Faith and Philosophy: Journal of the Society of Christian Philosophers

Volume 19 | Issue 2

Article 1

4-1-2002

Rowe's Probabilistic Argument from Evil

Richard Otte

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

Recommended Citation

Otte, Richard (2002) "Rowe's Probabilistic Argument from Evil," *Faith and Philosophy: Journal of the Society of Christian Philosophers*: Vol. 19 : Iss. 2 , Article 1.

DOI: 10.5840/faithphil200219217

Available at: <https://place.asburyseminary.edu/faithandphilosophy/vol19/iss2/1>

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

ROWE'S PROBABILISTIC ARGUMENT FROM EVIL

Richard Otte

In this article I investigate Rowe's recent probabilistic argument from evil. By using muddy Venn diagrams to present his argument, we see that although his argument is fallacious, it can be modified in a way that strengthens it considerably. I then discuss the recent exchange between Rowe and Plantinga over this argument. Although Rowe's argument is not an argument from degenerate evidence as Plantinga claimed, it is problematic because it is an argument from partitioned evidence. I conclude by discussing the modified argument and the epistemic framework Rowe is assuming in his argument.

Introduction

One of the most important arguments from evil has been developed by William Rowe. Recently, Rowe has presented an intriguing new probabilistic problem from evil and claims it is an improvement on his earlier arguments.¹ The conclusion of this argument is that our knowledge of certain evils lowers the probability of God's existing; that is, the probability of God's existing conditional on these evils is less than the probability of God's existing simpliciter (or conditional only on our background information). Rowe's argument for this is as complicated as it is clever and ingenious. After stating his premises, Rowe uses various theorems in the probability calculus to derive his conclusion. Unfortunately, this use of the probability calculus has prevented the essence of Rowe's argument from being clearly understood. In what follows I will present Rowe's argument geometrically instead of relying on theorems in the probability calculus. This will enable us to easily see the structure and essence of the argument that was hidden when presented using the probability calculus. Looking at the argument geometrically will also allow us to see that although Rowe's argument is fallacious, it can be modified in a way that strengthens the argument considerably. We will then look at the recent exchange between Alvin Plantinga and Rowe.² We will see that although Plantinga was incorrect in claiming Rowe's argument is an "argument from degenerate evidence," it does suffer from a related problem, in that it is "an argument from partitioned evidence." We will also see that my modification of Rowe's argument does not have this problem. In this exchange Rowe presents a complicated version of his argument that is not equivalent to either his original version or my modified version. Although Rowe is correct that

this more complicated argument is also not an argument from degenerate evidence, it is much more problematic than the modified version I present. I will conclude by looking at the crucial assumption of my modified version of Rowe's argument and more generally at the epistemic framework that Rowe is assuming for his discussion.

Rowe's Argument

Rowe (1996) begins his argument by describing two instances of evil, which he denotes by E1 and E2 (the details don't concern us). He then claims that theists are committed to the following principle P, and that P lowers the probability of God's existing:

P: No good we know of justifies an omnipotent, omniscient, perfectly good being in permitting E1 and E2. (1996: 263)

According to Rowe, in order for a good to justify God in permitting an evil, God must exist; it is not possible for God to not exist and a good justify God in permitting an evil. However, we also have a related concept of a good being good enough to justify God in permitting those evils which does not require that God exist. Rowe agrees, and says "[c]learly, if God doesn't exist, e [some evil] is actual, and g [some good] is not actual, it still may be true that if God were to exist and g and e were actual, he would be justified by g in permitting e. (1998: 549) We can express this concept in the following:

A good is *able* to justify God in permitting an evil if that good would justify God in permitting that evil if God were to exist.

According to this definition a good may be able to justify God in permitting an evil even if God does not exist. This is because if God were to exist, he would be justified by that good in permitting that evil. For ease of exposition let KA stand for there is a known good that is able to justify God in permitting E1 and E2, let UKA stand for there is an unknown good able to justify God in permitting E1 and E2, and let G stand for God exists. Given this, we can write Rowe's principle P as:

P': not-(G and KA)

It is easy to see that this is equivalent to Rowe's principle P. P implies P' because if no good we know of justifies God in permitting E1 and E2, then it will not be the case that God exists and there is a known good that is able to justify God in permitting those evils. And P' implies P because if it is not the case that both God exists and there is a known good able to justify God in permitting those evils, then there is no good we know of that actually justifies God in permitting those evils. Thus not-(G and KA) expresses Rowe's principle P; there is no known good that justifies God in permitting E1 and E2.³

At this point Rowe gives a complicated argument using the probability calculus for the claim that P lowers the probability of God's existing, which is the primary conclusion of his argument. Given this, he also argues for

the claim that *P* makes it likely that God does not exist. That is, he argues for the following two important claims:⁴

$$\begin{aligned} P(G/P) &< P(G) \\ P(G/P) &< .5 \end{aligned}$$

His argument is primarily concerned with the first claim. The second follows immediately from the first and the assumption that $P(G) = .5$. This is a very important assumption for Rowe because without it the probability of God's existing could be much greater than .5, even if all of our evidence lowered the probability of God's existing. But other than a short comment by Rowe, I know of no defense of this assumption in the literature. According to Rowe, we need to assume the probability of $P(G) = .5$ in order not to beg any questions: "In order not to beg any of these questions, I will assign a probability of 0.5 to $\text{Pr}(\neg G/k)$ and, of course, 0.5 to $\text{Pr}(G/k)$." (1996: 265) But this is mistaken. It is often thought that the best way to avoid taking a stand is to assign a proposition a probability of .5. After all, this means that one is not saying it is more likely to be true than not. But claiming a proposition has probability of .5 is making a very strong claim; it is not suspending judgment on the matter, which is what not begging any questions would amount to. For example, suppose you are given an alien die. You know nothing about this alien die except that it is used in games of chance by aliens. Suppose it even has two outcomes: *H*, *T*. Should you assign a probability of .5 to *H* and to *T*? I doubt it. If you know nothing about this die, you should not think it is as likely as not that *H* would result as would *T*; you should suspend judgment on this issue. But suspending judgment is not the same as assigning a probability of .5; you suspend judgement because you do not know if the outcomes are equally likely, or if *H* is twice as likely as *T*. The best way to represent complete suspension of judgment would be to assign *H*, and *T*, the interval $\langle 0,1 \rangle$. We simply do not know how likely *H* and *T* on this alien die are.

If, following Rowe, we do not want to beg any questions, we should not assign $P(G)$ and $P(\text{not-}G)$ a probability of .5, but instead should suspend judgment on those probabilities. If we suspend judgment on how likely *G* is, Rowe's argument will not show that *P* makes it likely that God does not exist. The only reason Rowe was able to draw that conclusion was because he assumed that the initial probability of *G* was .5. Given that *P* lowers the probability of *G*, and $P(G)$ is initially .5, after learning *P*, $P(G)$ must be less than .5. But if we don't know how likely *G* is initially, we won't know if *P* lowers it to less than .5. Thus I do not think Rowe's argument for his second claim is successful.

In this paper I will focus on Rowe's argument for his first claim, that $P(G/P) < P(G)$. Since I don't dispute Rowe's derivation of this claim in the probability calculus I won't repeat it here. However, in order to better understand Rowe's argument, I will present it geometrically instead of algebraically. Instead of following Rowe and presenting his argument by manipulations of the probability calculus, I will illustrate his claims using what van Fraassen calls 'muddy Venn diagrams.' These have the virtue of making the argument much easier to understand.

