insofar as it applies to our present concern. Suppose we have empirical grounds to think that where we find cases of massive causal ramification in the human sphere, the total good and bad consequences tend to cancel each other out in the long run. I concur with James Lenman’s assessment of this suggestion:

There are no cases of massive causal ramification of the kind to which identify-affecting actions are liable where we have empirical data adequate to any such conclusion, for the simple reason that, even if such ramification were easy to trace (in fact it is quite impossible), there are no such cases in which we have good grounds to suppose the ramification has yet come close to running its course.³⁵

But one might object.³⁶ We have empirical grounds to think that

P. All *observed* cases of massive causal ramification are such that the good and the bad consequences tend to cancel each other out.

In that case, what would be wrong with a straightforward enumerative induction to the conclusion that

C. All cases of massive causal ramification, including future cases, are such that the good and the bad consequences tend to cancel each other out?

With this conclusion in hand, we can reasonably ignore the unforeseeable consequences of intervention and nonintervention in Ashley’s case and focus on the foreseeable consequences as a basis for reasonably believing we should intervene.

What’s wrong with this inductive inference is that it is reasonable to make it only if it is reasonable to suppose that the observed cases of massive causal ramification constitute a representative sample of the total population of massive causal ramifications. But is it reasonable to suppose such a thing? If we had good reason to think that the observed cases were randomly selected from the total population, then we’d have good reason to suppose that they constitute a representative sample of the total population. But we *know* that’s not the case: we know that the observed cases are *not* selected in such a way that gives every member of the total population of massive causal ramifications an equal chance of being in our sample. So why suppose they are representative of the total population? (Perhaps you “just intuit” it? Perhaps it “just seems” representative to you?)³⁷

I see no a priori or empirical grounds to believe Indifference.

³⁵ Lenman 2000, 354.

³⁶ Thanks to Rob Epperson and an anonymous referee for suggesting this argument.

³⁷ I leave as homework for the reader what to make of the further suggestion that an abductive inference would work well here.

My third concern with the appeal to Indifference has to do with objections to Indifference itself. One is *Bertrand's Paradox*, a version of which is as follows.³⁸ Imagine a factory that randomly produces square tiles of different lengths, ranging anywhere from 0 to 10 centimeters.³⁹ What is the probability that the next tile to come out of the factory will have sides measuring 5 cm or less? The possible outcomes in this case correspond to all of the lengths from 0 cm to 10 cm—the next tile could have sides measuring 1 cm, or 4.5 cm, or 8 cm or 9.87654321 cm or... There are many possibilities here but, on the face of it, half of these possibilities are ones in which the sides are 5 cm or less since 5 is halfway between 0 and 10. Since we have no evidence favoring any of these outcomes, Indifference tells us that each is equally likely, so *the probability that the next tile to come out of the factory will have sides measuring 5 cm or less is 1/2*. Here is a second question: What is the probability that the surface area of the next tile will be 25 cm² or less? Well, if all of the tiles are squares, and if the lengths range from 0 cm to 10 cm, then the surface areas will range from 0 cm² to 100 cm². The possible outcomes in this case correspond to all of the different surface areas—the next tile could have a surface area of 1 cm² or 26 cm² or 62 cm² or 99.999 cm². Again, there are many possibilities here but, on the face of it, a quarter of them are ones in which the surface area is 25 cm² or less, since 25 is a quarter of the way between 0 and 100. Since we have no evidence favoring any of these outcomes, Indifference tells us that each is equally likely, so *the probability that the next tile to come out of the factory will have a surface area of 25 cm² or less is 1/4*. Here is the problem: *The next tile to come out of the factory will have a surface area of 25 cm² or less if and only if that tile has a length of 5 cm or less*. So the probability that the surface area will be 25 cm² or less *just is* the probability that the length will be 5 cm or less. In other words, $1/4 = 1/2$. Indifference leads to absurdity.

We might try to save the appeal to Indifference by modifying it. How we modify it will depend on what we see as problematic in cases like Bertrand's Paradox. Perhaps we think there is a nonarbitrary way to specify the possibilities. For example, perhaps we think that side length is more "natural" than surface area, say because a tile has its surface area in virtue of its side length and not conversely. Thus we might insist that if we have no evidence for any of n mutually exclusive possibilities, *and one of the ways of specifying the possibilities is more natural than the others*, we should assign them each a probability of $1/n$, under that specification. Or perhaps we think that, unlike the tile case, there are finitely many possibilities in Ashley's case. Thus we might insist that if we have no evidence supporting any of n mutually exclusive possibilities, we should assign them each a probability of $1/n$, *provided n is finite*. This is not the place to investigate these and other options. For my own part, I have yet to meet an *unobjectionable* restricted principle—including these two—that is *clearly applicable* to Ashley's case.

Bertrand's Paradox and other objections don't highlight the main problem with Indifference, namely that it codifies a way to get detailed information out of complete ignorance. Better that we assign vague or indeterminate probabilities, even the interval $[0,1]$ if needs be, or that we refrain from assigning any probabilities at all.

The upshot is this: if the appeal to Indifference is all we have to go on, we should be in doubt about whether we should intervene, given OMAC.⁴⁰

³⁸ This example is adapted from van Fraassen 1999, 303, paraphrased from Howard-Snyder, Howard-Snyder, and Wasserman 2009, 546-47.

³⁹ The numbers here are chosen to simplify the math—obviously there is a lower limit to how small we can make tiles.

⁴⁰ Thanks to Christian Lee and Ryan Wasserman for discussion of these points.

Perhaps friends of OMAC can ditch Indifference and argue as follows. Given the inscrutability of the distant future, in Ashley's case we have exactly as much reason to believe

A. The *unforeseeable* consequences of nonintervention outweigh the unforeseeable consequences of intervention

as we have to believe

B. The *unforeseeable* consequences of intervention outweigh the unforeseeable consequences of nonintervention.

Thus, we should—epistemically should—base our belief about what we should do on what reasons we have to believe

C. The foreseeable consequences of nonintervention outweigh the foreseeable consequences of intervention

and what reasons we have to believe

D. The foreseeable consequences of intervention outweigh the foreseeable consequences of nonintervention.

We have much more reason to believe D than we have to believe C. Thus, we should *not* be in doubt about whether we should intervene, given OMAC. Thus, OMAC is at home with premise 2 of the Objection.

