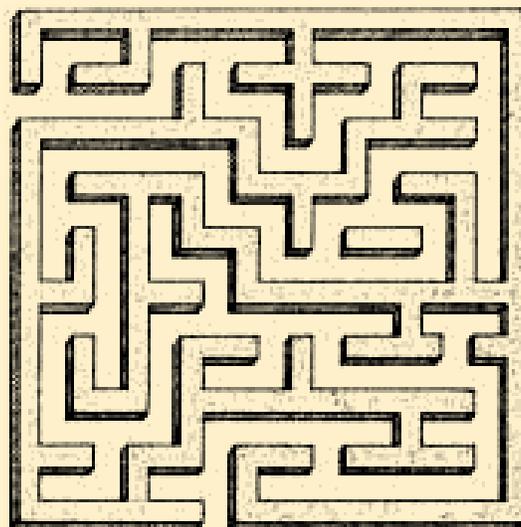


Labyrinth
Labyrinth
Labyrinth
Labyrinth
Labyrinth
Labyrinth



The Undergraduate Philosophy Journal of the
University of Iowa

Spring 2019
Vol. 5

In this issue:

Sam Lampe, “Why Fine-Tuning Does Not Support Theism”

Rana Hewezi, “Nozick’s The Experience Machine: The Preeminence of Authenticity in
Opposition to Stimulatory Experiences”

Benjamin Nelson, “Reconsidering Pascal’s Wager as a Bet Against Atheism”

table of contents

3	About the Journal
5	Sam Lampe, “Why Fine-Tuning Does Not Support Theism”
21	Rana Hewezi, “Nozick’s The Experience Machine: The Preeminence of Authenticity in Opposition to Stimulatory Experiences”
27	Benjamin Nelson, “Reconsidering Pascal’s Wager as a Bet Against Atheism”
33	Acknowledgements

Editor:

Kate Lohnes

about the journal

why “labyrinth”?

Plato uses the image of a winding labyrinth as a metaphor for the process of philosophical investigation in his dialogue *Euthydemus*. His image expresses his belief that, unlike the arts of rhetoric or sophistry—which rely on mere assertion and counter-assertion—philosophy absolutely requires that we retrace our steps in an argument and constantly re-examine our views in order to arrive at knowledge of the true and good. This is what we as students of philosophy aim to do, and this journal is meant to aid in that process.

our purpose

This journal was created with a threefold purpose. First, to provide undergraduate students with the experience and opportunity of publishing a paper in a philosophical journal. Secondly, to give an opportunity for students to be involved in peer to peer interaction through the process of creating the journal. Finally, to showcase the amazing philosophical work being done by undergraduate students at the University of Iowa.

papers

Why Fine-Tuning Does Not Support Theism

Sam Lampe

A version of the teleological argument for the existence of God that I have found interesting is one based less on analogy, like typical arguments from design, and more on probability. Often this likelihood-based version is called the fine-tuning argument. Generally, an advocate of the argument gives evidence that the range of life-permitting values in which the laws, constants, and initial conditions, etc., of the universe falls is small in comparison to what could have been. It makes it appear as though we were very lucky that life came to be in the way it is, and the suggestion that it came about randomly seems unlikely. Given that, we may be led to believe that it was more than just luck—there was a reason, and that reason is that God put us here. I wanted to find a more formal statement of this argument, however. It is not immediately obvious how it follows from the unlikelihood of this set of laws, conditions and constants given an atheistic cause that it is unlikely that said cause is what brought them about. More needs to be said about how we get from the former to the latter. We also need to know why theism does not succumb to the same considerations. Robin Collins offers an attempt at satisfying these considerations, which I will examine. In this paper, I will take a look at the support for the premises of his argument, including the likelihood principle and the indifference principle, and consider one of Collin’s responses to criticisms of this form of argument. I will ultimately conclude that this way of understanding the fine-tuning of the universe gives insufficient reason to favor the God hypothesis.

Collins begins his paper, “God, Design, and Fine-Tuning,” by offering a number of ways in which the universe is fine-tuned for life. One such example is that the region in which the strength of gravity must fall in order for life to exist is so exceedingly miniscule that it “amounts

to a fine-tuning of one part in 10^{31} " (Collins 3). He gives many more examples of fine-tuning throughout the paper. Counting up all these values, Collins says that "one could think of the values of the initial conditions of the universe and the constants of physics as coordinates on a dart board that fills the whole galaxy, and the conditions necessary for life to exist as an extremely small target, say less than a trillionth of an inch: unless the dart hits the target, complex life would be impossible" (4). If we take those estimates seriously, it seems that we are very lucky to be here, unless there were some other factor that caused all this.

To formalize the argument, Collins uses what he calls the "*prime principle of confirmation*," which I will call the "*likelihood principle*", as well as the "*principle of indifference*" (6,14). The likelihood principle states that, for any two explanations (hypotheses) for a piece of evidence, the evidence favors the explanation which makes that evidence more likely (6). This will be used to tie the premises to the conclusion. The indifference principle states that for that any two parameters, if we do not have any reason to think one value that it could have is more likely than any other, we should place equal likelihoods on each value or range of that parameter (14). This will be used to justify second premise. Collins formally states the argument as follows:

Premise 1. The existence of the fine-tuning is not improbable under theism.

Premise 2. The existence of the fine-tuning is very improbable under the atheistic single-universe hypothesis.