In a muddy Venn diagram we have one unit of mud evenly distributed on our tabletop, which is divided into various areas. The amount of mud on an area represents the probability of that area. We always keep one unit of mud on the table, and the probability of a certain region is always the proportion of the mud on that area. If we learn that some proposition X is not the case, we represent this in the muddy Venn diagram by wiping all the mud off of the area representing X on the table; we then distribute the mud proportionally over the remaining areas (or normalize so that the total mud left sums up to 1). If we learn that some proposition X is the case, we wipe all the mud off of the area representing not- X and distribute it evenly over the X area.

We can reconstruct Rowe's reasoning in our muddy Venn diagram as follows. Divide our tabletop into two sections, representing G and not- G .

Since Rowe thinks we need to assume each of these has a probability of .5, we'll give each of them equal area with an equal amount of mud on them. According to Rowe the theist is committed to either a known good justifying God in permitting $E1$ and $E2$, or some unknown good justifies him in permitting $E1$ and $E2$. This is equivalent to saying God exists and there is a known good able to justify those evils, or God exists and some unknown good is able to justify God in permitting those evils. Thus we now divide the G area into two sections, representing KA and UKA .

Although Rowe assigns (G and KA) and (G and UKA) equal probability, we will see that Rowe's argument doesn't depend upon how we divide these up, as long as KA gets some mud on it. We now need to look at the effect of learning that P is true, i.e., learning not-(G and KA). This has the effect of removing the (G and KA) section from the table, and redistributing the mud on the remaining areas. In the following diagram the mud on the blank area has been removed from the table and distributed evenly on the remaining areas. This represents the effect of learning that we don't know of any good that justifies God in permitting these evils.

We are interested in how likely G is before and after we learned that not-(G and KA). To do this we compare the areas or amount of mud on the G area before and after learning that information. It is clear that there is less mud on the G area after learning that not-(G and KA) than before; not-(G and KA) lowers the probability of G. Furthermore, if we began by assigning G probability .5, i.e., assigning it half the mud, we can see that G now has less than half the mud; thus the final probability of G is less than .5. This shows that Rowe's basic probabilistic reasoning is correct: the information that not-(G and KA) does lower the probability of God's existing (and lowers it to less than .5 if it was originally assumed to be .5). Rowe concludes from this that the facts about evil that we are aware of lower the probability of God's existing. He takes this to support that it is irrational to believe in God unless we have "strong evidence" for the existence of God.(1996; 282)

Analysis of the Argument

Although Rowe has given a valid argument for his two claims, we need now to think about how successful his general argument is. We've seen that Rowe has correctly pointed out the probabilistic implications of our learning this evidence, but we now need to investigate whether Rowe has presented a successful and convincing probabilistic argument from evil. Rowe is not clear about the epistemic framework he is assuming for his argument, and how various probabilities are supposed to be connected with what it is rational to believe. But whatever his framework, one thing that is certain is that we need to look at all of the relevant evidence in

deciding how evil affects rational belief in God. In confirmation theory one of the most widely accepted principles is the requirement of total evidence. But there is evidence about evil that Rowe ignores that renders his argument ineffectual. I will now show that Rowe's argument appeared successful only because it artificially excluded certain evidence we have.

Recall that Rowe's argument is based upon us learning principle P:

P: No good we know of justifies an omnipotent, omniscient, perfectly good being in permitting E1 and E2,

which we symbolized as:

P': not-(G and KA).

However, in addition to learning P, we have also learned that even if God does not exist, there is no good that is able to justify God in permitting those evils. In other words, we have learned that there is no good in the world we know of that is able to justify God, if he exists, in permitting those evils. We have learned that there is no good in the world we know of that is such that it *would* justify God in permitting these evils. As a result of learning this, the theist is committed to principle P, but P is only part of what the theist has learned. Rowe himself, in a footnote, brings up this possibility. He says that it is important not to confuse his principle P with principle P*:

P*: No good we know of *would* justify God (*if he exists*) in permitting E1 and E2. (original italics) (1996: 283)

Rowe doesn't give any reason for using P instead of P* in his argument, and he doesn't give any reason to think that P* is false. Since it is fallacious to ignore relevant evidence, in order for Rowe to use P instead of P* in his argument, Rowe must claim that the theist knows P and does not know P*. However, Rowe has given us no reason to believe this is true. This is a crucial assumption of his argument, but it appears false. Basically, we first learned that there is no known good able to justify God in permitting those evils. In other words, we've learned that P* is true as well as P. We didn't simply learn that it is false that both God exists and there is a known good that actually justifies him in permitting the evil, we also learned that it is false that there is a known good that could justify God in permitting those evils. Using our symbolization, we first learned not-KA instead of not-(G and KA). Of course, In learning not-KA we have also learned not-(G and KA), which is Rowe's principle P. Thus Rowe is correct that the theist is committed to P, but that is because the theist has normally learned something more, namely not-KA. Of course, it is possible for a theist to learn P without learning not-KA, and for this theist Rowe's argument may not ignore any relevant evidence. But that would be a highly unusual situation, and I don't know of any theist for whom this is true; almost all theists have learned not-KA as well as not-(G and KA). Rowe may think theists have learned P without learning P*, but he has given us no reason to think this is true; if he wishes us to accept his original argument and conditionalize only on P, he will need to

provide a reason to accept this assumption. Since we need to take account of all the relevant evidence that we've learned, we need to incorporate this evidence into Rowe's argument.

Let us look at how Rowe's probabilistic reasoning would look if we don't ignore this additional evidence. Returning to our muddy Venn diagram, we now must divide the area on the table representing not-G into three areas: a known good is able to justify God in permitting these evils (KA), an unknown good is able to justify God in permitting these evils (UKA), and no good is able to justify God in permitting those evils (NA). If God does not exist, one of these three possibilities must be correct. For now, let us postpone discussing what proportions of the not-G area each of these should get.

not-G	G
KA	KA
UKA	UKA
NA	

Taking account of all of our evidence, what we have learned is:

not-KA,

which is equivalent to

not-(G and KA) and not-(not-G and KA).

To represent this in our diagram, we need to erase the KA area and redistribute the mud.

not-G	G
KA	KA
	UKA
UKA	
NA	

When we do that, we are left with (G and UKA) or (not-G and UKA) or (not-G and NA). The important question is whether learning this information has lowered the probability of God's existing. Earlier we saw that learning only that not-(G and KA) had to lower the probability of G. But learning not-KA does not have this implication. Looking at the diagram, we need to determine whether the mud on the G area is less after learning not-KA than before. Since we are granting Rowe's assumption that G and not-G are equally probable, this would be true if more mud was taken off the G side than was taken off the not-G side. It is obvious that this depends on what the probabilities of (G and KA) and (not-G and KA) are. If $P(\text{G and KA})$ is greater than $P(\text{not-G and KA})$, then learning not-KA does lower the probability of G. But if $P(\text{G and KA})$ is less than $P(\text{not-G and KA})$, then learning not KA raises the probability of G. The situation is much more complicated than it was in Rowe's original argument. By restricting himself to only part of the evidence, Rowe was able to demonstrate that some (partial) evidence did lower the probability of God's existing. But when we take a more complete look at the evidence, we find that we cannot infer that the probability of God's existing is lowered. It follows only on the assumption that (G and KA) is more likely than (not-G and KA).

