A strike against a striking principle

Dan Baras

Philosophical Studies
An International Journal for Philosophy
in the Analytic Tradition

ISSN 0031-8116
Volume 177
Number 6

Philos Stud (2020) 177:1501-1514
DOI 10.1007/s11098-019-01265-5

Springer

Philos Stud (2020) 177:1501–1514
https://doi.org/10.1007/s11098-019-01265-5

A strike against a striking principle

Dan Baras[1]

Published online: 25 February 2019

Abstract Several authors believe that there are certain facts that are striking and cry out for explanation—for instance, a coin that is tossed many times and lands in the alternating sequence HTHTHTHTHTHT… (H = heads, T = tails). According to this view, we have prima facie reason to believe that such facts are not the result of chance. I call this view the *striking principle*. Based on this principle, some have argued for far-reaching conclusions, such as that our universe was created by intelligent design, that there are many universes other than the one we inhabit, and that there are no mathematical or normative facts. Appealing as the view may initially seem, I argue that we lack sufficient reason to accept it.

Keywords Calling for explanation · Strikingness · Probabilism · Epistemic principles · Fine-tuning arguments

1 The striking principle and its motivation

Many believe that some facts are *striking* and therefore *cry out for explanation*—that is, they believe that some facts have a certain property, call it *strikingness*, and that we have reason to believe that such facts have an explanation, perhaps a special kind of explanation (i.e., that they are not coincidental). This idea is used to support a variety of far-reaching conclusions. Here are a few examples.[1]

[1] The italics in the quotations were added by me for emphasis.

✉ Dan Baras
dan.baras@mail.huji.ac.il

[1] The Martin Buber Society of Fellows, The Hebrew University of Jerusalem, Mt. Scopus, 9190501 Jerusalem, Israel

Derek Parfit argues that there are many universes other than the one we inhabit, because

> Of the range of possible initial conditions, fewer than one in a billion billion would have produced a universe with the complexity that allows for life. If this claim is true, as I shall here assume, there is something that *cries out to be explained*. (Parfit 1998)

The alleged fact that the initial conditions of our universe call for explanation is used as a strong reason to disbelieve that those conditions are coincidental. He then argues that if many universes exist, the fact that there exists one with the complexity that allows life is explained.

Roger White argues from similar premises to a different conclusion, namely that God exists, because

> That our universe is hospitable to life *stands in need of explanation*. That God adjusted the constants in order to allow for life to develop provides a satisfactory explanation for why our universe is life-permitting. (White 2018, 30)[2]

Thomas Nagel argues that science should go through a paradigm shift, reintroducing Aristotelean teleological causation, because

> Regularities, patterns, and functional organization *call out for explanation*—the more so the more frequent they are…the causation of consciousness by brain activity being a prime example. (Nagel 2012, 47)

Hartry Field argues that mathematical Platonism must be false and there are no mind-independent mathematical facts, because

> [T]he correlation between mathematicians' belief states and the mathematical facts… is *so striking as to demand explanation*; it is not the sort of fact that is comfortably taken as brute. (Field 1989, 26)

And Sharon Street argues that there are no mind-independent moral facts, because

> There is a striking coincidence between the normative judgments we human beings think are true, and the normative judgments that evolutionary forces pushed us in the direction of making. I claim that the realist about normativity *owes us an explanation of this striking fact*, but has none. (Street 2008, 207)[3]

[2] Both Parfit and White are more or less repeating in brief Leslie's (1989) fine-tuning argument. Leslie is undecided between two possibilities, God or Multiverse, as explanations for the fine-tuning of our universe. Parfit rules out theism on the basis of the problem of evil, while White rules out multiple universes, arguing that it doesn't provide a satisfactory explanation for the fact that calls for explanation.

[3] Enoch (2011, 158–160) develops the argument against robust normative realism similarly, and Schechter (2010, 2018) develops a similar argument for skepticism about logic.

I call the shared premise of these arguments, the claim that there are facts that are striking and call for explanation, *the striking principle*.[4] Shortly, I will suggest a sharper definition.