Conclusion: From premises (1) and (2) and the prime principle of confirmation, it follows that the fine-tuning data provides strong evidence in favor of the design hypothesis over the atheistic single-universe hypothesis. (7-8)

By “the atheistic single-universe hypothesis” he means to contrast this alternative with the multiverse response, saying that this hypothesis gives just one universe, without God. I may occasionally call this the “chance hypothesis”. For “theism”, Collins seems to have in mind the typical all-good, all-powerful, all-knowing God, so when I refer to the “narrow God hypothesis” this is what I mean. If I name “Yahweh”, I am pointing out the same person. I will also discuss a “broad god hypothesis”, by which I then mean a less robust position which is open to any divine being (so long as it is the kind that would bring intelligent beings into existence).

In support of the first premise, Collins says that God, being all-good and all-knowing, would know that the existence of intelligent beings is good and would in all likelihood bring some into existence (8-9). From the earlier analogies, he thinks it should be clear that the universe was very unlikely to turn out this way under the atheistic hypothesis, though he does give a more formal defense later (9). He also mentions that some may try to erode the support for this premise by positing that it may have been necessary that the universe turned out this way, and I will describe and respond to his criticism later (12-13). First, I will describe the grounds for the likelihood principle, then move on to some criticisms.

In support of the likelihood principle, Collins asks his readers to imagine that he went hiking in the mountains and discovered an arrangement of rocks that states, “Welcome to the mountains Robin Collins” (7). We could hypothesize either that those rocks appeared there by chance or that Collins’ brother came along earlier and arranged them that way (assuming these are the only viable alternatives). Of course, all of us would prefer the brother hypothesis. Collins believes that this is because “it strikes us as extremely *improbable* that the rocks would be arranged that way by chance, but *not improbable* at all that my brother would place them in that

configuration” (7). This example could be extrapolated to any number of instances, so it seems like the likelihood principle is a reasonable summary of our intuitions.

Before going into any specific worries about the likelihood principle, I want to point out that this argument seems to work for any universe, no matter the appearance of that universe. Imagine, for a moment, that we discovered some mysterious window through which we could observe another universe. Imagine that in this universe there were many rapidly expanding and contracting gas clouds a few parsecs wide at their largest, and a few kilometers wide at their smallest—call these ‘crunchies’—and that no other notable features inhabit this universe. No doubt the ranges in which the values and constants of this universe fell would be narrow, comparable to our own. Thus, the indifference principle would lead us to believe that the probability of this token universe (and other universes with comparable structures) existing would be *very* small. Imagine, then, that someone offers the following hypothesis; perhaps there is an all-powerful, very knowledgeable, all-crunchies-loving god which created this universe. This hypothesis can maintain a near-certain likelihood of the evidence. This example can be expanded for every token universe—we can imagine a god which loves barren landscapes, a god which loves an empty void, a god which loves mangoes, etc. The point being that, no matter the outcome of the universe, this form of argument would strongly support the creation of that universe. There seems to be something off about this.

For another example which illustrates the same point, imagine that someone you know wins a lottery in which one million people participated. Assuming the lottery was fair, the probability of them winning was one in a million, however, if we postulate that they rigged the lottery, that person’s win is certain. Should this cause us to think they rigged it? If it did, we would have to say that about every outcome, no matter who won. Why do both these cases seem

suspect? I suggest it's because using the unlikelihood of an outcome based on a particular cause to hypothesize the unlikelihood of that cause existing is unfair when, given that cause, every outcome would have to be very unlikely.

Perhaps one might respond to the crunchies-containing universe, and the rest of the examples given, that the reason we would not take these hypotheses seriously is that the probability of gods existing who would desire crunchies or mangoes or an empty void is itself very low, such that the prior probability of their existence is not nearly as high as a god which desires the existence of intelligent beings like Yahweh. To this I say that we'd first need to know that if a god existed, it would actually value the things that we do. We would need to establish that life is either objectively valuable, such that a sufficiently intelligent being would want it (and that if a divine being existed, it would be intelligent enough in the relevant areas to observe this value), or we would need to establish that it would be subjectively valuable to such a being.

For the first option, this is certainly a debatable area of philosophy—existential nihilists and those who do not deny the existence of objective value in life, but are not convinced of its existence, would not take this as an obvious step, and any criticisms of objective value would also become criticisms of this argument. However, there are a decent number of non-theist realists who might accept this premise, so perhaps this at least somewhat plausible for them. Even still, one would have to posit that if a divine being existed, it would be intelligent enough to see this intrinsic value—for all we know, a divine being could be (literally) psychopathic, knowing much about some things (the creation of worlds) but nothing about existential value. For the second option, we must have some reason to believe that a god would value the same kinds of things we do, but since it is not clear that gods would share any of our cultural influences, any of our evolutionary history (meaning their celestial 'genes' may be very different

from ours), or any of our day-to-day experiences, it is not obvious that those things which are apparently valuable to us would be so to a god, nor that certain intellectual capacities to see objective value would be had by said god. Thus, we have yet to be provided with a good reason to believe that in the crunchies-containing universe, a crunchie-loving god would be unlikely, so this still serves as a relevant counterexample to the form of the argument.