The following two diagrams illustrate the importance of knowing whether KA is more likely on G than on not-G:

On the top, we have an example where KA is more likely on G than on not-G, and on the bottom we have an example where KA is more likely on not-G than on G. Of course, there are an unlimited number of different diagrams, and these are just illustrations. In the situation represented by the top diagram, learning not-KA lowers the probability of God's existing. But in the situation represented by the lower diagram, learning not-KA raises the probability of God's existing.

We thus see that in order to modify and develop a successful version of Rowe's argument we need to conditionalize on not-KA instead of on not-(G and KA), and we also need to assume that (G and KA) is more likely than (not-G and KA). We can state a modified version of Rowe's argument as follows:

1. not-KA (premise, what we have learned)
2. $P(G \ \& \ KA) > P(\text{not-G} \ \& \ KA)$ (premise)
3. The facts of evil we are aware of (not-KA) lower the probability of God's existing (from 1,2)

This argument is much better than Rowe's argument because it does not ignore relevant evidence we have. This modified argument does not ignore our information that it is false that (not-G and KA) as did Rowe's argument. However, this argument is highly dependent upon Rowe's assumption that $P(G) = P(\text{not-G}) = .5$. If we do not make that assumption, the above argument is not valid. Since I think there are good reasons not to assume $P(G) = P(\text{not-G}) = .5$, it would be preferable to have a version of Rowe's argument that did not depend on this assumption. The above argument was successful because our information removes mud from a greater proportion of the G area than the not-G area; since the G and not G area had equal parts of mud, to do this all that was needed was for more mud to be taken off of the G area. But if the G and not-G areas are not equal, taking more mud off the G area does not guarantee that a greater proportion of mud is taken off of the G area. We can fix this problem by requiring that KA be a greater proportion of G than it is of not-G. The best way to do this would be to replace premise 2 in the above argument by the following:

- 2'. $P(KA/G) > P(KA/\text{not-G})$.

This premise simply says that KA is a greater proportion of G than it is of not-G. If the probabilities of G and not G are equal, then this premise is equivalent to our original premise 2. The advantage of using this premise is that it gives us a modified argument from evil that will be valid even if we don't begin with the assumption that $P(G) = P(\text{not-G}) = .5$. No matter what initial probabilities of G and not-G we begin with, if premises 1 and 2' are correct, then the probability of God's existing is lowered by the information not-KA. Our final version of the modified argument is:

1. not-KA (premise)
- 2'. $P(KA/G) > P(KA/\text{not-G})$. (premise)
3. The facts of evil we are aware of (not-KA) lower the probability of God's existing (from 1,2)

To summarize, our modified version of Rowe's argument is different from Rowe's original argument in that it conditionalizes on not-KA instead of P, and so it does not ignore our information that it is false that (not-G and KA). Furthermore, the argument assumes as a premise that $P(KA/G) > P(KA/\text{not-G})$, and thus it does not require that we assume $P(G) = P(\text{not-G}) = .5$. We will soon see that this argument is stronger and more successful than Rowe's argument.

However, it is not obvious that the assumption that $P(KA/G) > P(KA/\text{not-G})$ is correct, and I suspect we have no way of knowing if it is true. There may be some reason to think $P(KA/G)$ is low, but I see no reason to think that $P(KA/\text{not-G})$ is lower. One problem is that it is not clear what we are conditionalizing on when we conditionalize on not-G, and as a result it is difficult to assign a probability to $P(KA/\text{not-G})$. Because of this a strong case can be made for withholding judgement on whether the assumption required by this modified version of Rowe's argument is correct. I will discuss this in more detail near the end of this paper, but for the purposes of argument let us grant Rowe this assumption in the modified version of his argument, keeping in mind that this assumption may have little, if any, support.

Plantinga's Response to Rowe's Argument

In response to Rowe, Alvin Plantinga claims that Rowe's argument is unsuccessful because its pattern of reasoning is fallacious. Plantinga draws our attention to what he calls 'arguments from degenerate evidence':

To give an argument from degenerate evidence, you propose to support a proposition A by showing that A is probable with respect to a part of your evidence which is such that there is an isomorphic part of your evidence with respect to which $\neg A$ is at least equally probable. (1998: 540)

Arguments from degenerate evidence are obviously unconvincing. As an example of an argument from degenerate evidence Plantinga asks us to consider the fact that you are now barefoot (B). We can easily see that the information that (not-G or B) lowers the probability of God's existing, and the information that (G or B) raises the probability of God's existing. Learning that (not-G or B) is equivalent to learning that it is false that (G and not-B). If we conditionalize on this, i.e., remove mud on the (G and not-B) area from the table, then it is clear that there is less mud on the G area than before. However, if we instead conditionalize on learning (G or B) and remove the mud on the (not-G and not-B) area, then there will be less mud on the not-G area than before. Consider the following muddy Venn diagrams where the blank area represents the area that the mud is removed from:

not-G	G
Not-B	Not-B
B	B

not-G	G
Not-B	Not-B
B	B

In the top table we conditionalize on the falsity of (G and not-B), which raises the probability of not-G, and in the bottom table we conditionalize on (not-G and not-B) being false, which raises the probability of G. These illustrate Plantinga's claim that both of these arguments are arguments from degenerate evidence and are unsuccessful. The argument that claims G is disconfirmed by (not-G or B) has a corresponding argument of equal strength that G is confirmed by (G or B).

From this we can easily see one way to construct an argument from degenerate evidence. Suppose our evidence is that some proposition Z is false. We divide the proposition Z into two partitions: one in which Z and what you are trying to support is true, and one in which Z and what you are trying to support is false. Our evidence is that Z is false, which is that both of the partitions are false. But in an argument from degenerate evidence we simply ignore the falsity of the partition in which Z and what you are trying to support is true. We conditionalize only on the falsity of the partition Z and what we are trying to prove is false. Another way of seeing this is that in partitioning the proposition Z, we are noting that Z is equivalent to $[(Z \ \& \ G) \vee (Z \ \& \ \text{not-G})]$. In learning that Z is false, we learn not-Z, which is equivalent to learning $[\text{not-(Z \ \& \ G)} \ \& \ \text{not-(Z \ \& \ not-G)}]$. In an argument from degenerate evidence we ignore one of these conjuncts, namely not-(Z & not-G) if we are trying to support atheism. We are then

left with conditionalizing on $\text{not}-(Z \ \& \ G)$, which is equivalent to $(\text{not-}Z \vee \text{not-}G)$. Notice that this is a disjunction of the evidence we learned with what we are trying to support. Thus in an argument from degenerate evidence we conditionalize only on one of the conjuncts (which is a disjunction of our evidence with what we are trying to support), and we ignore the other conjunct.

In the barefoot case, we divide the proposition that it is false that you are barefoot into two partitions: one in which you are not barefoot and God exists, and one in which you are not barefoot and God does not exist. Conditionalizing on only the first of these partitions being false lowers the probability of God's existing, and conditionalizing on only the second of these partitions being false raises the probability of God's existing. Another way of seeing this is that in the barefoot case we disjoin what we know (that you are barefoot) with what we are trying to prove, and conditionalize on it. For example, an atheist might propose to conditionalize on $(\text{not-}G \vee B)$. This has the effect of ignoring the evidence we have that it is false that $(\text{not-}G \ \& \ \text{not-}B)$. Similarly, if we were to give an argument for the existence of God by using $(G \vee B)$ as our evidence, this would ignore the evidence that it is false that $(G \ \& \ \text{not-}B)$. This illustrates that arguments from degenerate evidence essentially conditionalize on only part of the evidence we have.

