

# Sceptical Theism, the Butterfly Effect and Bracketing the Unknown

ALEXANDER R. PRUSS

## Abstract

Sceptical theism claims that we have vast ignorance about the realm of value and the connections, causal and modal, between goods and bads. This ignorance makes it reasonable for a theist to say that God has reasons beyond our ken for allowing the horrendous evils we observe. But if so, then does this not lead to moral paralysis when we need to prevent evils ourselves? For, for aught that we know, there are reasons beyond our ken for us to allow the evils, and so we should not prevent them. This paralysis argument, however, shall be argued to rest on a confusion between probabilities and expected utilities. A connection between this paralysis argument and Lenman's<sup>1</sup> discussion of the butterfly effect and chaos will be drawn, and the solution offered will apply in both cases.

## 1 Introduction

Weak sceptical theism (wst) holds that the existence, severity and distribution of observed evils does not noticeably decrease the probability of the existence of God given the massive extent of our ignorance about (a) the realm of value and (b) the connections, causal and modal, between goods and evils. This ignorance, it is claimed, makes it impossible to be confident that a perfect being would have done better to prevent an evil *E*, say because *E* might be necessary for some great good beyond our ken. The thesis that the existence, severity and distribution of observed evils do not *at all* decrease the probability of the existence of God because it is no more likely than not that God, if he existed, would on balance have reason to prevent *E* I will call "strong sceptical theism" (sst).

A presupposition of the discussion is that prior to the consideration of evil, the probability of the existence of God is neither 0 nor 1, and I shall assume this for the sake of the arguments. After all, if the probability of the existence of God is 0 or 1, then it won't be affected—at least in a Bayesian setting—by inductive evidence.

<sup>1</sup> James Lenman, 'Consequentialism and Cluelessness', *Philosophy and Public Affairs* 29 (2000), 342–370.

It is important not to overstate either weak or strong sceptical theism. The sceptical theist does not claim to be ignorant in general about what is and is not good or bad, or about ordinary causal and modal connections between goods and bads. She knows that poisoning someone causes death, and that death is something bad. But she denies that the known goods, bads and interconnections are likely to be representative of all of the ones there. For instance, for aught that she knows, while the death of a healthy innocent person is something bad, such a death might be, either in general or in a particular case, a necessary partly constitutive condition for some immense good.

It has been argued that sceptical theism leads to moral paralysis in the case of prevention of evil (e.g., Almeida and Oppy 2003). A standard formulation of this argument is this. Start with the story of Ashley Jones, a twelve-year-old girl raped and bludgeoned to death. We then imagine we could, with no danger to ourselves, have stopped this evil, and then we argue like this:

1. If [strong sceptical theism] 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 to prevent Ashley's suffering.
3. So, [strong sceptical theism] is false.<sup>2</sup>

The argument is valid and nobody in the debate denies (2). Thus, the discussion focuses on (1). One response to this influential type of argument is to claim that premises like (1) unduly depend on consequentialist reasoning (for the latest defense of this, see Daniel Howard-Snyder<sup>3</sup>). I shall argue that the extent of consequentialism that is needed for arguments like the above is no greater than the extent to which consequentialism is in fact true, and hence this response is inadequate. Instead, I will argue that the intuitions supporting (1) seem to make the mistake of using probabilities in practical reasoning where expected utilities are needed, and offer a very plausible consequence-focused *prima facie* moral principle that undercuts

<sup>2</sup> Cf. Daniel Howard-Snyder, 'Epistemic Humility, Arguments from Evil, and Moral Skepticism', *Oxford Studies in Philosophy of Religion* 2 (2010), 1–35.

<sup>3</sup> 'Agnosticism, the Moral Skepticism Objection, and Commonsense Morality', in: T. Dougherty and J. P. McBrayer (eds.), *Sceptical Theism: New Essays* (Oxford: Oxford University Press, 2014).

## Sceptical Theism, the Butterfly Effect

(1). If so, the argument against sceptical theism fails even on the consequentialist grounds on which it was made.

As noticed by Howard-Snyder<sup>4</sup>, the debate here parallels the argument in Lenman that the “chaos”, or great sensitivity to initial conditions, that we observe in the human world around us undercuts consequentialism. Along the way, I shall draw out this parallel, and my solution will also be a response to Lenman.