These authors seem to believe that our intuitions regarding a certain class of examples support the striking principle. Here are a few such examples[5]:

> Monkeys: Monkey1 is seated by a computer and types "mnev hwn ia sfiubafiuveil kjas n3i3 m I kcn"; that seems accidental. Monkey2, on the other hand, sits by a computer and types, "My name is Curious George, and The Man with the Yellow Hat is my friend." Surely it is implausible for Monkey2 to have typed that sentence by chance.

> John & Judy: Suppose I bring to your attention the fact that, throughout the past year, John and Judy were observed in close proximity to each other on countless occasions and in various locations, including at the movie theatre and at various cafés, clubs, and museums. The massive correlation between the locations of these two people over the past year seems too striking to leave unexplained. For instance, we may suspect that John and Judy are in a relationship.

> Pebbles: You walk along the shores of Alaska and notice that the pebbles there are all arranged in similar-looking rings. It seems that there must be some explanation for this phenomenon; it is unlikely to be the result of mere chance.

These examples nicely illustrate the most common types of facts that allegedly call for explanation: 1. Facts that have significance for human beings, such as a sentence in English; 2. correlations, like the correlation between John's and Judy's locations; and 3. patterns, such as rings or symmetrical repetitions. Most of the examples in the literature of facts that allegedly call for explanation share one or more of these features. The literature also includes examples that allegedly illustrate that there are facts that do not call for explanation. Monkeys typing nonsense, two people being located throughout the year in non-correlating locations and pebbles scattered in no identifiable pattern would not call for explanation.

While some version or other of this idea often appears in the literature, a clear definition of it is difficult to find. I suggest defining the striking principle as the conjunction of the following two claims:

(1) There is a distinction between facts that are striking and facts that are non-striking.[6]

[4] Mostly because they claim that there are facts that are striking, but also because I take the principle itself to be striking and mysterious.

[5] These examples are drawn, with some changes, from White (2005, 3) and Field (2001, 325). Many more fanciful examples can be found in Leslie (1989).

[6] More precisely, since strikingness, if it is a genuine property, is plausibly a graded property, the view is that facts can be striking or non-striking to various degrees. For simplicity, I will often just talk as if the view is that facts are either striking or non-striking simpliciter.

(2) To the extent that S is a purported fact that is striking, if theory T (or any conjunction of propositions) implies that S has no explanation,[7] then that is an epistemic reason to either reject T or reject the belief in S.

We can use any one of the above examples to illustrate the power of the striking principle. Monkey2, for instance, is a striking fact. If we received a report from a friend that a monkey sitting by a computer typed such a sentence, we should either conclude that the friend is lying (disbelieving the purported explanandum) or that the monkey is not typing randomly (rejecting a certain theory that we had assumed in the background). This inference seems intuitively plausible, and it therefore seems to support the striking principle, because the striking principle would, if true, explain why the inference is justified. Nevertheless, I argue that these intuitions do not, in fact, support that principle.[8]

Two notes on the distinction between striking and non-striking facts suggest that something is wrong with the striking principle. First, it is trivial and therefore redundant if all we can say about strikingness is the following:

> A purported fact S is *striking* iff and because: For any theory T, if T implies that S has no explanation, that is an epistemic reason to either reject T or reject S.

In order for the striking principle to be informative, there must be further constraints on what it takes for a fact to be striking. How, then, is strikingness to be understood? Is it to be understood as a fundamental irreducible property? That's a project in its own right, and it's one that striking theorists should engage. I note that an initial reason for skepticism about the striking principle is the difficulty of this project—it is unclear that there is any plausible non-trivial analysis of strikingness.[9]

Second, as illustrated by the above examples, when people think of facts that call for explanation, they usually think of very orderly facts. However, sometimes randomness calls for explanation as well. Suppose that I toss a stone, and instead of moving in a typical parabolic trajectory, it starts moving around randomly, up, down, right, left, and all around.[10] Intuitively, the crazy stone would call for explanation despite being quite the opposite of orderly. Moreover, it seems to call for explanation in the same sense that the aforementioned orderly sequences call for explanation.