Plantinga then claims that Rowe's argument is an argument from degenerate evidence. Earlier we saw that Rowe's principle P is equivalent to principle P' :

$P' \quad \text{not-}(G \text{ and } KA)$

But P' is equivalent to $(\text{not-}G \text{ or } \text{not-}KA)$. Like the barefoot argument, Rowe's argument from P conditionalizes on the disjunction of what he is trying to support with some evidence we've learned. The information P is logically weaker than the complete evidence we have, and ignores the information that it is false that $(\text{not-}G \text{ and } KA)$; in addition to knowing it is false that $(G \text{ and } KA)$, we also know it is false that $(\text{not-}G \text{ and } KA)$. This was evident from our earlier discussion, where we saw that Rowe's conclusion followed from his premise only because he stated the premise in such a way that it ignores evidence we have.

Plantinga has correctly shown that Rowe's argument has the same basic structure as the barefoot argument, in that it conditionalizes on the disjunction of our evidence with $\text{not-}G$. But this does not mean that Rowe's argument is an argument from degenerate evidence. In order for Rowe's argument to be an argument from degenerate evidence there must also be an isomorphic part of our evidence with respect to which $\text{not-}G$ is equally improbable. If the isomorphic part of our evidence is $(G \text{ or } \text{not-}KA)$, then the probability of God's not existing, conditional on $(G \text{ or } \text{not-}KA)$, would have to be less than or equal to the probability of God's existing, conditional on $(\text{not-}G \text{ or } \text{not-}KA)$. But this has not been shown.⁵ As we saw in the above muddy Venn diagrams, it will be true if $P(G \text{ and } KA)$ is less than or equal to $P(\text{not-}G \text{ and } KA)$, but Plantinga has given us no argument for that.

However, even if Plantinga has not shown that Rowe's argument is an argument from degenerate evidence, he has shown a serious problem with it. Let us say that an "argument from partitioned evidence" is an argument for H that instead of conditionalizing on our evidence that some proposition E is false, conditionalizes on only part of E being false, the part conjoined with not- H . Since E is equivalent to the disjunction $[(E \ \& \ H) \vee (E \ \& \ \text{not-}H)]$, an argument from partitioned evidence conditionalizes on only one of these disjuncts being false. Conditionalizing on the falsity of $(E \ \& \ \text{not-}H)$ is equivalent to conditionalizing on $(\text{not-}E \vee H)$, which is a disjunction of the evidence with what we are trying to support. Although this will raise the probability of H , there is a structurally similar argument for the opposite conclusion that will conditionalize on a different partition of the evidence being false, the part of E conjoined with H ; this argument conditionalizes on $(\text{not-}E \vee \text{not-}H)$. We have two partitions of the evidence we know to be false: $(E \ \& \ H)$ and $(E \ \& \ \text{not-}H)$. If we conditionalize on the first partition being false, we lower the probability of H , and if we conditionalize on the second being false, we lower the probability of not- H . These arguments may not be equally strong, which is why an argument from partitioned evidence may not be an argument from degenerate evidence. But even if it is not an argument from degenerate evidence, it is problematic because of the structurally similar argument for the opposite conclusion. The argument for lowering the probability of H is counterbalanced to some degree by the argument for lowering the probability of not- H . Furthermore, we can easily construct an argument from partitioned evidence to support anything we want.

To illustrate this, consider the following example of an argument from partitioned evidence. Suppose we have a lottery that can use either of two possible number generating machines. Machine M1 generates random numbers from 1 to 100, and machine M2 generates random numbers from 1 to 99. Since either of these machines could be used in the lottery, you assign each of these possibilities a probability of .5. Now suppose you learn that ticket number 37 won the lottery. To get the correct new probabilities of M1 and M2 being used in the lottery we conditionalize on the new information that ticket 37 won. The probability of M1 being the machine used given that ticket 37 won is lowered slightly to $99/199$, and the probability that M2 was used conditional on ticket 37 winning is raised slightly from .5 to $100/199$. But now suppose that instead of conditionalizing on our information that ticket 37 won, we create an argument from partitioned evidence by conditionalizing on the disjunction that M2 was used or ticket 37 won. This is equivalent to conditionalizing on it being false that M1 was used and some ticket other than 37 won; it ignores our evidence that it is false that M2 was used and some ticket other than 37 won. The probability that M2 was used given that M2 was used or ticket 37 won is $100/101$, which is almost 1.⁶ But surely this is not the correct probability to use in reasoning about how likely it is that M2 was used in the lottery. We saw the correct probability is close to .5 instead of being almost 1. And note that we could easily give an argument from partitioned evidence for the opposite conclusion. The probability that M1 was used, given that M1 was used or ticket 37 won, is $99/100$, which is also

almost 1. Our argument for partitioned evidence for M2 being the machine that was used is not an argument from degenerate evidence, because the isomorphic argument for the opposite conclusion that M1 was used is not quite as strong. But surely this small difference is not important. It does not matter whether we ignore evidence in an argument from partitioned evidence or in an argument from degenerate evidence; it is simply fallacious to ignore relevant evidence.

Another example is the following argument from partitioned evidence which supports theism. I do not know the population of China, but I believe it could be any number between 1.2 and 1.3 billion; I assign each of these possible populations an equal probability of .00000001. For each number that could be the population of China, I partition it into two partitions: one in which that number is the population of China and God exists, and one in which that number is the population of China and God does not exist. Now suppose I learn that the population of China is c , where $1.2 \text{ billion} \leq c \leq 1.3 \text{ billion}$. But instead of conditionalizing on this, I construct an argument from partitioned evidence that conditionalizes on (the population of China is c or God exists). If we grant Rowe's assumption that $P(G) = .5$, the probability of God existing, conditional on (the population of China is c or God exists), is .99999999. But surely this is not a convincing argument; learning the population of China should not raise the probability of God existing to almost 1.⁷

Plantinga has successfully shown that Rowe's argument is an argument from partitioned evidence. In order to show that Rowe's argument is also an argument from degenerate evidence we would have to show that the partition (not- G and KA) is at least as probable as (G and KA), which has not been done. Furthermore it is clear that my modification of Rowe's argument is neither an argument from degenerate evidence nor an argument from partitioned evidence. My modified argument conditionalizes on not- KA , and so it does not ignore our information that it is false that (not- G & KA). Nor does my modified argument disjoin what we are trying to prove with the evidence and conditionalize on that. Thus it does not have any of the problems that Plantinga raised for Rowe's original argument.

Rowe's Response to Plantinga

Rowe's first point is that his argument differs from Plantinga's barefoot argument because whether you are barefoot or not is irrelevant to whether God exists, but KA is 'evidentially relevant' to the existence of God. Suppose that by KA being evidentially relevant to G Rowe means that not- KA lowers the probability of God's existing.⁸ As we saw above, if he requires this, the analogous argument from the part of the evidence Rowe ignores will not be as strong as Rowe's argument, and Rowe's argument will not fit the definition of an argument from degenerate evidence. But his argument will still be an argument from partitioned evidence, and will be problematic for that reason.