On the face of it, this is a sensible line of thought.⁴¹ However, we can construct another argument for the opposite conclusion that, on the face of it, is equally sensible. Given that intervention and nonintervention have massive and inscrutable causal ramifications and given that the unforeseeable consequences swamp the foreseeable ones, in Ashley's case we have exactly as much reason to believe

E. The *total* consequences of nonintervention outweigh the total consequences of intervention

as we have to believe

F. The *total* consequences of intervention outweigh the total consequences of nonintervention.

Thus, we should—epistemically should—be in doubt about whether we should intervene, given OMAC. So, OMAC is *not* at home with premise 2 of the Objection.

Which of the two arguments is more sensible? As expected, I give the nod to the second. Here's why. I grant that, given the inscrutability of the distant future, we have exactly as much reason to believe A as we do B. But I deny that it follows that we should—epistemically should—base our belief about what we should do on what reasons we have to believe C and D.

⁴¹ Thanks to Dennis Whitcomb for suggesting it to me.

For that follows only if we have *more* reason to believe F than we have to believe E, despite the fact that intervention and nonintervention have massive and inscrutable causal ramifications and the unforeseeable consequences swamp the foreseeable ones. But given this fact, we do not have more reason to believe F than we have to believe E. To suppose otherwise is like supposing that we have more reason to believe that the consequences of intervening to prevent the execution of Socrates outweigh the consequences of nonintervention than we have to believe that the consequences of not intervening to prevent the execution of Socrates outweigh the consequences of intervention. It is like supposing that we have more reason to believe that the consequences of intervening to prevent Brutus's assassination of Caesar outweigh the consequences of nonintervention than we have to believe that the consequences of not intervening to prevent the assassination of Caesar outweigh the consequences of intervention. We are in no position to suppose such things.⁴²

Much more needs to be said about these matters. Suffice to say, as I see things, OMAC posits right- and wrong-making features of an act that should leave us in doubt about its moral status.

OMAC is not alone on this score. Consider

Objective Rossianism (OR). An agent's act is permissible solely in virtue of the fact that it has no less on balance *prima facie* rightness than any option open to him; otherwise it is impermissible.⁴³

Suppose that, as Ross said, one *prima facie* duty, the duty of Beneficence, is the duty to better other people, e.g. to help them achieve a greater degree of virtue, intelligence, pleasure, etc., and another is the duty of Non-Maleficence, the duty not to injure others. Now, for any act, whether it constitutes bettering or injuring others depends on its total consequences. (I injure the child whose leg is destroyed by a mine I planted fifty years earlier. You benefit me when the advice you gave me thirty years ago pays off.) So whether, on some particular occasion, intervening to prevent suffering has more on balance *prima facie* rightness than nonintervention depends on the total consequences of both. Or consider the moral imperative

Requirement R_o. Prevent suffering you can, unless *there is* better reason for you not to intervene.

Read *R_o* so that the total consequences of intervention and nonintervention are good reason not to intervene, even if we are ignorant of them. Then, whether or not on some particular occasion one's intervention violates *R_o* depends on the total consequences of intervention and nonintervention. Like OMAC, OR and *R_o* posit right- and wrong-making features of an act which that should leave us in doubt about its moral status. Or so I have argued.

Let us say that any moral theory or principle that, like these three, posits right- and wrong-making features of an act that that should leave us in doubt about its moral status an instance of *Moral Inaccessibilism*, or *Inaccessibilism* for short. I contend that, if we endorse an

⁴² Thanks to Frances Howard-Snyder for discussion of this argument.

⁴³ In order to deal with conflicting *prima facie* duties, Ross says that our duty *simpliciter* is that which has "the greatest balance of *prima facie* rightness"; see Ross 1930, 41.

instance of Inaccessibilism, we should be in doubt about whether we should have intervened to prevent Ashley's suffering, i.e. we should reject premise 2 of the Objection.⁴⁴

3. Taking consequences much, *much* less seriously

Suppose we want to avoid the skeptical implications of Moral Inaccessibilism. In that case, we might convert OMAC into

Subjective Maximizing Act Consequentialism (SMAC). An agent's act is permissible solely in virtue of the fact that she does not believe that its total consequences are overall worse than those of any option open to her; otherwise it is impermissible.

And we might convert OR into

Subjective Rossianism (SR). An agent's act is permissible solely in virtue of the fact that she does not believe that it has less on balance *prima facie* rightness than that of any option open to her; otherwise it is impermissible.⁴⁵

And we might convert Requirement R_o into the moral imperative

Requirement R_s . Intervene to prevent horrific evil you can prevent, unless you believe there is better reason for you not to intervene.

SMAC implies that even if the total consequences of our intervening to prevent Ashley's suffering are worse than the consequences of nonintervention, we should intervene, provided we don't believe that the total consequences of intervention are overall worse than those of nonintervention. Similarly, SR implies that even if the on balance *prima facie* rightness of intervention is less than that of nonintervention, we should intervene, provided we don't believe that the on balance *prima facie* rightness of intervention is less than that of nonintervention. And R_s implies that even if there is better reason not to intervene than there is to intervene, we should intervene, provided we don't believe that there is better reason not to intervene.

Notice that, on these views, the right- and wrong-making features of an act are typically accessible to us since our beliefs are typically accessible to us; consequently, we should not be in doubt about its moral status. Let's call any view that, like these three, posits right- and wrong-making features of an act that should not leave us in doubt about its moral status an instance of *Moral Accessibilism*, or *Accessibilism* for short.

⁴⁴ It has been my experience that some members of my audience, commentators and referees mainly (it wouldn't be nice to name names), suffer a temporary impairment of their auditory or visual capacities when they hear or read this sentence, or else they forget their lecture notes for their Logic 101 courses. For what they attribute to me is the thesis that we should be in doubt about whether we should have intervened to prevent Ashley's suffering. So, let me emphasize that I have *not* asserted that we should be in doubt about whether we should have intervened. I have asserted that *if we endorse an instance of Inaccessibilism*, then we should be in doubt about whether we should have intervened. In addition, I have asserted that *if we endorse OMAC, OR, or R_o* , then we should be in doubt about whether we should have intervened.

⁴⁵ In calling these views subjective maximizing consequentialism and subjective Rossianism, I do not mean to suggest that there are not other views that have a better claim to these labels. However, Ross himself moved from something OR-ish, in Ross 1930, to something SR-ish, in Ross 1939, due to what he deemed to be the undesirable epistemic implications of the former.