---

[7] Or, if chance counts as an explanation, then "no explanation" should be replaced with "no explanation of a particular type, a non-chancy explanation". This is how White seems to think, and he suggests that the particular type of explanation required is a stable explanation. See White (2005).

[8] My argument is close in spirit to a recent complaint voiced by Hawthorne and Isaacs against the idea that certain phenomena cry out for explanation. Hawthorne and Isaacs suggest that "that whole notion of 'crying out for explanation' is a potentially confusing way of getting at some coarse grained facts about probabilities" (Isaacs and Hawthorne 2018, 143). However, Hawthorne and Isaacs are only making a brief comment in passing. In this paper, I develop this idea into an elaborate argument against the view.

It is also close in spirit to Fumerton's (1980, 2018) argument against the fundamentality of IBE. Fumerton argues that any justified IBE can be explained by some combination of enumerative inductions and deductions.

[9] See Baras (2018; unpublished ms.).

[10] The example is taken from Baras (unpublished ms.).

One lesson from this example is that order is not a necessary condition for calling for explanation. However, a more important lesson for my current purposes is that strikingness must be relative to an evidential state. It is clear that the crazy stone calls for explanation because of our prior empirical knowledge that ordinary physical objects don't behave in this way. Without such empirical knowledge, I doubt that we would have the intuition that the crazy stone calls for explanation. Because the evidential state is not intrinsic to the potential explanandum, strikingness cannot be an intrinsic property of facts, which may suggest that something is fishy about the striking principle. This too is not yet an argument, just a reason to be initially suspicious. As we will see, the alternative presented shortly fits well with this datum.

2 The structure of the argument (and premise 1)

I now present my argument against the striking principle, the structure of which is as follows:

(1) If we are justified in accepting the striking principle, it is because it is implied by the best normative explanation of our intuitions about examples.[11]
(2) There is an alternative explanation for all of the examples.
(3) The alternative explanation does not imply the striking principle.
(4) The alternative explanation is better than any explanation that implies the striking principle.
(5) Therefore, the striking principle is not implied by the best normative explanation of our intuitions about examples. (from 2 to 4)
(6) Therefore, we are not justified in accepting the striking principle. (from 1, 5)

My defense of premise 1 is as follows: In the literature, the striking principle is always introduced and motivated using examples. It seems that authors accept the striking principle because they feel compelled by the examples that illustrate it. Therefore, if the examples do *not* support the striking principle, it seems that we lack reason to accept that principle. I can think of two potential alternative ways of justifying the principle, and it would be interesting to explore whether they can be developed. One would be to show that the striking principle fares better than alternatives in solving theoretical puzzles regarding inductive inferences. The second would be to claim that some of us just have a direct intuition that the striking principle is correct and that that fact in itself lends support to the principle. I am not aware of either of these routes being explicitly pursued in the literature, although White (2005) comes close to the former (unfortunately, I cannot discuss White in more detail here). Regarding the latter strategy, I note that arguing for or against direct intuitions is a difficult and often futile task. For the purpose of the present paper, I set these possibilities aside and take premise 1 for granted. My focus from

[11] For discussion of the idea of a normative explanation, see Schroeder (2005).

here on will be premises 2–4. At the very least, I hope to show that the examples are not a good reason to accept the striking principle.

3 The alternative explanation (premise 2)

To explain the second premise,[12] I start with two toy examples, which are also paradigmatic examples of facts that call for explanation, and then return to the opening examples. Consider the following:

> Die: An ordinary looking die is tossed a hundred times. It lands on six on almost every single toss.