Turning now to Rowe's main response to Plantinga's argument, we find that it is unusual in that when Rowe states his argument he appears to give a new version of his argument; the argument he now presents as his

argument is very different from any previously published version. Rowe does not clearly differentiate his original argument from various modifications of it, nor does he appear to recognize the problems with his original argument. In discussing what we learn about the goods we know of, Rowe does not give us sufficient conditions that a good must satisfy in order for God to exist and the good justify God in permitting those evils. However, he does accept certain requirements Plantinga presented as necessary conditions that a good *g* must satisfy in order for God to exist and the good justify God in permitting an evil. With respect to evils E1 and E2, those conditions are:

1. *g* is actual
2. *g* outweighs E1 and E2
3. a perfect being could not achieve *g* without permitting E1 and E2, and
4. no better world can be brought about if *g* and E1 and E2 are prevented by a perfect being.⁹ (1998: 549)

Let 'C' be the conjunction of those conditions. According to Rowe's new version of his argument, we first come to learn what he calls principle X: no known good satisfies those necessary conditions C. (1998: 550) It is important to realize the differences between X, P, and not-KA. We saw earlier that P is implied by not-G; P states that there is no good we know of that actually justifies God in permitting those evils. In contrast to this, neither X nor not-KA are implied by not-G, and they make a stronger claim than P's claim that no known good actually justifies God in permitting those evils. X and not-KA go further and claim that even if God doesn't exist, no good we're aware of would justify God in permitting those evils; the goods we know of either lack a necessary condition of or are not able to justify God in permitting those evils. Thus X and not-KA imply P, but P does not imply X or not-KA. Furthermore, X implies not-KA. If no known good satisfies C (X), it is false that there is a known good that is able to justify God in permitting those evils (not-KA). It is not the case that not-KA implies X. It is possible that no known good is able to justify God in permitting those evils (not-KA), even though some known good satisfies C (not-X). Merely satisfying the necessary conditions C does not imply that the good is able to justify God in permitting those evils.

After stating that we first come to learn X, Rowe then argues:

From this result (X) we directly infer

No good we know of justifies God in permitting E1 and E2.

Since this proposition is logically equivalent to P, it is entailed by $\neg G$, and thus lowers the probability of G." (1998: 550)¹⁰

In this passage Rowe's reasoning is quite different from what he has previously presented. He here claims that we first learn X. But instead of conditionalizing on that evidence, Rowe conditionalizes on something it logically implies, namely, P. But we've seen that P is logically weaker than X, and ignores some of the evidence included in X; P is equivalent to the dis-

junction [not-G or (G and not-KA)] and so it excludes the evidence that (not-G and KA) is false, even though this evidence is included in X (since X entails not-KA). As we saw earlier, conditionalizing on P erases the area (G and KA) off of the muddy Venn diagram, but unlike conditionalizing on X it leaves the area (not-G and KA) untouched. Some of the evidence included in X is simply ignored.

Since this version of Rowe's argument has many more steps than were initially presented it will be useful to look at it in more detail. First we get our evidence (X). Second, we show that this evidence is 'evidentially relevant' to G: $P(X/\text{not-G}) > P(X/G)$.¹¹ This step or requirement is what Rowe believes differentiates his argument from Plantinga's barefoot argument. Third, we note that this evidence (X) logically implies P. In the fourth step we then show that the conditional probability of God's existing on P is lower than the probability of God's existing. From the fact that P lowers the probability of God's existing, we conclude that the facts about evil we are aware of lower the probability of God's existing. On this new account of his reasoning we see that Rowe's original presentation of his argument presented only a portion of what he now claims is his argument; his previous presentation only presented the last part of this argument, and it ignored the previous steps. We can summarize his argument as follows:

- 1 X (premise)
- 2 $P(X/\text{not-G}) > P(X/G)$ (premise)
- 3 P (from line 1)
- 4 $P(G/P) < P(G)$ from probabilistic reasoning (see Venn diagrams)
- 5 The facts about evil we are aware of lower the probability of God's existing. (from 3,4)

There are several problems with this argument. First of all, it is an argument from partitioned evidence; it ignores some of the evidence contained in X. Second, it is important to realize how odd the structure of this argument is. One goal of Rowe's original argument was to show that our not knowing a justification for the evils in the world lowers the probability of God's existing. But if that is the goal, it is puzzling as to why Rowe even mentioned principle P in his argument. Rowe's conclusion follows immediately from premises 1 and 2. Premise 2 is equivalent to $P(G/X) < P(G)$, which says that X lowers the probability of God's existing; thus there is no need to bring in principle P. And coupled with the assumption that the initial probability of God's existing was .5, Rowe can infer his conclusion that X lowers the probability of G to less than .5. In a moment we will discuss a possible motivation for bringing in principle P, but it is important to see that it is not needed to derive his stated conclusion. Rowe could have ended his argument after the second premise, since the first and second premises imply his conclusion. Furthermore, this argument from premises 1 and 2 would not have been an argument from partitioned evidence. Notice also that the argument from premises 1 and 2 is very similar to the modified argument.

However, looking at the argument one immediately notices that after

being stated, premise 2 is not used in the remainder of the argument; Rowe does not use premise 2 to derive the conclusion. Principle P is derived from premise 1, and Rowe's conclusion follows from P and premise 4; no use is made of premise 2. Rowe only brought in premise 2 in response to Plantinga's claim that Rowe's original argument was an argument from degenerate evidence. Premise 2 implies that Rowe's argument is not an argument from degenerate evidence, so perhaps we should view premise 2 as an assumption that is necessary to prevent this argument from being an argument from degenerate evidence. Although Rowe could have used this assumption in place of P to present an argument that is not an argument from partitioned evidence, he does not do so and the argument he presents is still an argument from partitioned evidence.

Although Rowe's discussion of this argument hints at the modified argument I presented, he does not appear to see the advantages of it over both his original and new argument. The new argument he presents is clearly not the same as my modified argument, since his is still an argument from partitioned evidence. Even though some of Rowe's remarks may make it seem as if he has the modified argument in mind, other remarks cast doubt upon this. For example, in his response to Plantinga's criticisms Rowe notes that his argument "can be taken as an argument from X alone," (1998: 551) which is basically a version of the modified argument. But immediately following this, Rowe's comments cast doubt upon whether he had the modified argument in mind, or simply his original argument. Unlike the modified argument, the evidence (P) conditionalized on in Rowe's original argument is implied by not-G, and thus it ignores the part of our evidence that is not implied by not-G. Whether or not the evidence conditionalized on is implied by not-G is crucial because this determines whether X ignores evidence we have. But in this discussion Rowe is not clear about this crucial point, namely, whether not-G implies X. He writes: "It should be noted, however, that if $\neg G$ does not logically imply X, then even though X reduces the likelihood that God exists, it won't reduce it to quite the degree that P does. For P is logically implied by $\neg G$." (1998: 552) Rowe's original argument was an argument from partitioned evidence; since not-G implied P, P excluded certain evidence, namely, the falsity of (not-G & not-KA). If Rowe interprets X in such a way that it also is implied by not-G, then this argument will ignore similar evidence. One way to construe X so that it is implied by not-G would be to change condition 3 of C to:

- 3'. g cannot exist unless E1 and E2 are permitted to exist by a perfect being.¹²

If we let C' be the conjunction of conditions 1, 2, 3', and 4, and let X' be the claim that no known good satisfies the necessary conditions in C', then not-G implies X'.¹³ But since X' is implied by not-G, it ignores evidence we have; X' will not include that is false that both God does not exist and we know of a good that satisfies the necessary conditions in C. We first learned that X is true; from this we might infer that X' is true. But if we ignore X and focus on X', we ignore the evidence that is in X but not in X'.

Thus this argument will fail for the same reason that Rowe's original argument failed; it ignores relevant evidence we have. For this reason it is important to not base the argument on X' and to interpret X as not being implied by not- G ; otherwise the argument based on X is no better than the original argument. It would be fallacious because it would ignore relevant evidence.