Accessibilism is relevant to the Moral Skepticism Objection. For suppose we endorse Requirement R_s . (The arguments to follow can be made with SMAC and SR, *mutatis mutandis*.) In that case, we should accept premise 2. After all, we don't believe that there is better reason for us not to intervene than to intervene, and so the presumption in favor of intervention is not overridden. However, if we endorse R_s , we should reject premise 1. For, given R_s , there is no tension between saying we should not be in doubt about whether we should intervene, on the one hand, and saying we should be in doubt about whether there is some reason we don't know of that would justify someone else's nonintervention, on the other hand. That's because, given R_s , the fact that we should not be in doubt about whether we should intervene is grounded in the twin fact that we should prevent suffering we can prevent and we don't believe that there is better reason for us not to intervene. The fact that we should not be in doubt about whether we should intervene is not grounded in what anyone else believes or doesn't believe. So even if we should be in doubt about whether some reason we don't know of would justify *God's* nonintervention, it does not follow that we should be in doubt about whether *we* should intervene. Thus, premise 1 of the Objection is false *given Requirement R_s* .

I expect the following response. Suppose you are right and we should be in doubt about whether there is a reason that would justify God in permitting Ashley's suffering. In that case, we should be in doubt about whether we should intervene. For if there is a reason that would justify God in permitting Ashley's suffering, then, although it might, for all we can tell, be one that would justify God's nonintervention but not ours—say, in virtue of some special entitlement that He has but we lack—it might also, for all we can tell, be one that would justify our nonintervention as well. So, for all we can tell, there is better reason for us not to intervene. So, we should be in doubt about whether there is better reason for us not to intervene. So, we should be in doubt about whether we should intervene. Therefore, premise 1 of the Objection is true.

By way of reply, recall that we are supposing that R_s is true, according to which we should intervene to prevent Ashley's suffering unless we believe that there is better reason for us not to intervene. But even though we should be in doubt about whether there is better reason for us not to intervene, it does not follow from R_s that we should be in doubt about whether we should intervene. For even if we are in doubt about whether there is better reason for us not to intervene, we don't believe that there is better reason for us not to intervene. According to R_s , the presumption in favor of intervention is overridden by *believing* that there is better reason for us not to intervene, not by *being in doubt* about whether there is better reason for us not to intervene.

So far I have argued that if our moral theories or principles are instances of Inaccessibilism, then we should reject premise 2 of the Objection, but if they are instances of Accessibilism, then we should reject premise 1. More specifically, I have argued that if we endorse OMAC, OR, or R_o , then we should reject premise 2, but if we endorse SMAC, SR, or R_s , then we should reject premise 1. No doubt it will have occurred to some readers that the moral theories and principles on which I have focused are not plausible enough to decide the matter. As I mentioned at the outset, I will address this concern in section 6.

4. Is Agnosticism at home with atheism and theism?

There are other things we bring to an assessment of the Moral Skepticism Objection, notably our atheism or theism, as the case may be. Two questions arise. First, what bearing should these have on our assessment of the simple version of the Objection, and the negative evaluation of it that I have offered? Second, is there some tension between them and Agnosticism?

Let's begin with atheism. Adherents of atheism include naturalists and members of nontheistic religions. I will focus on naturalism.⁴⁶ Suppose we are naturalists who endorse an instance of Inaccessibilism, say OMAC. Should we accept premise 2 in that case, contrary to what I argued? No. Naturalism itself gives us no reason to think that intervention or nonintervention will not have massive and inscrutable causal ramifications, nor does it give us reason to be optimistic about using expected value, indifference, or anything else to eliminate doubt. Now suppose we are naturalists who endorse an instance of Accessibilism, say R_s. In that case, should we accept premise 1 of the Objection, contrary to what I argued? No. For consider the conjunction of naturalism and the antecedent of premise 1: there is no God but we should be in doubt about whether there is some reason that *would* justify God in permitting, say, Ashley's suffering (if only there were a God). But this gives us no reason to believe that there is better reason for our nonintervention than our intervention, and so it gives us no reason to think that we should be in doubt about whether we should have intervened. Thus, so far as I can see, naturalists who endorse OMAC and R_s, or one of the other theories and principles I have mentioned above, have no reason to reject my negative evaluation of the Objection, at least not by virtue of their naturalism.

The Agnostic commends Agnosticism about the Inference to those of us who are naturalists.

Now let's consider theism. Suppose we endorse both Agnosticism and theism. Then we'll think there *really is* some reason that justifies God's nonintervention in Ashley's case, and we'll also think that we should be in doubt about what that reason is, given the information available to us. In that case, one might argue, we should also be in doubt about whether we should intervene. After all, if we think there really is a reason, a reason that God is actively using, so to speak, and we have no idea what it is, then, for all we can tell, there is a reason for us not to intervene; indeed, for all we can tell, *God's* reason not to intervene is a reason for *our* nonintervention as well.⁴⁷

I have three things to say about this worry.

First, suppose we are *instructed theists*, that is we reasonably think God has instructed humankind to prevent suffering in general and that God permits a lot of it precisely because He intends for us to try to prevent it. In that case, we need not be troubled by this worry. For in that case, we have strong reason to think that God not only wants us to prevent Ashley's suffering but has directed us to do so as well. Consequently, unless we have some special reason to permit Ashley's suffering, a reason that is strong enough to override the presumption in favor of intervention as well as his instruction to prevent it, we should not be in doubt about whether we should intervene. Consequently, we—that is, we instructed theists—should not be in doubt about whether there is a morally sufficient reason for our nonintervention.⁴⁸

⁴⁶ According to Draper 2004, natural entities are physical entities, anything composed of physical entities, and anything caused to exist solely by physical entities. Non-natural entities include abstracta, souls, and God. A supernatural entity is one that is non-natural and has causal powers. Naturalism is the thesis that there are no supernatural entities, not the thesis that there are no non-natural entities. Thus, on this definition, naturalists can have abstracta (assuming the Benacerraf problem can be solved without abstracta having causal powers), but they can't have Cartesian souls or God. For dissent on defining naturalism, see Rea 2002.

⁴⁷ *Could* the reason that justifies God's nonintervention be a reason for our nonintervention as well? I think so. Suppose God is justified in letting an alcoholic "hit bottom" so that she can see more clearly the consequences of her drinking and seek help to amend her condition. Why couldn't that reason justify her husband in doing the same? Alcoholics Anonymous calls this "tough love," a crucial element in recovery from alcohol dependency.

⁴⁸ Cp. Maitzen 2007.