As Collins mentions, there are some notable counterexamples to the likelihood principle. In the Sober piece that Collins cites, the following example is given: “If you draw the six of spades from a deck of cards, the hypothesis that this was due to the intervention of an evil demon bent on having you draw that very card has a likelihood of unity, but few of us would regard this hypothesis as very plausible” (Sober 4). More can be said on these kinds of examples, but the basic idea is that we allow bizarre theories to be confirmed by the likelihood principle when it is not qualified.

Collins’ suggestion is that we restrict the likelihood principle to cases where the hypothesis in question “has some independent plausibility apart from evidence E or was at least not merely constructed to account for E” (Collins 6). He believes this is satisfied by the God hypothesis since “theism ... was believed long before the evidence for fine-tuning was known” (6). One wonders if this response is enough to distinguish the narrow God hypothesis from those kinds of counterexamples. Since belief in demons (and things like them) has existed throughout history (long before the six of spades was drawn), perhaps as long or longer than the modern conception of Yahweh, the demon-cards example would still fit even the restricted likelihood principle, suggesting we may need to restrict it further if we wish to save it, though that is beyond the scope of this paper.

Perhaps we can help Collins out. What other possible dis-analogies are there between the demon-cards example and the name-in-rocks example? One suggestion is geared specifically at the plausibility of the demon-cards example; why would the demon ever care what card was drawn? Sure, demons are typically associated with the number six, but such a menial sign does not seem like something even a being as arcane as a demon would care about. It seems like the initial plausibility of demons influencing card choices is incredibly low for that reason, but that the initial plausibility of an all-good being making intelligent beings is not nearly as low.

We may be able to sustain a different, but similar counterexample. Imagine that a poor farmer plants a crop of corn in the spring only to realize at harvest that the entire crop had what appeared to be a rare genetic disease that gives them extremely low yield. We could believe either that it was simply bad luck that led the seeds to come from the pairing of two recessive carriers of the diseased gene, or that a demon cursed his seeds, or influenced the pairing, or some such thing. In this case, given the character of demons, we know that they would love to bring pain on people, and the idea that they might do this to the farmer seems more expected than that of them influencing the drawing of a card, and should not fall to the same criticism the other example does. So, the likelihood principle still seems to have a class of potentially relevant counterexamples, in which case, it is not clear it can be reasonably used in its current form.

What I wish to explore further is exactly what it is that makes the demon-crops example less tenable than the name-in-rocks example. What is different? Collins' suggestion was that the hypothesis should be plausible independently of the given evidence. In other words, it is a *real possibility* that Collins' brother might spell words in rocks, for we are familiar with people who sometimes spell words with objects, including rocks. As for the demon examples, we have no reason at all to suspect that it is even possible for demons to exist, much less have familiarity

with the fact that they sometimes interact with humans to cause them grief. Thus, it is *not a real possibility* that demons influence crop yield or card draws.

What counts as a “real possibility” for use with the likelihood principle? In order for an explanation to be legitimate, the pieces of that explanation must at least be possible, as in, have a probability greater than zero. If the likelihood of a demon existing at all is zero, then no matter what evidence one can find in favor of a hypothesis which relies on it, it will never be a better hypothesis than any other. It is also worth pointing out that if we do not have an estimate of the relationship between the likelihoods of two hypotheses, we cannot say which one becomes more likely on the use of the evidence. The likelihood principle, given the fact that it only works with things known to have prior probability, should be thought of as something that impacts that prior probability. Even if the demon hypothesis has a nonzero likelihood, if it turns out that the estimate of the likelihood of demons existing is sufficiently low, then even a large impact on the hypotheses based on them may not make them more likely than their alternatives. Note that I am *not* saying that more evidence may counter the current evidence. Rather, what I am saying is that since we already need an initial estimate of likelihood, we must worry if that initial estimate is going to be high enough.

What makes the name-in-rocks example a real possibility is that, from our experience, the brother hypothesis has not only a nonzero probability, but also a close enough relationship in probability to the random formation of rocks such that the discovery of a particular formation like in the example would allow the brother hypothesis to outstrip the random hypothesis. As for the demon examples, we do not have any such ability to estimate the likelihood of the demon hypotheses, such that we could say not only that they are nonzero but also that they are close enough to the chance hypotheses to outstrip them given the evidence.

To illustrate the need for an initial estimate of likelihood more clearly, consider the following example. Imagine attending a soccer game that one of your relatives is participating in; call this relative Ethan. There are two coaches for this team, coach Alex and coach Barnes, who take turns coaching the games, neither of which you have seen yet. When the game starts, you notice that Ethan has been made a midfielder rather than a defender. You know that when Alex is coaching, he almost always makes Ethan a defender, perhaps because Ethan's parents tell you so. You also know that when Barnes coaches, he makes Ethan a defender and midfielder with roughly equal frequency. From this information, can you determine who is coaching this particular game? It doesn't seem so. It is entirely possible that they coach games with equal frequency, making the better choice Alex, but when we do not know how often they coach games, we can't say which is superior. For all we know, Barnes may coach games with ten times the frequency that Alex does, in which case, Barnes would be the better choice. Removing information about the frequency with which either one coaches undermines our ability to make a solid judgement about this case.