Strong sceptical theism is rather implausible. After all, suppose that we observed no evil at all. That would surely be evidence for the existence of God. But it is a Bayesian theorem that if a hypothesis  $H$  has a probability that's neither 0 nor 1, and  $D$  is evidence for  $H$ , then its negation,  $\sim D$ , is evidence against  $H$ . Hence, if the non-observation of evil is evidence for theism, then the observation of evil is evidence against theism. But it might be quite insignificant evidence for theism, and that is what the sceptical theist probably should say instead of saying that it is no evidence at all. Hence, *wst* is preferable to *sst*. However, in responding to the paralysis argument, I will work with *sst*, since the paralysis argument is more compelling on *sst* than *wst*. If my response on behalf of *sst* is successful, it is even more successful on behalf of *wst*.

I need to offer an important caveat, however. It has been argued that sceptical theism leads to scepticism *simpliciter* (e.g., Pruss<sup>5</sup>). If so, in particular it leads to scepticism about morality, and hence to moral paralysis. I shall assume for the sake of this paper that the argument that sceptical theism leads to scepticism in general fails, even though I am actually quite sympathetic to that argument.

### 2 An argument and a fallacy of moral reasoning

Why think that (1) is true? There is an intuitive line of thought:

4. If *sst* is true, it is no more likely that preventing Ashley's suffering will lead to an on-balance better outcome than that it will lead to an on-balance worse outcome.
5. If it is no more likely that preventing Ashley's suffering will lead to an on-balance better outcome than that it will lead to

<sup>4</sup> Op cit. (2010).

<sup>5</sup> Alexander R. Pruss, 'Sceptical Theism And Plantinga's Evolutionary Argument Against Naturalism' (2010) [http://prosblogion.ektopos.com/2010/05/22/sceptical\\_theism\\_and\\_platingas\\_evolutionary\\_argument\\_against\\_naturalism](http://prosblogion.ektopos.com/2010/05/22/sceptical_theism_and_platingas_evolutionary_argument_against_naturalism).

## Alexander R. Pruss

an on-balance worse outcome, consequence-based reasoning does not support preventing Ashley's suffering.

6. If consequence-based reasoning does not support preventing Ashley's suffering, then we should be in doubt whether we should have intervened to prevent Ashley's suffering.

And (1) follows immediately from (4)–(6) by a sequence of hypothetical syllogisms.

Howard-Snyder<sup>6</sup> challenges (6) by opting for a non-consequentialist moral theory. But it is important to note that our moral reasoning requires a significant component of consequence-based reasoning. Suppose two famine-relief organizations serve distant strangers, and the means they employ are morally on par. My decision which organization to give money to should be based precisely on answers to questions about consequences: How many lives can be saved by a donation of that size, how much suffering can be alleviated, and so on? And if an angel were to reveal to me that unbeknownst to me there are unfortunate side-effects of donating to both organizations that cancel out the benefits of the donation, then there would be no point to giving to either.

There will be versions of the Ashley case where (6) is false, say when we have a special relationship with Ashley, such as being a friend or a relative. It may even be the case that when Ashley is a stranger and turns her terrified eyes to us for help, that will constitute such a special relationship.

But suppose Ashley is a distant stranger and the only connection I have with her is that I know that if I press a button, her attacker will be teleported out of her home. Then consequence-based reasoning seems exactly the right kind of reasoning. For suppose I know for sure that the attacker will perpetrate the same heinous act on some other innocent stranger if I teleport him. Then, I argue, I should be at least in doubt whether I should intervene.

For, on the one hand, it could be intrinsically valuable to stand against rape and murder by teleporting the attacker. I would thereby be doing something like 'making a statement', maybe even if no one hears the statement. This provides a reason to teleport the attacker.

On the other hand, there is something deeply morally uncomfortable about deciding by a positive act which of two complete strangers will be raped and murdered. Suppose a trolley is heading for a fork in the track, and the switches are set so it will turn left. There is

<sup>6</sup> Op. cit. (2014).