Given this interpretation of X and given that the first two premises state that some of our evidence about evil lowers the probability of God's existing, the question arises as to why Rowe is concerned to show that P lowers the probability of G , and why he even used P in his argument. The answer is that Rowe appears to think that the amount that something implied by X , namely P , lowers the probability of G is interesting and important. (1996: 272-276) Assuming X is not implied by not- G , Rowe is correct that P will lower the probability of G more than X lowers it. Thus an argument based on P will appear to be stronger and lower the probability of G more than an argument based on X . But whether any logical consequence of what we know lowers the probability of G , and the degree to which it may lower that probability, is irrelevant to what it is rational to believe. We can easily find consequences of X that lower the probability of G much less than P does. Rowe himself notes that $(G \text{ or } X)$ is a consequence of X , and conditionalizing on that would raise the probability of G . (1998: 551). Arguments that conditionalize on consequences of our evidence X are arguments from partitioned evidence; they conditionalize on only some partition of our evidence and ignore other partitions. But clearly arguments from partitioned evidence are fallacious, and it is a mistake to think that the probability of G conditional on some partition of our evidence is relevant to what it is rational to believe.

We can easily see the fallaciousness of this reasoning, and all arguments from partitioned evidence, if we apply the same reasoning to a different example. Suppose we have an urn with red and black balls, and there are two possible distributions of balls in the urn, each with a probability of .5 (suppose a coin is flipped that determines which of these hypotheses describe the makeup of the urn). According to H_1 there are 990 red balls and 10 black balls, and according to H_2 there are 989 red balls and 11 black balls. Now suppose our evidence E is that we draw a black ball out of the urn. This evidence slightly lowers the probability of H_1 : $P(H_1/E) = 10/21$, and so the probability of H_1 is lowered from .5 to $10/21$. Now suppose someone were to point out that E implies $(E \text{ or } H_2)$, and this lowers the probability of H_1 even more than E does; it is easy to calculate that $P(H_1/E \text{ or } H_2) = 1/101$.¹⁴ The information $(E \text{ or } H_2)$ lowers the probability of H_1 much more than E lowers it. Indeed, the ratio of $P(H_1/E)$ to $P(H_1)$ is .9523; E lowers the probability of H_1 by about 5%. The ratio of $P(H_1/E \text{ or } H_2)$ to $P(H_1)$ is .0198; $(E \text{ or } H_2)$ lowers the probability of H_1 by about 98%. Clearly there is a big difference in the effect of E and of $(E \text{ or } H_2)$ on the probability of H_1 ; lowering the probability of a hypothesis by 98% is very different from lowering it by a mere 5%. But it is also obvious that the correct probability to use in this instance is $P(H_1/E)$, and that $P(H_1/E \text{ or } H_2)$ is irrelevant. Since we are interested in what the rational degree of belief in H_1 is after learning E it is clear that we should condition-

alize on E, and not conditionalize on (E or H2). Even though E implies (E or H2), this later information is not relevant to what it is rational to believe because it ignores evidence we have; it ignores the information that it is false that (not E and H2). The information E excludes the possibility of (not-E and H2), but (E or H2) is consistent with that possibility. In order to exclude this possibility that we know is false, we need to conditionalize on E, our evidence, instead of something that our evidence implies. Arguments from partitioned evidence are clearly fallacious and are not relevant to what it is rational to believe.¹⁵ Although I usually hate to gamble, if anyone seriously thinks the probability $P(H1 / E \text{ or } H2)$ is relevant to rational belief and is willing to bet on their beliefs, I would like very much to get some urns and make some bets with them; I've always wanted to buy a boat.

The reasoning in this example and in our earlier example with the lottery machines is identical to the reasoning in Rowe's argument; all of these are arguments from partitioned evidence. As in the urn example, we have two mutually exclusive hypotheses each with a probability of .5. We also have some evidence (X) that lowers the probability of one hypothesis (G), and we also know that (X or not-G) lowers the probability of G even more. But just as in the urn example, the fact that (X or not-G) lowers the probability of G, and the degree to which it lowers it, is not relevant to the strength of the disconfirmation of G by X, or to what our rational degree of belief in G should be. Thus we see that Rowe's more complicated argument provides no support for his original argument, and the modified version of his argument is the strongest version we have discussed.

Problems with the Modified Argument

We have seen that my modified version of Rowe's argument is the only version of his argument that is not fallacious. I now want to discuss two problems with this version: the assumption that it is more likely that we'd know of a justification for these evils given that God exists than given that God does not exist, and the general strategy and epistemic framework that Rowe's arguments assume. Although a complete discussion of how probability should be used in discussions of this sort is far beyond the scope of this paper, we will briefly discuss the assumptions Rowe makes in this regard.

Earlier we saw that in order to save Rowe's argument we needed to modify it and require as the central assumption that $P(KA/G) > P(KA/\text{not-G})$. Although Rowe does not seem aware of the importance of this assumption, in his discussion of his new argument he gives the following reason for thinking that $P(X/\text{not-G}) > P(X/G)$. Rowe thinks that if God does not exist, X is more likely:

One reason this is so is that no good *involving* God, such as the little girl's enjoying eternal felicity in the presence of God or (Plantinga's suggestion) the little girl enjoying God's gratitude in eternal felicity, is even a candidate for consideration, since such a good is *actual* only if God does exist. But if God does exist, these goods are at least can-

didates for the job of justifying him in permitting the little girl's suffering. (1998: 550)

Rowe is claiming that if God exists it might be true that if the girl suffers, she may eventually receive a much greater good from God. Evidently he thinks that if God exists there are more chances for there to be a good that is able to justify these evils than there are if God does not exist. But it is far from clear how this is supposed to support $P(X/\text{not-G}) > P(X/G)$, or $(P(KA/G) > P(KA/\text{not-G}))$. Suppose Rowe is correct and there are more candidates for there being a good able to justify God in permitting those evils if God exists; we might then conclude that conditional on God's existing it is more likely that there exists a good able to justify God in permitting those evils. But from this we cannot infer that conditional on God's existing it is more likely that we would *know* of a possible good that would justify God in permitting those evils. It is important to realize that throughout this discussion Rowe is referring to *possible goods*, and not merely to goods that we know to exist. In discussing what goods he is referring to, Rowe says "I mean to include goods that we have some grasp of, even though we have no knowledge at all that they have occurred or ever will occur." (1996: 264) This is important for at least two reasons. First, even if God does not actually exist, we could still know of a possible good involving the existence of God that would justify God in permitting those evils, if God were to exist. Second, even if God does not actually exist, we do not know that is the case. For all we know various goods involving God's existence, such as the girl enjoying eternal felicity in the presence of God, could justify God in permitting those evils. Thus even though these goods involve the existence of God, we need to consider them when we investigate how likely it is that we would know of a good able to justify God in permitting those evils, conditional on the nonexistence of God. It is conceivable that God does not exist, we do not know that God does not exist, and yet we know of some possible good *G* involving the existence of God that would justify God in permitting those evils if God were to exist. So even if God's existence implies there are more candidates for goods able to justify God in permitting those evils, it does not follow that we are more likely to know of those candidates if God exists than if he doesn't exist. But we are interested in how likely it is we'd know of a good able to justify God in permitting those evils, not in how likely it is that one exists; thus Rowe's brief argument does not support $P(KA/G) > P(KA/\text{not-G})$, which is the central assumption of the modified argument.