But what if we are *uninstructed* theists? Doesn't the worry arise in that case? No; at least not if R_s is the only moral principle relevant to the case. (This is the second thing I have to say.) For in that case, we should intervene to prevent Ashley's suffering unless we believe there is better reason for us not to intervene. But we don't believe that. We are in doubt about it. And, according to R_s , the presumption in favor of intervening to prevent suffering is overridden only by *believing* that there is better reason not to intervene, not by *being in doubt* about whether there is one. The same point applies to SMAC and SR, *mutatis mutandis*.

Third, given OMAC and the other instances of Inaccessibilism I have mentioned, we should already be in doubt about whether we should intervene to prevent Ashley's suffering. Does the theist who is a proponent of, say, OMAC have *further* reason to be in doubt by virtue of her additional endorsement of Agnosticism? Perhaps so; perhaps the conjunction of OMAC and Agnosticism gives the theist who endorses both even more reason to be in doubt than she would have if she endorsed only Agnosticism and not OMAC, although it is not entirely clear why this would be so. But even if it is so, it is crucial to keep two things in mind: (a) such a theist has *sufficient* reason to be in doubt about whether we should intervene solely by virtue of her endorsement of OMAC, and (b) absent her endorsement of OMAC, she has *no* reason to be in doubt about whether we should intervene, at least not if she endorses an instance of Accessibilism in place of OMAC like those I have mentioned. OMAC and other instances of Inaccessibilism are the engines of doubt, not Agnosticism.⁴⁹

The Agnostic commends Agnosticism about the Inference to those of us who are theists.

I have focused on a simple version of the Moral Skepticism Objection. I now turn to some real, live proponents of the Objection, Michael Almeida and Graham Oppy. I shall contend that the Agnostic's negative evaluation of the Simple Version of the Objection applies with equal force to their version.⁵⁰ To minimize repetition, I will not rehash the substance of sections 2-5; I will simply show where it has a bearing on their version.

5. Almeida's and Oppy's version of The Moral Skepticism Objection

Almeida and Oppy target Bergmann 2001, which resembles Agnosticism.⁵¹ They write:

...Suppose we take seriously the idea that it follows from our acceptance of [Bergmann's] ST1-ST3 that [it] is not unlikely that there are goods beyond our ken—or relations beyond our ken between goods and evils (which themselves may or may not be beyond our ken)—which justify a perfect being in not preventing E. Suppose further that we are, right now, witnesses to E, and that we could intervene to stop it at no personal cost. What we have just conceded is that, merely on the basis of our acceptance of ST1-ST3, we should insist that it is not unlikely that there is some good which, if we were smarter and better equipped, we could recognize as a reason for a perfect being's not intervening to stop E. Plainly, we should also concede—by parity of reason—that, merely on the basis of our acceptance of ST1-ST3, we should insist that it is not unlikely that [I] there is some good which, if we were smarter and better equipped, we could recognize as a reason for our not intervening to stop the event. That is, our previous concession surely forces us to allow that, given our acceptance of ST1-ST3, it is not unlikely that [II] it is

⁴⁹ Cp. Almeida and Oppy 2005, xx-xx, where the same worry is discussed.

⁵⁰ Elsewhere, I contend the same for Russell 1996b, Tooley 2004, Pereboom 2004, Jordan 2006, Piper 2007, and Maitzen unpublished.

⁵¹ Almeida and Oppy 2003. Other targets include Wykstra 1984, Alston 1991 and 1996, and Howard-Snyder 1996a.

for the best, all things considered, if we do not intervene. But, if we could easily intervene to stop the heinous crime, then it would be *appalling* for us to allow this consideration to stop us from intervening. Yet, if we take the thought seriously, how can we also maintain that [III] we are morally required to intervene? After all, as a result of our acceptance of ST1-ST3, we are allegedly committed to the claim that it is not unlikely that it would be for the best, all things considered, if we did not do so. (505-506, bracketed Roman numerals added)

What should we make of the line of thought in this passage?

Well, for starters, nothing properly called “skeptical theism” is at issue; agnosticism about the Inference is what is at issue. Moreover, Almeida and Oppy attribute the following claims to those who defend agnosticism (emphasis added):

- “[It] is *not unlikely* that there are goods beyond our ken—or relations beyond our ken between goods and evils (which themselves may or may not be beyond our ken)—which justify a perfect being in not preventing E”.
- “[I]t is *not unlikely* that there is some good which, if we were smarter and better equipped, we could recognize as a reason for a perfect being’s not intervening to stop E.”
- “[I]t is *not unlikely* that there are unknown goods which would justify a perfect being in not preventing E”.

But neither Bergmann, Alston, nor I assign a probability to these things, not even the interval associated with *not unlikely*. Rather, we say that *we should be in doubt about* the matter, and that we are in no position to assign a probability.⁵² In what follows, I will reconstruct the argument of this passage so that its target is not a straw man.

The quotations above suggest a familiar argument schema, namely this:

1. If Agnosticism is true, then we should be in doubt about whether _____.
2. We should not be in doubt about whether _____.
3. So, Agnosticism is not true.

But what fills the blanks? I see three options, indicated by the bracketed Roman numerals in the quotation above, which I set out here for convenience:

- I. “there is some good which, if we were smarter and better equipped, we could recognize as a reason for our not intervening to stop [E]”
- II. “it is for the best, all things considered, if we do not intervene [to prevent E]”
- III. “we are morally required to intervene [to prevent E]”

These do not mean the same thing, nor do they nontrivially entail each other. Thus, if we substitute them into the schema, we have substantially different arguments.

The substitution instance provided by (III) is the simple version of the Objection. I have nothing more to say about it. The substitution instance provided by (II) is this:

⁵² As Almeida and Oppy recognize at 505, note 16. If you hold a gun to the Agnostic’s head and say, “Assign or die,” she will reluctantly comply: “The interval [0,1]”. But if you hold a gun to the Agnostic’s head and say, “Assign a discrete value or die,” she will say “Die”.

1. If Agnosticism is true, then we should be in doubt about whether it is for the best, all things considered, that we do not intervene to stop E.
2. We should not be in doubt about whether it is for the best, all things considered, that we do not intervene to stop E.
3. So, Agnosticism is not true.

Here we face a hermeneutical difficulty. What exactly do Almeida and Oppy mean when they use the locution “it is for the best, *all* things considered, that p”? In particular, do they intend for us to take the quantifier seriously? I assume so. In that case, a question arises: what is the intended scope of the quantifier? Many options present themselves. I have space to consider two.