## Sceptical Theism, the Butterfly Effect

complete stranger lying on the leftward track and another on the rightward track. I could move the switch on the track so the trolley turns right, sparing the stranger on the left, but killing the one on the right. But this seems worse than pointless. Suppose I redirected the trolley to the right, and the mother of the stranger asked me why I killed her daughter. What would I say? 'This was an unintended side-effect of saving the life of the person on the left track.' But the next question could surely be: 'Why was the person on the left track more important to you?' If I said that the person on the left track was *my* daughter, I could imagine being understood. But if I said that they were both equally strangers, the mother's resentment at my meddling could be justified. On the other hand, if I left the trolley alone, I could imagine saying to the mother of the victim on the left track: 'I did nothing, because either way someone would die, and it was not my place to influence who that would be.'

In the case where Ashley is a distant stranger and teleporting the attacker leads to another attack, I started with an intuition that I should teleport the attacker out of Ashley's home. But upon imagining what I would have to say to a parent of the other victim, I lose that intuition.

In any case, it is not my point to argue that in that case I *shouldn't* save Ashley. It is only my point to argue that we should be in doubt when the consequences are balanced and Ashley is a stranger. Thus, when Ashley is a total stranger, (6) is rather plausible. The other premises of the argument for (1), as well as the argument (1)–(3), remain plausible when Ashley is a total stranger. Hence anti-consequentialism is not the way out of the argument.

The problem instead is that (5) is false. It is a fallacious form of consequence-based reasoning.

For suppose there are two deep swimming pools. You know there are ten people drowning in one of the pools, and none in the other. The pools are shrouded in mist and there is a lot of background noise, so you can't tell which pool is which. Suddenly you see a malfactor toss a fair coin, and on the basis of the coin toss decide into which pool to throw a tied-up innocent stranger. Despite the mist, you can see which pool the stranger is being thrown into. You can now press a button that will drain exactly one pool of your choice. (If you try to drain both, everything explodes.) Should you drain the pool with the tied-up stranger (pool *A*) or the one without (pool *B*)?

The probability that pool *A* has more people than pool *B* is 1/2: pool *A* has the tied-up stranger, but that won't make any difference to which pool has more people, since prior to the stranger's being

tossed in, there were ten in one pool and none in the other. Likewise, the probability that pool *B* has more people than pool *A* is also  $1/2$ . Thus, the probability that draining pool *A* produces the better result is  $1/2$ ; this is the same as the probability that draining pool *B* produces the better result. If an analogue of (5) held in this case, consequence-based reasoning would not support draining pool *A*.

But consequence-based reasoning clearly does support draining pool *A*: draining *A* has a better expectation. This is intuitively true, but can also be checked with an expected value calculation. The expected value of draining pool *A* is the saving of  $1+(1/2)(10) = 6$  lives. The expected value of draining pool *B* is the saving of  $(1/2)(10) = 5$  lives.

The problem with (5) is that consequence-based decisions should not be made on the basis of which action is more likely to be beneficial. They should be made on the basis of something more like an expected utility calculation.

While the pool case will be useful later, thinking about cases of safety checks makes the point even more forcefully. Sally gives tours of the crypt of an old church. At the end of every day, she checks to make sure that no one has been left behind in the crypt when she locks up for the night. It would be quite unpleasant to be locked up in a crypt for the night. Past experience shows her that she finds a person left behind in the crypt once every ten days. It's evening now, and she wants to go home. Sally's son has a minor illness, and she would like to get back to him a few minutes earlier, so she is thinking about skipping the check. She reasons:

If there is no one in the crypt, going home without checking is better, since it's good for my son that I get home sooner. If there is someone in the crypt, checking is better, since it's terrible to be locked up for the night in the crypt. Since the probability that there is someone in the crypt is only  $1/10$ , it is much more likely that not checking leads to the better outcome.

But clearly she should check. And while one might reasonably argue for this on the grounds of her special duties to tourists, it remains true that pure consequence reasoning also gives that answer. Being left in a crypt overnight is much more than ten times as bad as having one's mom come home a few minutes later when one is a little sick.