While he didn't make it essential to his argument, Collins mentioned a formula that can be useful for illustrating this problem further (7). Bayes' theorem, as it is called, goes as follows;

$$P(H|E) = \frac{P(H) * P(E|H)}{P(E)}$$

Reading the formula is simple enough. For example, the leftmost term reads $P(H|E)$, meaning the probability (P) of the specific hypothesis (H) being correct with respect to the observation of the evidence (E). Notice that in order to get our $P(H|E)$, we need $P(H)$, the general probability of the hypothesis occurring in the first place, regardless of the evidence in question. This is consistent with the examples provided earlier, showing that we need the general probability of the hypotheses in question before we can begin to say that one hypothesis is preferable to the other.

There is, however, something of value to be gained from this formula. First, we rearrange it so that it reads as follows:

$$P(H|E) = \frac{P(E|H)}{P(E)} * P(H)$$

We can now see that the formula is separated into three terms, from left to right; the probability that our hypothesis is correct with respect to the evidence, the impact of the evidence on the probability of the hypothesis, and the general likelihood of the correctness of the hypothesis. We then divide each side of the formula with its equivalent from another instantiation Bayes' theorem dealing with the same evidence but a different hypothesis to get the version that Collins gives:

$$\frac{P(H_1|E)}{P(H_2|E)} = \frac{P(E|H_1)}{P(E|H_2)} * \frac{P(H_1)}{P(H_2)}$$

The terms now are, from left to right; the ratio of the probability of the first hypothesis with respect to the evidence to that of the second hypothesis with respect to the evidence, the ratio of the probability of the evidence coming about with respect to the first hypothesis' truth to the probability of the evidence coming about with respect to the second hypothesis' truth, and the ratio of the general probability of the first hypothesis to the general probability of the second hypothesis.

If we assume that H_1 is the narrow God hypothesis and H_2 is the chance hypothesis, this formula can tell us something about what the evidence does to our hypotheses' relationship. We know that the probability of this universe coming about with respect to a random cause, $P(E|H_2)$, is incredibly low, and that the probability of this universe coming about with respect to a divine cause, $P(E|H_1)$, is nearly certain. The impact on the ratio of the general likelihoods of the two hypotheses, $\frac{P(E|H_1)}{P(E|H_2)}$, will therefore be incredibly large. That is something that can be

ascertained from this formula, though it still leaves us in the dark as to which hypothesis is better. Whenever we have a general relationship between two values, even an enormous impact on that relationship cannot tell us which of the two becomes larger since such a determination would require knowing how they were related in the first place.

Perhaps one could suggest for the two-coaches example, and the teleological argument in general, that the indifference principle provides the initial estimate of likelihood that is needed. It seems intuitive to say that, given these two coaches, there is no better reason to think that one is busier than the other, and as such, we should be indifferent and place equal likelihoods on their being the correct coach before the evidence. Similarly, it might be said, that given our two kinds of hypotheses, theistic and atheistic, we can say that they have equal likelihoods absent any evidence to the contrary.

To this I say that I'm not sure such an instantiation of the indifference principle would be justified, since, in all cases where the indifference principle can obviously be used, there is some relevant similarity between the options in virtue of which we can assign them equal likelihoods. The statement Collins gives of the indifference principle seems to include this consideration, describing the principle by saying "*when we have no reason to prefer any one value of a parameter over any other, we should assign equal probabilities to equal ranges of the parameter, given that the parameter in question directly corresponds to a natural parameter*" (14). In his formulation, that similarity would be the shared property of being a parameter in a range of natural values.

For the two-coaches example, we can say that with respect to the fact that one hypothesis picks one of two coaches, it is equally likely to choose the one of the two that has greater availability in his schedule and is therefore open to coaching games. Collins gives an example

which holds to this rule as well—he uses a fifty-sided die, saying that the indifference principle is what would account for our intuitions that the die is equally likely to land on each side (15). This bears the same kind of similarity between the alternatives—with respect to the fact that we have picked a certain side on a fifty-sided die, it is equally likely that it is weighted as any other side, and as such, it is equally likely to come up given the current information. If, for the two-coaches example, we replaced coach Alex with team captain Charles, it is not so clear that the indifference principle applies, and if we change the die so that the sides do not have identical surfaces, it too would not seem to fit the indifference principle.

Can the same be said of the narrow God hypothesis and the atheistic hypothesis? With respect to what fact do they have the same probability of coming up? They are not both ranges of the same parameter, as Collins' version of the principle governs over. Can we say the fact that they are both hypotheses which we can articulate is a good enough reason? Such a loose application of the indifference principle would also work for both the demon-cards example and the demon-crops example. In fact, that would overcome the problems facing the demon hypotheses' use with the likelihood principle. Evidently, we need more concrete similarities between our hypotheses, and it is not clear that we have this for our atheistic single universe vs. theism hypotheses.

Let's now move on to an objection to the general form of the fine-tuning argument, which involves the indifference principle. Collins argues that to respond to the teleological argument by saying “there could be a more fundamental law under which the constants of physics *must* have the values they do. Thus, given such a law, it is not improbable that the known constants of physics fall within the life-permitting range” is insufficient, as this does not alleviate the problem of the universe being very unlikely (9). Instead, it merely makes the existence of that

more fundamental law very unlikely, since the indifference principle applies to it as well. Rather than looking at the myriad of different values the universe could've had and concluding that it would be very unlikely for them to turn out the way they did, we instead look at the myriad of different possible fundamental laws that could have necessitated their various universes, and the few that would lead to a universe like this, making it still incredibly unlikely. The form of the argument changes, as we are no longer saying that the evidence makes the outcome unlikely, since the outcome is guaranteed. Rather, we are saying directly of the hypothesized cause that it is unlikely to exist.