The fact that the die lands on six on most tosses would surely surprise us, and we should infer that it is no coincidence. That is, I accept the verdict of the striking principle in this case. However, I offer an explanation for this verdict that does not invoke the striking principle: Prior to having tossed the die, because it looks ordinary, we should be quite confident that each toss is equally likely to land on any side of the die and that the tosses are independent of each other. This belief is based on our previous observations of dice tosses and relevant background knowledge. That is, we believe a conjunction of propositions about the circumstances, C, that implies that the tosses are equal chance and independent. Given C, the probability of the result of the tosses, E, is exceedingly low.

However, we should not be fully confident in C. There is some probability that it is not really an ordinary die. For instance, it might be weighted on one of the sides. After tossing the die, we should disbelieve C because there is an alternative to C—call it H—that the die is weighted toward landing on six. H has some prior probability, even if tiny. However, conditional on H, E is much more probable than conditional on C. Therefore, E strongly confirms H.[13] After learning that E, we should believe H and therefore disbelieve C.

Things are not always that simple, however. In many cases, we should make such an inference even if we did not really know what precisely could explain the result. Consider another example:

> Coin: Imagine that you take a seemingly ordinary coin, toss it a hundred times and it lands HTHTHTHTHTHT…

[12] This section draws some inspiration from Horwich (1982, 100–104). Note though that Horwich's argument is very different than mine, and he is arguing against a principle that is different from the striking principle, though perhaps close in spirit, "that verification of relatively surprising consequences of a theory has especially great evidential value" (p. 100).

[13] One formal way to represent the degree to which E confirms H is as the ratio between the posterior of H and the prior. Because $\frac{P(H|E)}{P(H)} = \frac{P(E|H)}{P(E)}$ (a version of Bayes's law) and (focusing on the right side) E is much more likely conditional on H than the prior of E (when C is initially most probable), this ratio (on the left) is high as well.

Intuitively, prior to tossing the coin, you should have believed that the coin was an ordinary symmetrical coin, implying that for any given toss, it would have an approximately equal chance of landing on either of the two sides and that the results of individual tosses would have been independent of each other. We'll represent this belief again as C. However, after the tosses, you should revise this belief. Now you should believe that the tosses are dependent and that, for any given toss, it is highly likely to land on the opposite side of the previous toss. In this case, we have no idea what mechanism might be causing this behavior. Still, I argue, even if we don't have such a precise hypothesis, there is a disjunction of possibilities, or just the cruder claim that there is some mechanism involved, that would make the result, E, more likely if true. There might be some natural constraint that makes the coin land in such a pattern, or it might somehow be manipulated by some person.

Generally speaking, simple patterns are often produced by natural constraints and human manipulations tend to produce things that are of significance to us. We should be open to the possibility—that is, we should have some non-zero credence in the possibility—that there is something naturally constraining the coin tosses or that they are somehow being manipulated. Among the possible results of a hundred coin tosses, few will display simple patterns or sequences with human significance. Therefore, if the coin lands in such a pattern, we should be much less confident that it is ordinary and much more confident that there is some such explanation of its behavior.

Another way to put this is that, given our background knowledge, we should think it much more likely that if the coin is not an ordinary coin (¬C), then the possible sequences are not equiprobable. Rather, some sequences are more likely than others—namely, those that conform to patterns, are in correlation with some other factor, or are of particular significance.

The argument can be presented more formally as follows. Our intuitive judgment is that, in this case, learning E lowers the probability of C to below ½. What would it take for this to be the case? The following is a theorem of probability:

$$P(C|E) < \frac{1}{2} \quad \textit{iff} \quad \frac{P(C)}{P(\neg C)} < \frac{P(E|\neg C)}{P(E|C)}$$

Because the prior probability of C is high, in order for the inference to be justified, the probability of E conditional on ¬C must be significantly higher than the probability of E conditional on C. How much higher? That depends on how confident we initially should be of C. In our example, we know the value of P(E|C). There are 2^{100} possible results of tossing a coin a hundred times, and given C, they are equiprobable; hence, the probability of E is $1/2^{100}$. Now, in order for the inference to be justified, we need to be justified in believing that P(E|¬C) is higher. That means that we must be justified in believing prior to tossing the coin that if the coin and tosses are biased, then among the 2^{100} possible results, certain results are more likely than others, and E must be among the more likely ones. I suggest that we are justified in believing that this is the case because there are certain general hypotheses, incompatible with C and with some non-zero prior probability, such that if true, would make E significantly more probable than if C were true.