Safety check procedures often violate analogues of (5). It is morally worth making a small effort—even one that imposes small but morally significant costs on others—to prevent an unlikely great bad, even though most likely the better outcome will eventuate if one skips the small effort. In the case of safety checks, one expects

## Sceptical Theism, the Butterfly Effect

that analogues of (5) in fact are quite a serious blight on society. In any instance of a safety check, one can be pretty confident that a better outcome would eventuate from not checking. (Finding problems when doing safety checks surely should be an exception rather than the rule.) But if a safety check is rightly instituted, as it often is, the expected utility of checking is positive.

Of course, a further problem with omitting a safety check is that it leads to a habit of not-checking. However, even if a habit were not being formed, a failure to check simply because the probability of an accident is less than  $1/2$  would be an egregious breach of responsibility—so egregious, in fact, that even people who omit safety checks typically don't justify their omission by saying that the probability of the check being needed in the given case is less than or equal to  $1/2$ , as an analogue of (5) would have it, but claim that the probability is negligible or something like that (and are sometimes literally fatally wrong about the negligibility).

A specific formulation of a principle like (5) in the literature is given by Almeida and Oppy<sup>7</sup> in their version of the paralysis argument:

if we do believe that it is not unlikely that there are unknown goods which would justify us in not preventing [an evil]  $E$ , then it is very hard to see how we could fail to be justified in not preventing  $E$ .

I take it that Almeida and Oppy are committing to this principle:

7. If it is not unlikely that there are unknown goods which would justify us in not preventing an evil  $E$  were the goods known to us, then we are justified in not preventing  $E$ .

In an important way this claim is stronger than in (5), since it is easier to meet Almeida and Oppy's 'not unlikely' condition than the probability  $1/2$  condition in (5). And the swimming pool story is a counterexample to (7), just as it tells against (5). Let  $E$  be the evil of the drowning of the innocent person we see thrown into pool  $A$ . Further, it is not unlikely—indeed, it has probability  $1/2$ —that there are ten people drowning in swimming pool  $B$ . If so, the good of saving their lives would, if known, justify us in not preventing  $E$ , since if we knew that there are ten people drowning in pool  $B$ ,

<sup>7</sup> Michael J. Almeida and Graham Oppy (2003), 'Sceptical Theism and Evidential Arguments from Evil', *Australasian Journal of Philosophy* 81 (2003) 496–516, 507.

we would drain pool *B*, not pool *A*. Hence by (7) we are justified in not draining pool *A*. But clearly we are obligated to drain pool *A*.

Just as (5) does, (7) uses probabilities where something more like expected values are needed.

### **3 Can the mistake be avoided?**

We looked at two arguments for the claim that strong sceptical theism implies that, at least on some versions of the story, consequence-based reasoning leads to doubt about whether we should stop Ashley's attacker. Both confused probabilities with expected utilities. The natural question is whether there is some way of avoiding this mistake and still getting to the conclusion along similar lines.

I will argue that the answer is likely negative. Here is the rough intuition. There are known features of the action and unknown ones. The known ones clearly require stopping the attacker, assuming we can do so without undue risk to self or others. The unknown features, were we to know them, would be just as likely to strengthen the case for stopping the attacker as to weaken it, and for any degree by which they could weaken the case for intervention, they could equally well strengthen it to that degree. Thus our information about the unknown features neither favors intervention nor nonintervention. Given that the known features require intervention, and the unknown make no difference given our information, we should still intervene. I will now give a more elaborate version of this argument.

Actions are evaluated on the basis of what I will call 'morally relevant features' or just 'features' for short. These features may include consequences, intentions, relationships, normative statuses of the agent and patient, etc. Consequences include causal consequences, but also constitutive ones—i.e., morally relevant states of affairs partly or wholly constituted by the action and/or its causal consequences. If consequentialism is true, then only consequences are among the 'features'.

An action is required, justified or unjustified *simpliciter* or *on balance* provided that it is required, justified or unjustified (respectively) in the light of all its features. But an action may also be required, justified or unjustified in the light of features of type *F*. For instance, an action may be justified in light of its consequences, or unjustified in light of its intentions. Whether an action justified or unjustified in light of features of type *F* is justified *simpliciter* depends on how it fares with respect to features beyond *F*.