My criticism here is that this response is a double-edged sword, especially when a specific god, like Yahweh, is chosen, since much the same can be said about the narrow God hypothesis. Given the myriad of possible gods that could in theory exist, positing one of them merely moves the unlikelihood of the universe existing the way it does up a level, to the unlikelihood that one particular kind of god exists. While this hypothesis still guarantees the outcome, the probability of it being true in the first place is now incredibly unlikely, because of the indifference principle. If, however, we gave up the principle to avoid this problem, we would lose our support for the second premise, that the chance hypothesis makes this universe unlikely, but if we keep it, the prior probability of the hypothesis becomes extremely low. Either way, we would have no reason to favor the god hypothesis given this evidence.

It could be responded to this that the teleological argument need not posit so specific a god as Yahweh, and merely argues for *some* kind of personal creator which is inclined to bring intelligent beings into existence, meaning the broad god hypothesis. This is not Collins' argument, but we can imagine someone re-formulating it this way. Therefore, we may be able to mitigate this problem by weakening our claim. However, the atheistic single-universe hypothesis

may be able to respond in the same way. The atheist has no specific idea of what natural cause would bring this universe about: perhaps it is a random-universe-generating cause, perhaps it is a more fundamental law, perhaps a pseudo-random one (weighted cosmic die), perhaps it is a divine wind or divine robot (neither of which would be theistic options, since those are not beings), etc. Nor do any of these options have to specify but one kind of life—they're all perfectly open to exotic life, leaving the possibilities for atheistic causes at least as open as the possibilities for theistic causes.

Moreover, it is not even clear that this avoids the moving-the-unlikelihood-up-a-level problem the indifference principle presents. Even if we are not positing only *one individual* form of god, we would still be positing only *one class* of gods. Specifically, it would be the class of gods who either recognize the inherent value in the existence of intelligent beings or subjectively values intelligent beings. Yet, there are many more classes of potential gods. What of the class of gods which loves crunchies? Some may like big crunchies, some prefer small ones, some which crunch quickly and others which crunch slowly—there are at least as many theoretically possible crunchie-loving gods as there are intelligent-being-loving gods. The same could be said for any other arbitrary object which a god can be offered for. In which case, the indifference principle would lead us to believe that the likelihood that one particular class of gods would have one or more of its members instantiated in the world would be incredibly low, given the myriad of classes which exist, thus being no greater detriment to the atheistic necessity hypothesis than to the broad god hypothesis, though certainly this criticism is stronger for the narrow God hypothesis.

It may be possible to avoid this consideration altogether by making the hypothesis even less strict, saying only that it is meant to hypothesize a god of some kind, not even going so far

as to say that this god would necessarily like or bring about intelligent beings. That said, it is not clear that this extremely minimal god hypothesis actually makes the fine-tuning likely at all, given that this would be susceptible to a criticism from earlier—if gods existed, why would we think that they would value the same things we do? Thus, we merely trade unlikelihood of the prior probability for unlikelihood of the evidence coming about given the hypothesis.

What can be said, then, of this argument for theism? In order for the likelihood principle to work, we need to know that the chance of God existing in general is not only greater than zero, but also great enough that the impact the fine-tuning has on the probability of his existence is enough to raise it above the chance hypothesis. It is unclear that we know this, and Collins' suggestion about the prior existence of belief in God does not seem satisfactory. It is important to note that Collins was not trying to argue that the likelihood principle alone makes the narrow God hypothesis more likely than the alternative, for he thinks all the evidence needs be considered for such a feat (8). That said, if "evidence" simply means "impact on initial likelihood," then we still need a reasonably high initial likelihood in order for the all the evidence together to make a difference. If we take Collins to be making a sufficiently minimal claim, his argument *may* be perfectly valid—that the change in the ratio of the likelihood of one hypothesis to the likelihood of the other has been modified very heavily, assuming we can avoid the problem the indifference principle causes for the prior probabilities of the narrow and broad hypotheses. Even then, if we wish to ever use this modification of the ratio of likelihoods to support belief in some kind of god, we still need to have a reasonably high estimate of those hypotheses' initial likelihoods, and if we cannot provide them, then we cannot conclude anything of interest from the fine-tuning argument.

Works Cited

- Collins, Robin. "God, Design, and Fine-Tuning". *God Matters: Readings in the Philosophy of Religion*, 2002. *CiteSeerX*, doi:10.1.1.470.7166. Accessed November 3, 2018.
- Sober, Elliot. "Bayesianism – Its Scope and Limits." *Bayesian*, stats.org.uk/bayesian/Sober2002.pdf. Accessed November 6, 2018.

Nozick's The Experience Machine: The Preeminence of Authenticity in Opposition to Stimulatory Experiences

Rana Hewezi

Robert Nozick's, "The Experience Machine," is a hypothetical scenario that Nozick employs as evidence to claim that individuals desire far greater than their own intrinsic pleasure. In doing so, Nozick applies the Experience Machine to ultimately reject hedonistic utilitarianism, while simultaneously highlighting how stimulatory experiences in themselves, though pleasurable, do not overcome the vast majority's preference to genuine experiences that are of subordinate quality. Through this, Nozick makes a fundamental distinction between what individuals want to achieve and what they want to experience.