The explanations of the inferences in coin tosses and dice throws employ the basic principles of probability theory: the axioms, the standard analysis of conditional probability, and the Bayesian updating principle. The prior probabilities are, to a great extent, determined by our observations of other dice, coins and nature in general. For this, we need a principle that connects our beliefs regarding objective probabilities to subjective credences. The principle is sometimes called Miller's Principle (Strevens 2017, 34); David Lewis (1981) calls it the Principal Principle. We have well-established independent reasons to accept all of these principles. The striking principle, however, falls out of the picture.

A similar probabilistic explanation, I argue, is available for all of the examples in the literature that are supposed to motivate the striking principle. Given our background knowledge regarding monkeys, we initially have very low credence that the monkey will type any meaningful sentence, including "My name is Curious George…". We initially believe that the monkey chooses which key to press quite randomly. The number of possible sequences that the monkey can type is huge, so the probability of each is miniscule. However, we can attribute some prior probability—low but non-zero—to the possibility that the monkey's typing is somehow manipulated by another human. Perhaps the monkey was trained somehow, or there is some other way of manipulating monkeys of which we aren't aware. If the monkey types something that is of significance to humans, this fact would confirm the hypothesis that the monkey is being manipulated, as imprecise as the hypothesis may be.

In John & Judy, the hypothesis being confirmed is obvious: the hypothesis is that they are in a relationship. More generally, though, we can say that correlations tend to suggest a causal connection, if the context does not rule out such a causal connection. Finally, the rings in Alaska: Even if we have no idea how they are formed, we know that natural constraints tend to conform to simple patterns. The vague hypothesis that there is some natural constraint on how the pebbles rest is confirmed. I suggest that similar things can be said of all of the examples in the literature.

Finally, why is it that if a monkey types "mnev hwn ia sfiubafiuveil kjas n3i3 m I kcn", or two people are located throughout the year in non-correlating locations, or pebbles are scattered in no identifiable pattern, then we should not conclude that these facts are non-coincidental, that is, we should think that they are the result of chance? Take the monkey for example. True, the probability of a monkey typing this particular sequence conditional on it being an ordinary monkey is extremely low. And true, there are some hypotheses conditional on which the monkey typing this sequence would be significantly higher. For instance, if there existed an omnipotent demon that had a special liking of the particular sequence "mnev hwn ia sfiubafiuveil kjas n3i3 m I kcn", then it would be likely that this demon would take over the monkey's brain and bring it about that the monkey typed the loved sequence. However, this hypothesis, and any other hypothesis that comes to mind that would make this sequence more probable, has such low prior probability, that the posterior probability remains low as well. Or, coming at it from a different angle, we have no reason to believe that if the monkey is not an ordinary monkey, then it is more likely to type "mnev hwn ia sfiubafiuveil kjas n3i3 m I kcn" than other sequences. Therefore, the fact that the monkey typed "mnev hwn ia sfiubafiuveil kjas n3i3 m I kcn" is not a reason to doubt that the monkey is just an ordinary monkey typing at random.

Now, the fact that it is possible to explain these examples without invoking the striking principle is not yet a strong argument against that principle. It might still be the case that an explanation invoking the striking principle is a superior explanation, or that the probabilistic explanation is compatible with the explanations offered by strikingness theorists. I examine these possibilities in following sections.

4 A genuine alternative (premises 3–4)

Now that premise (2) has been established, I move to premise (3): that the striking principle is not implied by the alternative explanation. Taken at face value, the alternative explanation says nothing about strikingness or calling for explanation, so there's no obvious reason to believe that it does imply the striking principle. However, you may think that the striking principle itself can be given a Bayesian interpretation. That is, you may think that the distinction between facts that call for explanation and those that do not—and the claim about how this distinction can serve as a reason to reject some theories—can be fully reduced to less-mysterious probabilistic principles. If reducing the distinction in this way is possible, then it might turn out that the probabilistic explanation described earlier is actually compatible with and, in fact, implies the striking principle.