The first option is fully unrestricted scope. We might imaginatively engage this option by supposing that there’s an atemporal perspective on what happens in the flow of time. Boethius ascribed such a perspective to God. He thought that God’s eternity was properly conceived of as an atemporal mode of existence. God exists outside of time and everything that happened in time was cognitively accessible to God in one fell swoop, in one “eternal now”. We need not suppose that such a perspective is in fact occupied in order to consider what would be available to an occupant of that perspective. So suppose there is such a perspective and consider what would be available to one who enjoyed that point of view, one who took in everything all at once *sub species aeternitatis*. Now, when we say that “it is for the best, *all* things considered, that p” we might mean (in the present context) that, from an atemporal perspective on everything, p is for the best; or that, taking into account everything, past, present, and future, it is for the best that p; or, more poetically, that it is for the best, all-things-considered-*sub-species-aeternitatis*, that p. Such is the fully unrestricted scope option. Suppose Almeida and Oppy intend something along these lines.

In that case, premise 2 of their argument is more clearly expressed as the claim that

- 2s. We should not be in doubt about whether it is for the best, all-things-considered-*sub-species-aeternitatis*, that we do not intervene to stop E.

2s is false, however. For it is for the best, all-things-considered-*sub-species-aeternitatis*, that we do not intervene only if there is nothing in the total consequences of our nonintervention that makes it better than the total consequences of our intervention. But intervention and nonintervention have massive and inscrutable causal ramifications and the unforeseeable consequences swamp the foreseeable ones. So we should be in doubt about whether there is nothing in the total consequences of our nonintervention that makes it better than the total consequences of our intervention. Therefore, we should be in doubt about whether it is for the best, all-things-considered-*sub-species-aeternitatis*, that we do not intervene to stop E.

If we want a true premise 2, we’ll need to restrict the scope of the quantifier. There are many options here. I will consider one. Suppose we restrict the quantifier to things that we have in fact considered. That is, suppose that when we use the locution that “it is for the best, *all* things considered, that p” we mean (in the present context) that it is for the best, all-things-*we-have-considered*, that p. The idea is that taking into account just those things that we have in fact considered, p is for the best. Suppose that Almeida and Oppy meant this.

In that case, premise 1 of their argument is better understood as the claim that

- 1w. If Agnosticism is true, then we should be in doubt about whether it is for the best, all-things-*we-have-considered*, that we do not intervene to stop E.

1w is false, however. The proposition that we should *not* be in doubt about whether intervention is for the best given just those things *we* have considered is in no tension whatsoever with the proposition that we should be in doubt about whether intervention is for the best given just those things *God* has considered.

Perhaps there is a way to restrict the scope of the quantifier in “it is for the best, *all* things considered, that p” so that both premise 1 and 2 are true. One referee made the following suggestion: to say that it is for the best, all things considered, that I perform a certain action is just to say that it is for the best, taking into account all of the *morally relevant* information, that I perform that act. This is unhelpful because it is uninformative. We need to be told what information counts as morally relevant. If taking into account *all* of the morally relevant information requires occupying a position akin to taking in everything all at once *sub species aeternitatis*, then the modified premise 2 will be false. If taking into account all of the morally relevant information only requires consulting what we believe, then the modified premise 1 will be false. Without substantive guidance as what counts as morally relevant information, the proffered suggestion is vacuous and thus unhelpful.

The helpful and informative candidates for what might be meant by the locution “it is for the best, *all* things considered, that p” that I can think of are such that if they imply that we should *not* be in doubt about whether it is for the best, all things considered, that we do not intervene, then Agnosticism does *not* imply that we should be in doubt about whether it is for the best, all things considered, that we do not intervene. But perhaps I’m missing something.

Let us now turn to the substitution instance we get with (I):

1. If Agnosticism is true, then we should be in doubt about whether there is some good which, if we were smarter and better equipped, we could recognize as a reason for our not intervening to stop E.
2. We should not be in doubt about whether there is some good which, if we were smarter and better equipped, we could recognize as a reason for our not intervening to stop E.
3. So, Agnosticism is not true.

What should we make of this argument?

That depends on what we bring to our assessment of it. Suppose we are instructed theists. Then, for reasons adumbrated earlier, we’ll deny premise 1.

But what if we are *uninstructed* theists or, better yet, *atheists*? Well, suppose we are, and suppose we also endorse OMAC. Then, for reasons adumbrated earlier, we should deny premise 2.

But what if we endorse R_s instead? In that case, interestingly enough, we’ll accept premise 1. For, according to Agnosticism, a sufficient condition for being in doubt about whether there’s a God-justifying good is being in doubt about whether the goods we know of are a representative sample of the whole lot. In that case, why shouldn’t that be sufficient for being in doubt about whether there’s a *human*-justifying good as well (if only we were smarter and better equipped to discern it)? It should. And neither uninstructed theism nor atheism should affect our

judgment on this score. Thus, if we are uninstructed theists or atheists who endorse both R_s and Agnosticism, we will also endorse:

- ~2. We should be in doubt about whether there is some good which, if we were smarter and better equipped, we could recognize as a reason for our not intervening to stop E.

But what's objectionable about that?

According to Almeida and Oppy, "it would be *appalling* for us to allow this consideration to stop us from intervening" (their emphasis). We agree. It would be appalling. But we're not going to allow our endorsement of ~2 to stop us from intervening. Nor should we. For we endorse R_s , according to which we should intervene to prevent E unless we believe there is better reason for us not to intervene. But we don't believe that. True enough, we are in doubt about whether there is some good which, if we were smarter and better equipped, we could recognize it as a reason for our not intervening to stop E. But that doesn't add up to *believing* there is a better reason for us not to intervene. Not by a long shot. Thus, our endorsement of ~2 will not stop us from intervening. (Similarly for SMAC and SR, *mutatis mutandis*.) As for Almeida and Oppy's rhetorical question, "if we take [~2] seriously, how can we also maintain that we are morally required to intervene?," well, *that's* how—if we endorse R_s .