Utilitarianism is a normative, ethical perspective that determines right and wrong by maximizing the total amount of goodness itself. Various subgroups arise from this broad viewpoint, such as hedonistic and ideal utilitarianism, differing on what constitutes goodness and the property of what is good. Both outlooks accept that the most ideal action is the one that leads to the most net value, in relation to all the other alternatives, and is achieved by summing up individual consequences. However, hedonistic utilitarianism adopts a more singular definition of goodness, in which pleasure is the only intrinsic good, and its complement, pain, the only intrinsic bad. Despite their differences, hedonistic utilitarianism is considered to be a species of ideal utilitarianism.

However, ideal utilitarianism, often called generic utilitarianism, denies that the sole object of moral concern is maximizing pleasure. Alternatively, this perspective holds that intrinsic goodness is multifaceted. Therefore, generic utilitarianism believes goodness to be much more complex and can consequently be characterized by numerous intrinsic desires, such

as love, family and knowledge. Furthermore, Nozick himself sides with this view and attempts to defend generic utilitarianism by derailing hedonistic utilitarianism through the scheme of the Experience Machine.

The Experience Machine is a stimulating thought experiment that generates exclusively, pleasurable experiences handpicked by the individual. Although these hallucinatory experiences maximize goodness and conceal the fraudulent conditions from the participant, Nozick argues that most individuals would choose not to participate. He claims that humanity desires to perform actual experiences in addition to having experienced them. Moreover, Nozick defends this point by revealing that in the matter of specific experiences, one must first crave achieving the act in order to desire experiencing it. Additionally, Nozick believes that individuals have a preconception of the type of person they yearn to be. However, when one enters the machine, all semblance of character is erased, seeing as it is hard to distinguish between the person they are inherently, in contrast to the person they are experiencing themselves to be.

Despite the numerous versions of the Experience Machine, including the added condition of an individual's willingness to execute their entire family, in which the memory would be immediately forsaken, in exchange for continued stimulation of maximized pleasure, or a transformative machine that converts an individual to the type of person they desire to be while preserving their original character, Nozick maintains the position that human desire extends beyond our own intrinsic pleasure. Therefore, to Nozick, the machine is incompetent in achieving a fundamental human desire: an ability to actively live in contact with reality.

In this way, Nozick demonstrates that one's own intrinsic pleasure is not the only thing humanity, as a whole, desires because, otherwise, more people would cooperate. Furthermore, Nozick reveals that the experience of unflinching pleasure, in addition to the liberation of future

pain, though mere fabrications of the mind, is not enough to convince individuals to partake in this experiment. By doing so, Nozick not only suggests that the fact that most individuals are repelled by the mere idea of partaking in the machine is strong evidence against hedonism, but also hints at a support for a more general view of what is good. Through this, hedonistic utilitarianism faces a big obstacle in maintaining pleasure, or lack thereof, as the leading contributor to action.

In connection, Nozick's argument instead supports generic utilitarianism because he indicates that individual desire exceeds one's innate emotional state. By doing so, Nozick suggests that humanity craves more than experiences, but interactive contact with a greater reality. This implication that individuals favor to stay acquainted with the larger and potentially harmful complexities of reality in contrast to the fixed simulation of the machine implies that humanity longs for far greater than simple happiness in the absence of pain that the machine would be able to supply. This again supports Nozick's claim that individuals want to participate in the physical act of doing things, rather than simply experiencing them. However, the simple, natural tendency to act is stripped away in the virtual reality that constitutes the Experience Machine.

Likewise, Nozick further claims that fulfillment in one's life may also be uncovered within deeper meaning outside of simply pursuing happiness. This is a direct opposition to hedonistic utilitarianism, which views happiness, in the absence of pain, as humanity's fundamental objective. Moreover, authentic happiness is thwarted in the hallucinatory thoughts that are induced in the Experience Machine because the experiences one encounters are deceptive to the harsh realities that lay outside the walls of the experiment. Therefore, the emphasis on real experiences reflects humanity's demand to connect within a deeper segment of

reality, where one can experience both hardship and loss, which function to make happiness discrete and immensely recognizable. Ultimately, Nozick challenges the public to question the complexity of what they consider to be significant, in relation to their experiences.

Correspondingly, George Edward Moore, whose view is aligned with generic utilitarianism, raises similar arguments against hedonistic utilitarianism, although indirectly. Moore stresses that goodness is a simple, unanalyzable and indefinable property which cannot be described through naturalistic properties. Therefore, Moore does not believe that what is good and pleasure can be equated, seeing as it refutes his claim that goodness is indefinable. Hence, to Moore, objects of nature may still be considered good, but one cannot make the reverse argument. However, hedonism defines intrinsic good, or good in itself, in terms of a natural quality, pleasure, and thus, according to Moore, violates the naturalistic fallacy. In doing so, Moore may be implying that individuals who identify with a naturalistic view will be urged towards an incorrect, singular view of value, rather than the pluralist perspective he attempts to defend.

Consequently, when one rejects naturalism, Moore's perspective becomes more apparent: that there are countless indecomposable goods. Therefore, Moore adheres to a pluralistic utilitarianism, otherwise known as ideal utilitarianism. Considering that he takes goodness to be a property whose connotation is independent of others, it can be embodied in various states or entities. In this way, multiple objects may possess intrinsic good instead of merely states of pleasure, as generic utilitarianism has it.