However, I do not think that this is the best way to think of the striking principle. Although proponents of the striking principle are not always clear about it, I believe that a more internally coherent interpretation of their view is that the striking principle is a fundamental epistemic principle that is not reducible to standard probabilistic principles.[14]

One reason to reject the suggestion that the striking principle is reducible is that if it *is* reducible, then fine-tuning arguments premised on the striking principle are vulnerable to a Humean objection.[15] This objection states that although we have observations and relevant background knowledge that support inferences regarding coins, dice, monkeys, and pebble formations, we have no such observations regarding the formation of universes. If strikingness is reducible to probabilistic considerations such as those that explain the guiding examples, it will not be easy to apply them to the boundary conditions of the universe. To the extent that some

[14] Another way to frame the question is to ask whether the striking principle, if true, is a fundamental epistemic principle or is derived from other, more fundamental principles. A similar question is debated regarding inference to the best explanation (IBE) and enumerative induction. The debate was started by Harman (1965), who argues that induction is reducible to IBE. Fumerton (1980) defends the opposite view, that IBE is reducible to induction. For more recent contributions to this debate, see McCain and Poston (2018).

[15] That is, an objection inspired by Hume's objection to a different version of the argument from design, as described in his *Dialogues Concerning Natural Religion* chap. 2.

advocates of the striking principle were trying to avoid such vulnerability, that is a reason not to interpret them in this reductive way.

Moreover, if the striking principle is reducible to the kind of probabilistic reasoning offered in the alternative explanation, then we are better off forgetting about the striking principle altogether and focusing our attention and reasoning on the less-mysterious probabilistic principles that appear in the alternative explanation. If the striking principle is to be of interest, then, for the following three related reasons, it must be an independent, non-reducible epistemic principle.

The first reason is that, for practical reasons, we should strive for theories with less principles. Anything achieved by the reductive version of the striking principle can be achieved by the more basic inductive principles, so it is difficult to see what benefit we gain by adding the striking principle to our list of epistemic principles. For that reason, theoretical parsimony (in the practical sense) suggests that we leave the striking principle out of our epistemic theories.

Second, the striking principle relies on terms with disturbingly unclear meanings. For example, it is unclear what it means to say that a fact is striking, or how we should identify when that is the case. Some attempts have been made to account for strikingness, however all existing attempts face difficulties.[16] Additionally, it isn't clear what it means to claim that a theory implies that a fact is unexplainable—and it is not a simple task to make precise the relevant meaning of "explainable". There are reasons to believe that not just any lack of explanation counts in this context. For example, a conjunction of all the positions and velocities of the individual tosses might count as some kind of explanation for the fact that a coin landed HTHTHT…, however, that's not the kind of explanation that is meant when it is said that such a sequence calls for explanation.[17]

Thirdly and relatedly, the striking principle can readily lead to confusion. On its face, it seems to support inferences that are not based on past observations or experience in the way that the alternative explanation is. For example, it seems that we can infer that the fine-tuning of our universe calls for explanation regardless of whether we have any relevant experience (say, of beginnings of universes). However, if the striking principle is to be understood as reducible to the kind of reasoning presented in the alternative explanation, then much more needs to be said to justify its application to the initial conditions of our universe than just that those conditions are striking. In order to avoid potential misapplications of the striking principle, we are better off tossing the principle in the theoretical garbage bin and sticking to the more precise axioms of probability and Miller's Principle.

For these reasons, I believe that the striking principle should be interpreted as a candidate for a fundamental epistemic principle. Interpreted as such a principle, the striking principle is not implied by the alternative explanation.