If we restrict ourselves to possible goods that do not involve the existence of God, we also find that there is no reason to believe that the existence of God makes it more likely that we'd know of a good that is able to justify God in permitting those evils. On the contrary, it may even be the case that the existence of God makes it less likely that we'd know of a good that is able to justify God in permitting those evils. For example, suppose some sort of cosmic karma is true, and if the girl in Rowe's example suffers horribly now she will eventually receive a much greater good. It is conceivable that it is more likely that we'd know about this cosmic karma if God does not exist than if he exists. One can easily imagine all sorts of sit-

uations to illustrate this possibility. For example, if God did not exist, perhaps Zeus, Athena, Baal, or some other being would exist and would be more likely than God to tell us about this cosmic karma. In this case our not knowing a good that is able to justify God in permitting those evils would raise the probability of God existing, because we are more likely to know of such a good if God does not exist. The crucial point is that we have been given no argument to show that the existence of God makes it more likely that we'd know about possible goods that are able to justify God in permitting those evils. The only discussion of this in the literature is the comment by Rowe quoted above, and much more investigation is needed into this.¹⁶ For this reason we have no reason to believe that $P(KA/G)$ is greater than $P(KA/\text{not-}G)$. But if we should withhold judgement on this, then we withhold judgement on the crucial premise in the modified version of Rowe's argument. We then have no reason to accept the conclusion of the modified argument, because we do not know if its central assumption is correct.

A second problem with the modified argument is that although this argument is not an argument from partitioned evidence, even this argument does not treat all the evidence we have equally; it gives the evidence of evil that Rowe focuses on special treatment and does not consider other relevant evidence equally. This special treatment of evidence is built into Rowe's methodology. It is well known that probability is highly dependent upon the background information. For the background information in this argument, Rowe proposed that we take beliefs common to all parties in the dispute and subtract from it our belief that we know of no justification for the evil in the world. This background evidence will contain all of the other evidence for and against the existence of God that theists and non-theists agree on. Given this we assign $P(G) = .5$ with respect to this background information, even though neither theists nor non-theists believe the probability of God's existing is .5, and even though this doesn't reflect the strength of the evidence for or against the existence of God contained in this background information. We then generate conditional probabilities of this evidence conditional on God's existing and God's not existing. It is not clear how we are supposed to generate the probability of $(G \text{ and } KA)$, and of $(\text{not-}G \text{ and } KA)$ with respect to this background information. Assuming Rowe is correct, both the theists and atheists agree that KA is false; but that makes it very difficult to see how we get those conditional probabilities.¹⁷ Assuming we can justify generating some values for these probabilities, we then see what the effect of conditionalizing on the evidence (that we previously subtracted from our beliefs) has on the probability of God's existing. From this Rowe concludes that it is irrational to believe in God unless we have strong evidence for the existence of God. It is clear that this method focuses only on certain evidence we have, and it is difficult to see why this evidence should get this special treatment. If it is permissible to single out and focus only on some evidence we have, then we can easily develop an analogous argument to support any position we want.

For example, we are aware of other evidence relevant to the existence of God; besides our not knowing a justification for God to permit those evils, we know that some people believe in God, some attend church, some

act altruistically, and that some feel a deep impulse to believe in God. I am not claiming that the phenomena in this list cannot be accounted for on the assumption that God does not exist, but one can plausibly claim they are more likely given that God exists.¹⁸ Following Rowe's methodology, we could get our background information by subtracting this evidence from the beliefs theists and non-theists have in common. The facts about evil Rowe discusses would be in this background information, and we would assume that the probability of God existing conditional on this background information was .5 (even though the background information contained all the facts about evil that Rowe mentioned). We might then notice that this evidence raises the probability of God's existing relative to this background information, and conclude that it is irrational to disbelieve in God unless we have strong evidence against the existence of God. This argument has the exact same structure as Rowe's argument, and illustrates that by focussing only on certain evidence we can construct an argument to support any position.

From this we see the difficulty in drawing any interesting conclusions from arguments that focus on only part of our evidence. Since Rowe's argument focuses on evil and does not treat our other evidence equally, it is difficult to see what epistemic conclusions follow from it. A successful probabilistic argument against the existence of God would need to consider and assess all of the relevant evidence.¹⁹ To summarize, Rowe's methodology is not obviously correct, and on the contrary, it should be rejected unless good reasons can be given for accepting it. But since no defense of it has been given, we should not accept the general epistemic framework that his argument assumes.

Conclusion

In conclusion we have seen that Rowe's original argument is fallacious because it ignores evidence that we have; it is an argument from partitioned evidence and only looks at a certain partition of our evidence. We also saw that Plantinga's claim that Rowe's argument is an argument from degenerate evidence depends upon whether a certain probability relation is false. As a result, Plantinga did not show that Rowe's argument is an argument from degenerate evidence, although his discussion shows Rowe's argument is an argument from partitioned evidence. Rowe's newer version of his argument simply adds an assumption, which prevents it from being an argument from degenerate evidence. But this new argument is still an argument from partitioned evidence, and it is fallacious to think the degree of confirmation in an argument from partitioned evidence is relevant to what it is rational to believe. I presented a modified version of Rowe's original argument which is neither an argument from degenerate evidence nor an argument from partitioned evidence. This argument is simpler than Rowe's arguments and is the strongest 'Rowe-type' argument from evil that has been presented. However this argument relies on a controversial assumption, and I argued that we do not know the probabilities involved in it. Since it is rational to withhold judgement on those probabilities, this modified argument will not be successful against

those who do not accept the assumption. However, with the exception of a very short passage by Rowe, this crucial assumption has not been discussed in the literature, and I expect future debate to focus on it. We then saw that the way probabilities are used in the epistemic framework Rowe assumes is quite controversial. Without a defense of his methodology there is no reason to think the pattern of reasoning in these arguments is legitimate. Future discussion of the probabilistic problem of evil will need to address these issues.²⁰

University of California at Santa Cruz

NOTES

1. William Rowe, "The Evidential Argument from Evil: A Second Look," in Daniel Howard-Snyder, ed., *The Evidential Argument From Evil* (Bloomington, IN: Indiana University Press 1996) pp. 262-285.

2. Alvin Plantinga, "Degenerate Evidence and Rowe's New Evidential Argument from Evil," *Noûs* 32 (1998), pp. 531-544, and William Rowe, "Reply to Plantinga," *Noûs* 32 (1998), pp. 545-552.

3. Rowe (1998) doesn't discuss his original formulation of P, but instead presents the following as P:

P: God does not exist or (God exists and no known good justifies him in permitting E1 and E2)

This is logically equivalent to his original P and to P'.

4. See Rowe (1996: p. 266). Here and throughout this paper I have deleted Rowe's reference to background information from the probability formulae in order to make them easier to read. The probability formulae used in this paper should always be understood as being relative to some specified background information. The background information is the basis of some interesting discussion by Plantinga (1998), but I will not discuss these issues in this paper.

5. Plantinga (1998: 542) does claim that conditionalizing on (G or no known good is known to justify E1 or E2) provides a counterbalancing argument of equal strength to Rowe's argument. But this depends on the assumption that (there being no known good known to justify E1 or E2) is equally likely on G and on not-G. This assumption is not stated, nor are we given any reason to accept it.

6. $P(M1)=P(M2)=.5$, $P(37 \vee M2)=101/200$, $P(37 \vee M2/M2)=1$. We can put these values into a version of Bayes' theorem to get: $P(M2/37 \vee M2) = P(M2)P(37 \vee M2/M2) / P(37 \vee M2) = 100/101$. We also know that $P(37 \vee M1) = 100/198$, and thus $P(M1/37 \vee M1) = P(M1)P(37 \vee M1/M1) / P(37 \vee M1) = 99/100$.