A hedonist might respond to Moore's argument by insisting that others are mistaken in holding things outside of pleasure as being intrinsically good, or good for its own sake. Instead, a hedonist would claim that these other goods, such as love, family, and knowledge, aside from

others, are only instrumentally good, or good as a means. Nevertheless, Moore would counter by acknowledging that the hedonist doctrine is misled due to confusion on the origin and cause of pleasure. For instance, to Moore, something, mainly beauty, may be considered valuable regardless of whether premeditated pleasure may be derived from it. In this way, pleasure and intrinsic good are not necessarily one and the same, nor are they interchangeable entities.

Likewise, as a premise, if Moore conceives that knowledge is a form of intrinsic desire, and its counterpart, ignorance, is a form of intrinsic bad, then being deprived of it is equivalent to an entity that is morally bad. As a result, the generation of hallucinatory stimulations counteract the individual's access to direct knowledge by obstructing their reality, and therefore, provokes an ignorance that mirrors suffering. Accordingly, Moore could then argue that the Experience Machine is not plausible since it is contingent on its ability to only provide pleasure. Hence, Nozick's condition of the absence of suffering throughout the duration of his theoretical thought experiment cannot hold. Furthermore, if a participant partakes in Nozick's experiment for the sole purpose of escaping the inevitable adversity that awaits every individual, Moore would maintain that the very act of willingly discarding one's awareness is the definition of suffering in itself. Therefore, the Experience Machine, in this specific instance, becomes obsolete.

If this premise is readily accepted by the public, as well as humanity in general, and pleasure is not taken to be the only thing that is intrinsically good, then Moore would hold that the inclusion of the extermination of one's family is an unnecessary condition for the experiment because the bereavement of authentic experiences, conscious awareness, and greater knowledge would be sufficient in deterring people from participating. From this, it is apparent that individuals do not simply yearn for what they experience, but rather what they physically do. In this way, humanity values authenticity instead of just the preservation of authenticity through

stimulatory means. Still, the original argument that one does not solely desire their own pleasure as an end continues to apply since the larger public prefers real, absolute and concrete pleasure in contrast to artificial pleasure.

In extension, if two individuals were to have identical experiences with the sole exception of one individual possessing real experiences, while the other receives stimulatory experiences, the individual that obtains authentic experiences has the advantage. This may be due to humanity's preference to action over internal state, as well as the compulsion to attain pleasure amongst the pain of reality, in contrast to deceitful means. These deceitful means omit the pain that serves to fundamentally delineate genuine pleasure—or pleasure that is crafted by the hands of the individual instead of that of the simulator. Additionally, deceit that takes on the form of masked awareness and consciousness whilst in the Experiment Machine further augments our entranced afflictions. By applying both Nozick's and Moore's argument, authentic experiences encompass insightful knowledge, that while bearing vulnerability to pain, hinder the larger threat of ignorant suffering.

Reconsidering Pascal's Wager as a Bet Against Atheism

Benjamin Nelson

In this paper I will review Blaise Pascal's argument to motivate the belief in god based on the expected utility of the believer in the afterlife. I will consider a few objections to his argument. I reject the wager's argument as one for motivating the belief in a Christian god, or any particular god or gods. As we will see later, its strongest objection is that it does not specify any particular god (or gods) and so becomes worthless. Ultimately, I show that the logic of Pascal's Wager, which is in unsuccessful and motivating belief in a particular god, can may be compelling at motivating belief against not believing in a god or gods.

Pascal's Argument

Blaise Pascal, writing in the 15th century, attempted to show that there is sufficient reason for believing in God based off a probability alone. The basic argument can be shown with the following table from Simon Blackburn (Cahn 183):

	God exists	God does not
Believe in God	+infinity	0
No belief/Do not believe	-infinity	0

The table can be read as follows. If I believe in God, and he exists, then the happiness that I will enjoy in the afterlife will be positively infinite in heaven. If I believe in God, and he does not exist, then when I die I will receive no added happiness and will just be a corpse in the ground. If I do not believe in God, and he exists, then I will suffer a utility of negative infinity in

hell. And lastly, if I do not believe in God and he does not exist then I will not suffer or be happy but, like the believer in the case God does not exist, I will just be a corpse in the ground.

Pascal argues that belief in a god maximizes the expected utility of one's afterlife. If I do not believe in God there is a chance he does not exist resulting in zero utility, but there is also a chance that he does exist that would give a negatively infinite amount of utility. If I choose to not believe, then zero utility is the best outcome I could have. Alternatively, if I choose to believe in God there is a possibility for infinitely positive utility if he exists and zero utility if he does not exist. Where zero utility was the best potential outcome for someone who does not believe, it is the worst potential outcome for someone who does believe.

At this point a word should be said about probabilities. It does not matter the particular probabilities for the existence/non-existence of God in this case. Even if the chance God exists is very small, the fact remains that if he exists and if I believe in him I will receive infinitely positive utility (and the opposite result if I do not believe), then my expected utility from believing will still be greater than the expected utility of not believing. Some theologians may also disagree with the use of positive/negative infinity when describing the outcomes of heaven or hell. All that matters here is that the absolute value of whatever number we attach is very very very large.

Objections to Pascal

There are quite a few objections to Pascal's argument for motivating belief based of probability. Linda Zagzebski summarizes three objections roughly as the "no chosen belief" objection, the "low view" objection, and the "many gods" objection (Cahn 184).