I move now to premise (4). Premise (4) is that the alternative explanation is better than any explanation that involves the striking principle. To defend this premise, I will explain why I think the alternative explanation is better than the explanation

[16] See above, footnote 9.

[17] For proposals, see Field (1996, 370), White (2005, 4) and Baras (2017).

that striking theorists provide for the opening examples. First, it equally well explains our intuitions regarding the examples that motivate the striking principle. Second, it draws on principles that we have independent reasons to accept, whereas the striking principle lacks independent support. Finally, for the same reason, the alternative explanation is more theoretically parsimonious, as it does not force us to add a new epistemic principle to our list of accepted principles. For these reasons, the alternative explanation described in Sect. 3 is superior to any explanation that relies on the striking principle.

This concludes my case for the premises of the argument. If those premises are correct, then the striking principle lacks an epistemic foothold, and striking theorists misidentify the epistemic principle that explains our different judgments of examples. Therefore, arguments premised on the striking principle are unsound.

Is it possible to revise those arguments so that they maintain their initial motivating intuitions but rely on a more plausible epistemic principle? I address this worry in the final section.

5 A final worry: an easy fix to the striking arguments?

I began this article by mentioning a number of arguments premised on the striking principle. You may worry that even if the argument of this paper is sound, and we reject the striking principle, that doesn't make much of a difference to the arguments mentioned. Perhaps those arguments as formulated are unsound; however, there may be a simple way to reformulate them all such that they do not rely on the striking principle. If so, though it's helpful to know that we should reject the striking principle, the conclusions of the arguments in which the striking principle was invoked do not lose much of their epistemic support.

One possible way to do that is to reformulate the arguments probabilistically. However, there are obstacles awaiting such reformulations that require additional argumentative moves. In the case of fine-tuning arguments, as is well known, it is a difficult task to figure out what it means to claim that a certain condition of our universe has a certain probability and how such claims can be justified.[18] In the case of reliability challenges such as Field's argument against mathematical Platonism, a case can be made that it was highly probable that we would arrive at the relevant set of true beliefs and this makes it difficult to understand precisely what the reformulated argument might be (Clarke-Doane 2016, 29–30; Baras 2017). Perhaps part of why the mentioned authors avoided probabilistic arguments is a wish to avoid such worries. I cannot explore these issues here. My brief comments should suffice though as an explanation of why rejecting the striking principle advances the relevant debates even if probabilistic arguments are available.

[18] For a recent probabilistic fine-tuning argument, see Isaacs and Hawthorne (2018). Notice how complicated their argument is in relation to White's (2018) argument, for example.

A different strategy, closer in spirit to the original arguments, is to rely on a version of inference to the best explanation that does not rely on the distinction between striking and non-striking facts. In particular, it seems plausible that if there are two competing theories, one of which explains a fact that the other does not explain, then that is some reason to believe the former. Or more generally, if there is a theory that implies that a certain fact has no explanation, and there are alternative theories that would, if true, explain that same fact, that is a reason to reject the former theory.[19] Thus, for example, there is some reason to reject the undesigned mechanistic single-universe theory because this theory implies that the initial conditions of our universe are coincidental whereas other theories, such as the multiverse theory or theism do not have this implication. Similarly, there is some reason to reject mathematical Platonism because the theory implies that the correlation between the beliefs of mathematicians and the mathematical truths is coincidental, whereas structuralism and fictionalism don't have this implication. The same conclusions are supported without distinguishing between striking and non-striking facts.

My response is that the aforementioned authors do not just want *some* epistemic reason; they want a strong epistemic reason, strong enough for us to accept their conclusion, not just raise our confidence by a tiny bit. Even if we accept the general idea that failing to explain a fact counts against theories, the strength of the epistemic reason need not always be significant. After all, all theories leave some facts unexplained.[20]

Consider the following die example. A die is tossed 15 times and lands 621626526626412. One possibility is that the die is an ordinary die and the sequence was the result of pure chance. An alternative theory is that there's an invisible powerful being with the ability to control the die, and this being has a special liking to the sequence 621626526626412. While it is true that the latter theory would explain the sequence if true, whereas the former would not, it doesn't seem like that's much of a reason to accept the latter theory or even to reject the former.