Almeida and Oppy express their "key point" differently in another passage:

Here is another way of making our key point. Suppose that we try to give a rational reconstruction of the moral reasoning that we undertake when we reach the decision to intervene in the case in which we can easily prevent rape and murder. The reconstruction will have to go something like this:

- (1) There is *pro tanto* reason for me to intervene to prevent E. (Indeed, I have a *pro tanto* duty to intervene to prevent E.) (Premise)
- (2) I have found no *pro tanto* reason for me not to intervene to prevent E. (Premise)
- (3) (Hence) There is no *pro tanto* reason for me not to intervene to prevent E. (From 2)
- (4) (Hence) I have all things considered reason to intervene to prevent E. (From 1, 3)

If we like, we can make this reconstruction look even more like the evidential argument from evil, by casting it in terms of reasons why the world would not be non-arbitrarily improved if I were to prevent E. However, even the version which we have given makes the point clearly enough: our reasoning from *pro tanto* reasons to all things considered reasons always relies upon a 'noseeum' inference of just the kind which appears in our evidential argument from evil. If sceptical theism is sufficient to block 'noseeum' inferences about value, then we lose our ability to reason to all things considered conclusions about what to do. (507)

To repeat: nothing properly called "skeptical theism" is at issue; agnosticism about the Inference is. One can be an agnostic about the Inference even if one is an intellectually fulfilled atheist. At

least I hope so. It would be a real shame if the intellectual credentials of atheism hung on the propriety of the Inference.

What should we make of the “key point” here? If I understand it correctly, there are two basic claims. First, a “rational reconstruction” of the reasoning ordinary intelligent people use to decide that they should intervene to prevent suffering “will have to go something like” (1)-(4). Second, if Agnosticism is sufficient to block the Inference, we cannot sensibly reason along the lines of (1)-(4) to the conclusion that we should intervene. One might worry about the first claim but I will focus on the second.

Once again, we face hermeneutical obstacles. What sorts of things count as a *pro tanto* reason to perform an act? More importantly, what is an “all things considered reason” to perform an act? How seriously are we supposed to take the quantifier? What is its scope? What is it for someone to “have” an all things considered reason? What’s the relationship between *pro tanto* reasons and all things considered reason? And what does any of this have to do with what we *should* do? There are many alternatives here. It isn’t at all clear which of them Almeida and Oppy have in mind. In what follows, I’ll consider two options that occupy opposite ends of a spectrum of answers to these questions.

Option 1. What I should do is what I have all things considered reason to do. Moreover, I have all things considered reason to intervene to prevent Ashley’s suffering if and only if there is all things considered reason for me to intervene. It follows that I should intervene if and only if there is all things considered reason for me to intervene. Now take the quantifier in “all things considered reason” in a fully unrestricted fashion. It ranges over everything, including both reasons to intervene and reasons not to intervene. Finally, reasons to intervene and reasons not to intervene include the total consequences of intervention and nonintervention.

That’s the first option. Suppose it’s correct. Now: according to Almeida and Oppy, “something like” (1)-(4) constitutes the only rational reconstruction of our moral reasoning. Suppose they are right. Then, contrary to what they say, I cannot sensibly reason to the conclusion that I should intervene to prevent Ashley’s suffering. For intervention and nonintervention have massive and inscrutable causal ramifications and the unforeseeable consequences of intervention and nonintervention swamp the foreseeable ones. So I should be in doubt about whether the consequences of nonintervention are significantly better than those of intervention. So I should be in doubt about whether there is all things considered reason for me to intervene. So I should be in doubt about what I should do. On the first option, if Almeida and Oppy are right that “something like” (1)-(4) constitutes the only rational reconstruction of our moral reasoning, I cannot sensibly reason to the conclusion that I should intervene. On the first option, Almeida’s and Oppy’s rational reconstruction of our moral reasoning *all by itself* suffices to make it the case that I cannot sensibly reason to the conclusion that I should intervene. Agnosticism just comes along for the ride.

Option 2. What I should do is what I have all things considered reason to do. Moreover, I have all things considered reason to intervene to prevent Ashley’s suffering if and only if what I believe supports my intervention. It follows that I should intervene if and only if what I believe supports my intervention. Finally, in so far as reasons to intervene and reasons not to intervene include the consequences of intervention and nonintervention, it is only those that I believe to be the case that are morally relevant.

That’s the second option. Suppose it’s correct. Then even if Agnosticism is sufficient to block the Inference, it does not follow that I cannot sensibly reason along the lines of (1)-(4) to the conclusion that I should intervene. For I can sensibly conclude on the basis of (1)-(3) that

what I believe supports my intervention, in which case I have all things considered reason to intervene, and so I should intervene. This line of thought is in no tension whatsoever with my admitting that nonintervention is permissible for others, say God. After all, what God believes might be different from what I believe, in which case what he believes might well support his nonintervention, and render his nonintervention permissible. So I can reason along the lines of (1)-(4) to the conclusion that I should intervene despite the fact that Agnosticism is sufficient to block the Inference.

If the argument of the “key point” passage is to have any bite, we cannot endorse the principles adumbrated by Options 1 and 2. So what principles or theories might we endorse so that the argument of this passage and the Moral Skepticism Objection more generally will have more of a bite?

6. Conclusion

Good question. At the outset of this paper, I mentioned that several people have put forward the following objection to the overall argument of this paper. “In the course of your overall argument, you consider moral theories and principles that are tailor-made to make the argument succeed. However, these theories and principles are massively implausible. If you had considered *my* favored theory and principles, which, of course, are highly plausible, your argument would have been an obvious failure,...

Well, what do I have to say for myself?

This: guilty as charged! If I had had the foresight, I would have contacted you. Please accept my apologies. But perhaps you can help out, even at this late date. Here’s what I would find helpful.

Specify your favored theory of the right- and wrong-making features of an act, or specify a moral principle that you think governs the prevention of horrific evil. If it contains terms of art like “all things considered reason,” as in “we have an all things considered reason to prevent Ashley’s suffering, or “all things considered duty” and “grounded in virtue” as in “we have an all things considered duty—grounded in virtue—to prevent Ashley’s suffering,” and the like, give a helpful and informative account of what they mean. Then do two things.

First, explain how it is that, on your theory or principles, we should not be in doubt about whether we should intervene to prevent Ashley’s suffering. When you give your explanation, be sure to take into account the fact that most of what we do is either directly or indirectly identity-affecting, and thus that most of what we do has massive causal ramifications. If you deny this fact, explain why. If you don’t deny it, explain how it is that, despite this fact, we should not be in doubt about whether we should intervene, given your theory or principle. If your explanation appeals to expected value, indifference, intuition, virtue, duties, or the tea leaves in your kitchen sink, explain why objections to your explanation have no force.

Second, explain how it is that, given your theory or principles, Agnosticism implies that we should be in doubt about whether we should intervene to prevent Ashley’s suffering. And whatever you say on that score, make it plain why it is that your own theory or principles aren’t really driving the doubt and not Agnosticism.