## Sceptical Theism, the Butterfly Effect

I shall take reasons, as well as requirement, justification or unjustification, to be relative to an agent and her information state. In other words, the interest is in the internalist concepts. For the interest in this paper is whether sceptical theism paralyses an agent morally, and that seems to be an internalist question.

One may have moral reasons for or against an action in light of a feature or a type of feature. I will use 'reason' to mean moral reason. Reasons provided by different features combine and there is, no doubt, some complex relationship between the reasons provided by the features of an action and its requirement and justification statuses. Clearly, if all the features of an action provide reasons for the action, the action is justified, and, very likely, if they all provide reasons against the action, it is unjustified. But how the relationship works in the in-between case, where some features provide reasons in favor of the action and some against, is difficult. Perhaps, for instance, an action is required if and only if on-balance its features provide reason for it.

There is hard work in moral theory to be done in identifying morally relevant features and working out the relationship between requiredness, justification and reasons. But the following is a platitude: reasons—and remember that I am only talking of moral reasons—in favor of an action contribute to its being required and to its being justified while reasons against an action contribute to its being unjustified. This platitude makes this moral principle plausible:

8. Suppose that the features of an action  $A$  are divided into two types,  $X$  and  $Y$ , and suppose that in the light of the features of type  $X$ , the action is required (or, respectively, justified) for an agent  $x$ , and that features of type  $Y$  do not provide  $x$  with a reason against  $A$ . It defeasibly follows that action  $A$  is required (respectively, justified) for  $x$ .

The phenomenon of exclusionary reasons requires the defeasibility qualification. Suppose a judge is in a position where recusal is impossible (justice needs to be served, and no other judge is available), but the judge's decision impacts the well-being of the judge's family in such a way that solely in light of the well-being of the judge's family, judicial action  $A$  should be taken. Suppose that all the other features of the case mildly support  $A$ , but do not require it. Let  $X$  be features relevant to the well-being of the judge's family. Let  $Y$  be all other features. Among other things,  $Y$  will include requirements of procedural justice that require the judge to exclude features of type  $X$  from consideration. By itself,  $Y$  does not tell against  $A$ . But

because of these requirements of procedural justice, *Y* type reasons neutralize the (moral) reason-giving force of *X*. Consequently, if *X* were all there was, the judge would be required to do *A*, but once we throw the reasons of type *Y* into the mix, even though the *Y* reasons on their own mildly favor *A*, the judge is no longer required to do *A*.

More generally, higher-order reasons affect which first-order reasons count, and how much they count for. In themselves, higher-order reasons in *Y* may not favor or disfavor an action *A*, but by affecting how the first-order count they may shift the evaluation for or against *A*. The possibility of such reasons being found in *Y* requires the defeasibility qualification in (8).

Now let's go back to Ashley's case. Let *X* be the facts about what Ashley's sufferings would be like if we do not intervene as well as the other known features of the case. Let *Y* be all the unknown features. Remember that the reasons we are interested in are internalist reasons in light of the available evidence. We do not have any evidence either way about the unknowns, or so the sceptical theist insists. Given this, the unknowns internalistically favor neither prevention nor non-prevention. On the other hand, *X* clearly requires us to prevent the crime. So, by (8), we can defeasibly conclude that we should prevent Ashley's sufferings.

The crucial question now is whether sceptical theism provides a defeater. And we need to be clear on what that defeater would be like. The unknown features do not provide an internalist reason against preventing the crime. Depending on what the unknowns actually are, they might provide an externalist reason against prevention, or an additional externalist reason for prevention. However, the defeasibility in (8) comes not from this, but from the possibility of *Y* containing higher-order reasons that do not themselves favor non-prevention but that exclude some or all the reasons in *X*.

Higher-order reasons come up in special contexts such as commands, promises, official roles or special relationships. The case of Ashley does not *appear* to be at all like any such context. It may seem odd to rely on this non-appearance, however, in a defense of sceptical theism. After all, the sceptical theist specifically wants to block inferences from the absence of the appearance of a reason to the absence of a reason. But the contexts are different. The sceptical theist's scepticism encompasses the realm of value as well as causal and constitutive connections between localized states of affairs and other morally relevant states of affairs. This scepticism lies at the level of first-order reasons. Here, however, what we are relying on is our ability to know what higher-order reasons there may be. I

## Sceptical Theism, the Butterfly Effect

will argue that we should not have qualms derived from sceptical theist with regard to higher-order reasons in *Y* when deciding whether to help Ashley.