7. Consider another example: suppose there are 9 planets, but we realize that for any integer n, God could have made n planets. Let us assume that on the assumption that God exists, every natural number has the same equal probability of being the number of planets that God made. To represent this situation we need to partition the areas representing G into an infinite number of equal areas, each representing the probability of God making exactly n planets, for each n. Each of these partitions will receive a probability less than any finite number. Now suppose we learn that God made exactly 9 planets. This

implies: God did not make exactly 0 planets, God did not make exactly 1 planet, God did not make exactly two planets, ..., God did not make exactly 8 planets, God did not make exactly 10 planets, God did not make exactly 11 planets, If we take account of this evidence, we find that the probability of God's existing given that he made exactly 9 planets is less than any positive finite number. By carefully choosing what evidence to focus on we lowered the probability of God's existing to 0 or an infinitesimal, even though the evidence we conditionalized on implies that God exists.

8. If Rowe allows KA to be evidentially relevant to G if KA raises the probability of G, then his argument will be an argument from degenerate evidence, because there is an argument from part of the evidence that disconfirms not-G even more than Rowe's argument disconfirms G. In order for Rowe to claim that his argument is not an argument from degenerate evidence, he must require not only that KA be evidentially relevant to G, he must further require that KA be negatively relevant to G: $P(KA/\text{not-G}) < P(KA/G)$.

9. Although Rowe's conditions 1, 2, and 4 are equivalent to the ones Plantinga presented (1998: 541), Rowe's condition 3 is slightly different. The version stated in the text is Plantinga's version of this condition. Although Rowe's condition 3 may look equivalent to Plantinga's condition 3, there is a subtle but important difference between them. Rowe's condition 3 is:

g cannot exist unless E1 and E2 are permitted to exist by a perfect being.

Since it appears that Rowe intended to adopt Plantinga's conditions, I will use his version of condition 3 instead of the version that Rowe presented. Later I will discuss Rowe's version of condition 3. As we shall soon see, the choice of which version of condition 3 is used determines the type of problem that his argument faces.

10. In a footnote to the proposition directly inferred from (X), Rowe says: "... this proposition is logically equivalent to P: God does not exist or (God exists and no known good justifies him in permitting E1 and E2)." (1998: 552)

11. Rowe is not clear about what he means by 'evidentially relevant'; he says "...what we learn initially in this reasoning is something (X) that is evidentially very relevant to the question of whether God exists....(X) itself makes $\neg G$ more likely than G. For X is more likely given $\neg G$ than it is given G.... So for this reason and others, X is much more likely if God does not exist than it is if God does exist." (1998: 550) It initially appears that Rowe is claiming that $P(\text{not-G}/X) > P(G/X)$ on the basis of $P(X/\text{not-G}) > P(X/G)$. However, this in general is not a correct inference, and the conclusion implies nothing about whether X is evidentially relevant to G; it could be true even if X is irrelevant to G and does not affect the probability of G or of not-G. However, if we assume that $P(G) = P(\text{not-G}) = .5$ this inference holds, and the conclusion would be relevant. Later in the quoted passage Rowe appears to adopt a more traditional account of evidential relevance, whereby X is evidentially relevant to G if $P(X/\text{not-G}) > P(X/G)$. This tells us that the probability of X is lowered by G, which implies that X lowers the probability of G. I will use this second interpretation of Rowe's claim in my argument, because it is the standard account of evidential relevance and does not require that $P(G) = P(\text{not-G}) = .5$.

12. Condition 3' is the third condition that Rowe actually presented (1998: 549), and thus Rowe actually proposed X' as instead of X. As I stated in note 9, there is reason to believe Rowe intended to present condition 3 instead of 3', and that is why I initially discussed X instead of X'. Furthermore, the argument based on X' is immediately problematic because it ignores relevant evidence. However, it is still unclear as to whether Rowe was endorsing X or X'.

13. Condition 1 says that *g* is actual, and condition 3' says that *g* cannot exist unless *E*1 and *E*2 are permitted to exist by a perfect being. But a perfect being cannot permit *E*1 and *E*2 to exist unless it exists, and so *C*' implies that a perfect being exists. Assuming that a perfect being is God, if God does not exist, conditions 1 and 3' cannot be satisfied. Thus not-*G* implies that no known good satisfies *C*', which means that not-*G* implies *X*'.

14. $P(H1) = P(H2) = .5$. $P(EvH2/H1) = 10/1000$. $P(EvH2/H2) = 1$. Plug these values into: $P(H1/E \vee H2) = P(H1)P(EvH2/H1) / [P(H1)P(EvH2/H1) + P(H2)P(EvH2/H2)] = 1/101$.

15. Our previous example involving the lottery with machines *M*1 and *M*2 also illustrates this point well. In that example the correct probability to use in reasoning about how likely a certain machine was used remained close to .5, even though it was raised to almost 1 by the argument from partitioned evidence.

16. One might attempt to support this by developing an argument based on Rowe's claim that $P(P/G)$ is low. I do not think this strategy will succeed, and leave the details to the reader.

17. The modified argument uses the conditional probabilities $P(KA/G)$ and $P(KA/\text{not-}G)$, which are equal to $P(KA \& G)/P(G)$ and $P(KA \& \text{not-}G)/P(\text{not-}G)$ respectively. Thus our use of those conditional probabilities depends upon our having values for $P(KA \& G)$ and $P(KA \& \text{not-}G)$. But both theists and non-theists agree that both of these probabilities are equal to 0, because we know that *KA* is false; since $P(KA) = 0$, $P(G \& KA) = P(\text{not-}G \& KA) = 0$. As a result, the conditional probabilities are also equal to 0. In order to avoid this problem Rowe proposes that we remove the information that $P(KA) = 0$ from our background knowledge. Presumably we also remove $P(G \& KA) = 0$ and $P(\text{not-}G \& KA) = 0$ from our background knowledge as well. But then the question of what value to assign to $P(G \& KA)$ and $P(\text{not-}G \& KA)$ relative to our background knowledge arises. In order to get the conditional probabilities that the modified argument relies on we need to assign values to these probabilities relative to the new background knowledge. Any choice of values to assign to these probabilities is not determined by the background knowledge, and must be based on other considerations. Since these considerations are not included in the background knowledge, the assignment of values to these probabilities appears arbitrary. Furthermore, since these considerations are not included in the background information they are not agreed upon by both theists and non-theists. This is a version of the well known problem of old evidence. Interestingly, I know of no proposed solution to this problem that supports the procedure Rowe is proposing.

18. I find it difficult to assign values to these conditional probabilities, as well as to the ones Rowe uses involving evil. But if Rowe is correct that our awareness of evil lowers the probability of God's existing, I suspect the impulse to believe in God would raise that probability. If one wishes to withhold judgement on these probabilities they should also withhold judgement on the probabilities in Rowe's argument.

19. Rowe might argue that he is not ignoring evidence, because other evidence is allowed to come in at a later stage as a possible defeater for this argument. For problems with this approach see Richard Otte, "Evidential Arguments from Evil," *International Journal for Philosophy of Religion* 48 (2000), pp. 1-10.

20. I would like to thank members of Notre Dame's Center for Philosophy of Religion discussion group, William Hasker and the very diligent referees for *Faith and Philosophy*, Michael Bergman, Michael Peterson, David Hemp, and especially Alvin Plantinga and William Rowe, for very helpful comments on various drafts of this paper.