I have considered only six theories or principles in this paper, none of which are up to these two tasks, by my lights. My *tentative* hypothesis is that your theory or principles won’t be up to them either. But I may well be wrong about that.⁵³

⁵³ For comments on previous drafts, I thank Michael Almeida, Nick Beckstead, Michael Bergmann, Rebecca Copenhaver, Trent Dougherty, Robert Epperson, Frances Howard-Snyder, Hud Hudson, Shieva Kleinschmidt,

Christian Lee, Wes Morriston, Michael Murray, Robert Pasnau, Michael Rea, John Schellenberg, Michael Tooley, Ryan Wasserman, Dennis Whitcomb, and three anonymous referees.

REFERENCES

- Michael Almeida and Graham Oppy, 2003. "Sceptical Theism and Evidential Arguments from Evil," *Australasian Journal of Philosophy*, 496-516.
- Michael Almeida and Graham Oppy, 2005, "Reply to Trakakis and Nagasawa," *Ars Disputandi* 5.
- Michael Almeida and Graham Oppy, 2005. "Evidential Arguments from Evil and Sceptical Theist Responses," *Philo* 8, 84-94.
- William Alston, 1991. "The Inductive Argument from Evil and the Human Cognitive Condition," *Philosophical Perspectives* 5, 29-67, collected in Howard-Snyder 1996b.
- William Alston, 1996. "Some (Temporarily) Final Thoughts on Evidential Arguments from Evil," in Howard-Snyder 1996b, 311-32.
- David Anderson, 2005. *Epistemic Humility and William Rowe's Evidential Argument from Evil*. Unpublished honor's thesis at Acadia University.
- David Anderson, "Skeptical Theism and Ordinary Moral Reasoning," unpublished.
- Michael Bergmann, 2001. "Skeptical Theism and Rowe's New Evidential Argument from Evil," *Noûs*, 278-96.
- Michael Bergmann and Michael Rea, 2005. "In Defense of Skeptical Theism: A Reply to Almeida and Oppy," *Australasian Journal of Philosophy* 83, 241-51.
- Michael Bergmann, 2009. "Skeptical Theism and the Problem of Evil," in Thomas Flint and Michael Rea, eds., *The Oxford Handbook of Philosophical Theology*, New York: Oxford University Press.
- Terry Christlieb, 1992. "Which Theisms Face an Evidential Problem of Evil?," *Faith and Philosophy*, 45-64.
- Terence Cuneo, 2007, *The Normative Web: An Argument for Moral Realism*, New York: Oxford University Press.
- Paul Draper, 1989. "Pain and Pleasure: An Evidential Problem for Theists," *Nous* 23, 331-50.
- Paul Draper, 1992, "Probabilistic Arguments from Evil," *Religious Studies* 28, 303-17.
- Paul Draper, 2004a. "More Pain and Pleasure: A Reply to Otte," in Peter van Inwagen, editor, *Christian Faith and the Problem of Evil*, Grand Rapids, Michigan: Eerdmans Publishing Co.
- Paul Draper, 2004b. "God, Science, and Naturalism," in William Wainwright, editor, *Oxford Handbook of Philosophy of Religion*, Oxford: Oxford University Press, xxx-xxx.
- Fred Feldman, 2006. "Actual Utility, the Objection from Impracticality, and the Move to Expected Utility," *Philosophical Studies*, 49-79.
- Richard Gale, 2007. "Evil and Alvin Plantinga," in Deane-Peter Baker, editor, *Alvin Plantinga*, New York: Cambridge University Press, 48-70.
- James Hawthorne, 2004. "Inductive Logic," in Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, URL = <http://plato.stanford.edu/archives/fall2004/entries/logic-inductive/>
- Daniel Howard-Snyder, 1996a. "The Argument from Inscrutable Evil," in Howard-Snyder 1996b.
- Daniel Howard-Snyder, ed., 1996b. *The Evidential Argument from Evil*, Bloomington: Indiana University Press.
- Daniel Howard-Snyder, 2003. "On Hume's Philosophical Case Against Miracles," in *God Matters* (Longmans 2003), eds. Christopher Bernard and Raymond Martin, 395-411. For an updated version with corrections, see <http://www.ac.wvu.edu/~howardd/miracles.pdf>

- Daniel Howard-Snyder, "The Argument from Horrors," unpublished.
- Daniel Howard-Snyder and Michael Bergman, 2001a. "Grounds for Belief in God Aside, Does Evil Make Atheism more Reasonable than Theism?," in William Rowe, ed., *God and the Problem of Evil*, Malden, Massachusetts: Blackwell Publishing, 140-55.
- Daniel Howard-Snyder and Michael Bergmann, 2001b. "Reply to Rowe," in William Rowe, ed., *God and the Problem of Evil*, Malden, Massachusetts: Blackwell Publishing, 137-40.
- Frances Howard-Snyder, Daniel Howard-Snyder, and Ryan Wasserman, 2009. *The Power of Logic*, 4th edition, New York: McGraw-Hill.
- Frances Howard-Snyder, 1997. "The Rejection of Objective Consequentialism," *Utilitas*, 241-48.
- Jeff Jordan, 2006. "Does Skeptical Theism Imply Moral Skepticism?," *Philosophy and Phenomenological Research*, 403-417.
- Keynes, J.M., 1921. *A Treatise on Probability*. London: Macmillan (reprinted 1963).
- James Lenman, 2000. "Consequentialism and Cluelessness," *Philosophy and Public Affairs* 29, 342-70.
- Isaac Levi, 1974. "On Indeterminate Probabilities," *Journal of Philosophy* 71 5-29.
- Shelly Kagan, 1998. *Normative Ethics*, Boulder: Colorado, Westview Press.
- J.L. Mackie, 1977. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Stephen Maitzen, 2007. "Skeptical Theism and God's Commands," *Sophia* 35, 235-41.
- Stephen Maitzen, unpublished. "Skeptical Theism and Moral Obligation," presented at the Eastern Division Meeting of the American Philosophical Association, 2007.
- Elinor Mason, 2004. "Consequentialism and the Principle of Indifference," *Utilitas*, 316-21.
- G. E. Moore, 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- Wes Morriston, 2004. "The Evidential Argument from Goodness," *The Southern Journal of Philosophy*, XLII, 87-101.
- Alastair Norcross, 1990. "Consequentialism and the Unforeseeable Future," *Analysis* 50, 253-56.
- Graham Oddie and Peter Menzies, 1992. "An Objectivist's Guide to Subjective Value," *Ethics*, 512-533.
- Derk Pereboom, 2004. "The Problem of Evil," in William E. Mann, editor, *The Blackwell Guide to Philosophy of Religion*, Malden, Massachusetts: Blackwell Publishing.
- Mark Piper, 2007. "Skeptical Theism and the Problem of Moral Aporia," *International Journal for Philosophy of Religion* 62, 65-79.
- Alvin Plantinga, 1998. "Epistemic Probability and Evil," *Archivio di filosofia* (Italy) 56, 557-84, collected in Howard-Snyder 1996b.
- Michael Rea, 2002. *World without Design: The Ontological Consequences of Naturalism*, Oxford: Clarendon Press.
- W.D. Ross, 1930. *The Right and the Good*, Oxford: Oxford University Press.
- W.D. Ross, 1939. *Foundations of Ethics*, Oxford: Clarendon Press.
- Rowe, William L. 1978. *Philosophy of Religion: An Introduction*, 1st edition, Encino, California: Dickenson Publishing Company.
- William Rowe, 1979. "The Problem of Evil and Some Varieties of Atheism," *American Philosophical Quarterly* 16, 335-41, collected in Howard-Snyder 1996b.
- Rowe, William L. 1986. "The Empirical Argument from Evil," in Robert Audi and William Wainwright, eds., *Rationality, Religious Belief, and Moral Commitment*, Ithaca, New York: Cornell University Press, 227-47.
- Rowe, William L. 1988. "Evil and Theodicy," *Philosophical Topics* 16, 119-32.