First, higher-order reasons appear to always be reasons for a particular agent in a particular context. They are reasons for a particular agent to evaluate her first-order reasons in a particular way. That I am commanded by my commander to take yonder hill entirely excludes reasons of personal convenience from my consideration, rather than just outweighing them<sup>8</sup>, and hence the command provides *me* with a second-order exclusionary reason. But it does not exclude reasons of my (or your) personal convenience from *your* consideration when you are not subject to that command. In this way, higher-order reasons may well differ from many first-order reasons, since the fact that something is a first-order good for me arguably gives a (de-feasible) first-order reason to *every* agent to provide it.

Now, the sceptical theist only professes scepticism about *God's* reasons. In the case of first-order reasons, these may have sufficient overlap with our reasons so as to raise the paralysis problem that we are considering. But scepticism about God's *higher-order* reasons does not with any plausibility lead to scepticism about ours.

Second, we can imagine cases where we don't know about the force of higher-order reasons. These will be rather contrived cases. Let's say you are now in an obvious bit of a difficulty. That gives me reason to help you. But I notice I have amnesia. So for all I know, I promised to leave you to your own devices in a case like this. But the mere chance that I made a promise, with no actual evidence, gives me no reason to refrain from helping. Moreover, just as there is a chance that I promised not to help, there is a chance that I promised to help. The unknown externalist higher-order reasons might readjust the evaluation of the known first-order reasons in favor of not helping but they can also readjust in favor of helping, and so these opposed higher-order unknowns wash out. Thus not knowing whether there might not be a relevant externalist higher-order reason does not defeat the inference in (8).

In summary, by (8), we have a defeasible internalist reason to help Ashley. It looks like the only defeater for (8) is something that would change how reasons combine, and that would have to be a second-order reason. But neither known nor unknown externalist higher-order reasons provide a defeater for the inference.

<sup>8</sup> Cf. Joseph Raz, *Practical Reason and Norms*, 2nd ed. (Princeton: Princeton University Press, 1990).

Note, too, that the present solution works *better* if consequentialism is true. Higher-order reasons are not a notion consequentialism is particularly friendly to. And when we have only first-order reasons in play, the defeasibility in (8) seems to disappear.

#### **4 The butterfly effect**

Sceptical theists are not the only people who have the difficulty that unknown factors morally swamp the known factors. Suppose we take the “butterfly effect” hypothesis from chaos theory really seriously, so that we think relatively small causes, such as a butterfly’s wing flutter, can have morally enormous effects down the road, like an earthquake in Japan in ten years.

In his discussion of the butterfly effect, Lenman<sup>9</sup> focuses on identity-involving actions, ones that change who the people populating the earth are. Which of millions of sperm meets up with an ovum is very likely to be different if the timing of intercourse is slightly changed. And it is not hard to change the timing of intercourse. Suppose that Alois Hitler went shopping in the afternoon, and in the evening, together with Klara, conceived Adolf. A small shift in when some other customer entered the store hours earlier could easily have resulted in a shift in when Alois was served, and in turn changed the timing of conception, and brought it about Adolf was never conceived.

Of course, whether Adolf Hitler or—on the positive side—Jonas Salk existed has vast repercussions for the identities of the earth’s denizens. But the Hitlers and Salks of the world are not the only ones whose lives have vast repercussions. The same is true for most people—it may just take longer. For a typical person’s actions are likely to affect the timing of intercourse for a number of people over a lifetime, not to mention the choices of partners. The extent of the identity-effect will likely then grow exponentially from generation to generation, especially now that our world is so interconnected globally.

And once the identities of much of the earth’s population are affected, this will have vast effects on people’s wellbeing compared to which the direct effects of our daily actions are likely to pale. What awful dictators will come into existence? Will a great medical researcher who finds a cure for cancer be born in the 22nd century or the 23rd? The wellbeing of millions or even billions depends on this.