The "no chosen belief" objection criticizes Pascal by suggesting that we cannot really choose our belief, and that if we can choose to believe something then our motives might be

selfish (and, according to Zagzebski, morally suspect). Suppose a rich old man falls in love with a poor woman but she initially feels neither positive nor negative toward him. If she can get herself to love him then they will get married and she will certainly receive a lot of money as inheritance. Pascal might say, as he does for choosing to believe in God, that in order to fall in love with the man the woman should do all the things that lovers do; go on picnics, romantic walks, arguments in the car, etc. Now her chances of success vary with the circumstances, but it is her motivation that Zagzebski also has trouble with. To her, it seems that reprehensible that a person should fall in love with a person just for their money. Here I cannot say anything other than that while her motivation may begin with money in mind, by definition of her falling in love with the man she will value him for more than is monetary wealth and so I do not see an issue.

The second objection is that the wager presupposes a low view of god. This view sees the wager as against the spirit of religiosity in that it is selfish. As before with the poor woman falling in love with the rich man, Pascal does not see the wager as the end, but rather as a means to genuine belief. The one making the wage may fail to believe but they will not know unless they try.

The final objection, and the one that I think is most damning Pascal, is the “many gods” objection. The concern is that Pascal falsely assumes a dichotomy of choice between believing in the Christian God or not believing in a god at all. Of course this is not the case. It is possible to become a Jew, a Muslim, a Hindu, a pagan, practice an indigenous religion, and so on and believe in those respective gods. And so ultimately it seems that Pascal’s Wager does not motivate evidence for believe in the Christian God. However, given this evidence, I do think that Pascal’s train of thought should be abandoned. If the wager does not provide reason for believing

in the Christian god (or any god for that matter), it does seem to provide reason to not believe there is no god.

An Argument Against Non-Belief

Let us consider a version of the utility table that grants the possible existence of gods from different religious traditions (though not at the same time). The table as I present it is as follows:

	G1 exists	G2 exists	No god exists
Believe in G1	+infinity	-infinity	0
Believe in G2	-infinity	+infinity	0
No belief/Do not believe	-infinity	-infinity	0

Like the table before, it is read in the same way. If I believe in G1 and G1 exists then I will receive infinitely positive utility. If I believe in G1 and G2 exists then presumably I am banished to hell for believing false gods and receive negative infinite utility. Lastly if I believe in G1 and no god exists, then I receive 0 utility since I would just be corpse in the ground. The same line of reasoning is repeated in the row for believing in G2. And as we see in the row for not believing or withholding belief the result is infinitely negative utility in the event either G1 or G2 exist, and zero utility if god does not.

Further, the table only shows a G1 and a G2 that serve as blank stand-ins for any religion, but it could also be expanded to any number of options to choose from. Adding a third G3 column, etc. would not change the outcomes of my argument. This point is expanded upon in response to a particular criticism later below.

Now some particular points should be made about the argument. This argument is strictly against atheism and agnosticism.

This is not an argument meant to persuade toward the belief of any particular religion's god or gods. Per our table, it is possible to believe in a false god and still receive infinitely negative utility, but, without knowing what is in store in the afterlife there is still the possibility of being right and receiving infinitely positive utility. I leave the case for persuasion of belief for a particular religion for those particular missionaries.

This also is not an argument for the existence of any particular god or gods. It very well may be the case that no god exists and we rot in the grave for eternity. In which case every belief, even the nonbelief, would receive a utility of zero. But if we chose nonbelief we would be wagering the likelihood of zero utility against the likelihood of negatively infinite utility for every case. I leave the case for proof of existence to scientists and the theologians of particular religions.

What about the objection that this sort of self-interested belief is against god's wishes? Well my reply is that that objection is outside of the scope of the argument. First it is a theological question and not a philosophical one, but we can still consider it to defend my thesis. We could add another column to the table, say "G₁ Doesn't like selfish believers" and we could add another row for "Non-selfishly believes G₁" in which case they would receive infinitely positive utility in both the case of G₁ and G₁ Doesn't like selfish believers. So in fact it may be the case that the expected utility of genuine belief might be greater than the expected utility of a selfish belief, but the response does not give an outcome of infinitely positive utility to the nonbeliever. And so even selfish belief, based on expected utility, is still superior to nonbelief.

I create a G1 with subscripts here to emphasize the point that whether or not the god would react positively to a selfish believer is a theological question and we are not considered a separate god entirely. But, as indicated earlier in this paper, the same outcome would occur for the non-believer regardless if we are discussing various theological approaches (columns of Gs with various subscripts indicating different theological interpretations) or separate religions entirely (columns of Gs with different numbers).

Acknowledgements

I would like to thank everyone who submitted a paper for consideration both for the magazine and for our spring colloquium. Without the dedication of the undergraduate philosophy students, we would have no journal or colloquium to organize. I would also like to thank those members of Philosophy Club who helped me to organize and curate this magazine. Being a part of this was a wonderful and rewarding experience, and I'm honored to have had the opportunity.

I would also like to thank Professor Carrie Swanson, who has helped us each semester to produce this magazine as well as to hold our colloquium.

I encourage any and all undergraduates to submit papers in the future. Keep an eye out for a call for submissions both fall and spring semesters.

- Kate Lohnes