- Rowe, William L. 1991. "Ruminations about Evil," *Philosophical Perspectives* 5, 69-88.
- Rowe, William L. 1995. "William Alston on the Problem of Evil," in Thomas D. Senior, ed., *The Rationality of Belief and the Plurality of Faith: Essays in Honor of William P. Alston*. Ithaca, New York: Cornell University Press, 71-93.
- Rowe, William L. 1996. "The Evidential Argument from Evil: A Second Look," in Howard-Snyder 1996b, 262-85.
- Rowe, William L. 2001a. "Grounds for Belief Aside, Does Evil Make Atheism More Reasonable than Theism," in William Rowe, ed., *God and the Problem of Evil*, Malden, Massachusetts: Blackwell Publishing, 124-37.
- Rowe, William L. 2001b. "Reply to Howard-Snyder and Bergmann," in William Rowe, ed., *God and the Problem of Evil*, Malden, Massachusetts: Blackwell Publishing, 155-58.
- Rowe, William L. 2006. "Friendly Atheism, Skeptical Theism, and the Problem of Evil," *International Journal for the Philosophy of Religion* 59, 79-92.
- Bruce Russell and Stephen Wykstra, 1988. "The 'Inductive' Problem of Evil: A Dialogue," in *Philosophical Topics* 16, 119-60.
- Bruce Russell, 1989. "The Persistent Problem of Evil," *Faith and Philosophy* 6, 121-39.
- Bruce Russell, 1996a. "Defenseless," in Howard-Snyder 1996b, xxx-xxx.
- Bruce Russell, 1996b. "Why Doesn't God Intervene to Prevent Evil?," in Louis Pojman, ed., *Philosophy: The Quest for Truth*, 3rd edition, Belmont, California: Wadsworth, 74-80, and at http://www.infidels.org/library/modern/bruce_russell/intervene.html
- John Schellenberg, 2000. "Stalemate and Strategy: Rethinking the Evidential Argument from Evil," in *American Philosophical Quarterly* 37, 405-19.
- John Schellenberg, 2007. *The Wisdom to Doubt: A Justification of Religious Skepticism*, Ithaca, New York: Cornell University Press.
- Ira M. Schnall, 2007. "Sceptical Theism and Moral Skepticism," *Religious Studies* 43, 49-69.
- James Sennett, 1991. "The Inscrutable Evil Defense," *Faith and Philosophy*, 340-53.
- J.C.C. Smart and Bernard Williams, 1973. *Utilitarianism: For and Against*. Cambridge: Cambridge University Press.
- Daniel Snyder, 1990. "Surplus Evil," *Philosophical Quarterly*, 78-86.
- Jim Stone, 2003. "Evidential Atheism," *Philosophical Studies* 114, 253-77.
- Richard Swinburne, 1988. "Does Theism Need a Theodicy," *Canadian Journal of Philosophy* 18, 287-312.
- Mark Timmons, 2002, *Moral Theory: An Introduction*, New York: Rowman and Littlefield.
- Michael Tooley, 1991. "The Argument from Evil," *Philosophical Perspectives*, 89-134.
- Michael Tooley, 2004. "Comments on Daniel Howard-Snyder's Paper," delivered at the *Theology Forum*, University of Colorado, October 2004.
- Michael Tooley, 2008. "The Problem of Evil," *The Stanford Encyclopedia of Philosophy*. URL: <http://plato.stanford.edu/entries/evil/>
- Nick Trakakis, 2003. "What No Eye Has Seen: The Skeptical Theist Response to Rowe's Evidential Argument from Evil," *Philo*: 6, 250-66.
- Nick Trakakis and Yujin Nagasawa, 2004. "Skeptical Theism and Moral Skepticism: A Reply to Almeida and Oppy," *Ars Disputandi* 4.
- Nick Trakakis, 2007. *The God Beyond Belief: In Defence of William Rowe's Evidential Argument from Evil*, New York: Springer Publishing.
- Bas van Fraassen, 1990. "Figures in a Probability Landscape," in J. Dunn and A. Gupta, eds., *Truth or Consequences*. Dordrecht: Kluwer.

- Bas van Fraassen, 1998. *Laws and Symmetry*, Oxford: Oxford University Press.
- Roger White, forthcoming. "Evidential Symmetry and Mushy Credence," *Oxford Studies in Epistemology*.
- Stephen Wykstra, 1984. "The Humean Obstacle to Evidential Arguments from Suffering: On Avoiding the Evils of 'Appearance'," *International Journal for the Philosophy of Religion*, 73-94.
- Stephen Wykstra, 1996. "Rowe's Noseeum Argument from Evil," in Howard-Snyder 1996b, 126-50.
- Linda Zagzebski, 2007. *Philosophy of Religion: A Historical Introduction*, Malden, Massachusetts: Blackwell Publishing.