<sup>9</sup> Op. cit. (2000).

## Sceptical Theism, the Butterfly Effect

Given plausible assumptions about the chaotic nature of our world, our actions have unpredictable consequences of very large magnitude in the future. That magnitude is likely to be so large that it will completely swamp the predictable short-term consequences.

This leads to a paralysis argument exactly analogous to the one for sceptical theism. Lenman defends this paralysis argument, and Howard-Snyder<sup>10</sup> observes the analogy between the two arguments. The two arguments are indeed analogous, and both are brought down by principle (8). The unknown chaotic consequences can be put into  $Y$ , and they neither favor action nor inaction, and so the decision should be made on the basis of the known factors which we can put into  $X$ . In fact, the butterfly effect case is easier to handle than the sceptical theism case, because *ex hypothesi* all we are worrying about are consequences, since chaos doesn't provide any mysterious higher-order reasons that might yield a defeater in (8).

### 5 Weakening the force of reasons

But there is an objection to lines of thought like this. To make the objection clearer, simplify our story by supposing consequentialism is true and that we are deciding between actions  $A$  and  $B$ . For each action  $X$  (where  $X$  is  $A$  or  $B$ ), there is a known utility  $U_{X,1}$  of consequences and an unknown utility  $U_{X,2}$  of consequences, and it is known that the unknown utility swamps the known in the sense that  $|U_{A,2} - U_{B,2}| \gg |U_{A,1} - U_{B,1}|$ . Moreover, the total utility is just the sum of the two:  $U_X = U_{X,1} + U_{X,2}$ . Hence if we knew  $U_{A,2}$  and  $U_{B,2}$ , we could make our decision solely on their basis. But we don't. Instead all we know is that  $U_{A,1} > U_{B,1}$ . And this difference is swamped. How, then, can we decide?

A simplified version of the response that relies on (8) is that we should decide on the basis of the epistemically expected values  $E(U_{A,1} + U_{A,2})$  and  $E(U_{B,1} + U_{B,2})$ . Since  $U_{A,1}$  and  $U_{B,1}$  are simply known—let's say they are known to be equal to some numbers  $\alpha$  and  $\beta$ —these two values respectively come to:  $\alpha + E(U_{A,2})$  and  $\beta + E(U_{B,2})$ . But now in the complete absence of information about the unknown effects, we have no way to distinguish  $U_{A,2}$  and  $U_{B,2}$  epistemically, and hence it seems reasonable to say  $E(U_{A,2}) = E(U_{B,2})$  as we are dealing with epistemic expectations. Since we

<sup>10</sup> Op. cit. (2010).

know that  $\alpha > \beta$ , we then have  $E(U_{A,1} + U_{A,2}) > E(U_{B,1} + U_{B,2})$ , and we should do *A*. There is no paralysis.

So far so good. But now Lenman<sup>11</sup> raises this worry. Given that the difference between  $\alpha$  and  $\beta$  is much smaller than the difference between  $U_{A,2}$  and  $U_{B,2}$ , the force of the reason to do *A* instead of *B* is pretty small. We are deciding something that has enormous repercussions on the basis of something very minor. Lenman<sup>12</sup> offers the analogy about deciding between two possible landing locations for the Allies on the continent, and in the absence of any data as to which is better, deciding on the basis of the wellbeing of a particular dog known to be at one of these locations.

One answer is that the defeasible conclusion in (8) can be strengthened. If the *Y* features don't provide a reason against *A*, then the strength of the reason for *A* overall is no less than that provided by the *X* features. But while this is very plausible, it does not address the intuitions behind Lenman's worry.

The dog case is probably easy to handle. It is insulting to the soldiers to make the decision on the basis of the life of a dog. But vary the case so it's not insulting. Suppose that somehow (due to a prophecy?) we know that a landing at *A* will cost 100,000 lives and a landing at *B* will cost 100,001 lives. If all who would die on *A* would also die on *B*, and if we know who that 100,001st victim on *B* would be, say Jim, we clearly ought to go for a landing on *A*. And the strength of the reason for landing on *A* is exactly the strength of the reason for saving Jim's life. The fact that 100,000 people are going to die at Jim's side does nothing to weaken the value of Jim's life.

But what if we don't know who the extra survivor on landing *A* is, and we keep everything the same? This, too, should make no difference. The value of that anonymous (that is, anonymous to us: but a mother, brother, sister, friend, daughter, etc. to those close to him or her) extra survivor is no less for the hundred thousand dying at his or her side. Recall the swimming pool case. The sense in which we knew who the extra person is was very thin: we saw a particular stranger in the distance, but we didn't know him or her from Adam, as we say.

Finally, suppose that the identities of those who would die on either landing are completely different. In case *A*, 100,000 people—anonymous to us—will die. In case *B*, 100,001 will. Is there a significant difference in the strength of reasons to land on *A*? Without being able to identify particular individuals across the

<sup>11</sup> Op. cit. 356–358.

<sup>12</sup> Op. cit. 357.

## Sceptical Theism, the Butterfly Effect

two scenarios, it may be difficult to *feel* the difference. But let's switch cases. Alice is about to unleash 100,000 doses of a poison gas, where each dose will kill a random person in New York. Her friend Bob can't stop her, but thinks to himself: 'There is not much difference between 100,000 anonymous people and 100,001. If I load another dose in the canister, probably quite a different collection of hundred thousand people will die.' And he adds a dose.

Alice is a murderer a hundred thousand times over. But Bob's action is morally on par with murder, too. It is no less on par with murder for the fact that the identities of the victims were likely switched as a result of his action. But even though Bob is effectively a murderer, there may not be a person whom he murdered, since the dose Bob added presumably got mixed with all the others, and contributed a little to the deaths of many. (The law will get Bob for conspiracy to murder, but that is just a pragmatic solution.)

Still, even if we can only call Bob's action 'on par with murder' and not really a 'murder', he is no less bad than a murderer. To make the point perhaps clearer, suppose that 100,000 people each with malice contribute one dose of poison gas to the canister. Each is as bad as if they had released that dose in some place where it wasn't mixed with other doses but killed one individual.<sup>13</sup>

We can make the same point on the side of saving lives. Carl is about to release an antidote to a poison that was set to kill a million people. Unfortunately, Carl only has 100,000 doses of the antidote, so he is going to save only 100,000 lives. Dale comes by with an extra dose. Dale's action is morally on par with saving one more life, even if the identities changed.

How hard should we fight to keep Bob from adding his dose of poison gas to Alice's canister? Exactly as hard as we would to prevent one random murder. How hard should Dale try to come up with that extra dose? As hard as it would be worth trying to save one life. We can, after all, imagine 100,000 people each working hard to contribute a dose of antidote. They should each work as hard as one should to save one life. The strength of a reason isn't measured by feeling, but by how hard it makes it rational to follow the reason—what cost the reason makes it rational to accept.

And by the same token, the reason to land on landing *A* instead of *B*, when one fewer soldier will die, is as strong as the reason to save a life. It doesn't *feel* as strong. But the mistake here is a non-financial version of the well-known mistake by which someone will drive

<sup>13</sup> Cf. the cases in Derek Parfit, *Reasons and Persons*, Oxford: Oxford University Press, 1987), chapter 3.

## Alexander R. Pruss

across town to get a free \$5 bottle of shampoo but who would not bother to cross the street to a different dealership to buy a car for \$5 less. The apparent decreases in the strength of reasons in light of the butterfly effect or sceptical theist hypotheses are a kind of moral illusion, akin to this financial illusion.

## 6 Conclusions

The paralysis argument against sceptical theism confuses reasoning about which action is more likely to result in the better outcome with reasoning about which action is a better bet. Resolving this confusion does not require going beyond consequence-based reasoning, but also does not require committing to consequence-based reasoning. I offer a defeasible moral principle, (8), that suggests that whether consequentialism is true or not, it is very unlikely that there is a way of repairing the paralysis argument. This also solves the related, but not specifically theistic, problem of alleged paralysis coming from the butterfly effect. Sceptical theism and chaos may be problematic, but not for reasons of paralysis.

*Baylor University*  
*[alexander\\_pruss@baylor.edu](mailto:alexander_pruss@baylor.edu)*