This shows that even when a theory fails to explain some fact, that doesn't always count significantly against it. Now there's a theoretical challenge to explain when and why a failure of explanations counts *significantly* against a theory.[21] The striking principle proposes an answer to this question. It says that the more a fact is

[19] I thank an anonymous referee for posing this challenge to me.

[20] Field, for example, explicitly agrees that there are certain facts that Platonists can plausibly accept as brute. He claims however that there is something special about the reliability of mathematicians that seems "altogether too much to swallow" (Field 1989, 26). In the fine-tuning arguments, where we're discussing the fundamental conditions of the universe, every theory will have to posit some brute facts. Therefore, to make any fine-tuning argument, a distinction between facts that we can leave unexplained and facts that we cannot leave unexplained is needed.

[21] Counting the number of facts that a theory leaves unexplained won't provide an answer. For instance, many coins were flipped since the first coin was invented, forming many random sequences of tosses. Each of these sequences is a fact. Yet it doesn't seem to count significantly against a theory if it doesn't explain this very large number of facts. Arguably, there are infinite such uninteresting facts that are unexplained (in the relevant sense) by all of our best theories.

striking, the more a failure to explain it counts against a theory. Once we reject the striking principle, further argument is needed to support the conclusion that we have a significant reason to reject, say, the undesigned mechanistic single-universe theory just because it implies that the initial conditions of the universe are unexplained.

This, I believe, is why the mentioned authors did not just say that the theory they are arguing against fails to explain some fact. They felt a need to say that the explanandum is striking and calls for explanation, thinking that that makes the degree of disconfirmation especially strong. Therefore, if the argument of this paper is sound, much more is needed to rescue their conclusions. That there is some fact unexplained by a certain theory is not enough of a reason to reject that theory, even if it is some epistemic reason to decrease confidence.

Acknowledgements Many thanks to Sharon Berry, David Enoch, Yehuda Gellman, Eli Pitcovski, Joshua Schechter, Ruth Weintraub and Preston Werner for helpful comments on previous drafts, and to Roger White, Miriam Schoenfield as well as participants in my presentations at Tel Hai College and at the Haifa Conference on the Philosophy of Religion (2018) for helpful discussion.

References

Baras, D. (unpublished ms.) Do extraordinary types cry out for explanation?.

Baras, D. (2017). Our reliability is in principle explainable. *Episteme, 14*(02), 197–211. https://doi.org/10.1017/epi.2016.5.

Baras, D. (2018). Why do certain states of affairs call out for explanation? A critique of two Horwichian accounts. *Philosophia*. https://doi.org/10.1007/s11406-018-0047-x.

Clarke-Doane, J. (2016). What is the Benacerraf problem? In F. Pataut (Ed.), *New perspectives on the philosophy of Paul Benacerraf: Truth, objects, infinity* (pp. 17–43). Dordrecht: Springer. https://doi.org/10.1007/978-3-319-45980-6_2.

Enoch, D. (2011). *Taking morality seriously*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199579969.001.0001.

Field, H. (1989). *Realism, mathematics, and modality*. Oxford: Basil Blackwell.

Field, H. (1996). The a prioricity of logic. *Proceedings of the Aristotelian Society, 96,* 359–379.

Field, H. (2001). *Truth and the absence of fact*. Oxford: Oxford University Press. https://doi.org/10.1093/0199242895.001.0001.

Fumerton, R. (1980). Induction and reasoning to the best explanation. *Philosophy of Science, 47*(4), 589–600. https://doi.org/10.1086/288959.

Fumerton, R. (2018). Reasoning to the best explanation. In *Best explanations* (pp. 65–78). Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198746904.003.0005.

Harman, G. H. (1965). The inference to the best explanation. *The Philosophical Review, 74*(1), 88–95. https://doi.org/10.2307/2183532.

Horwich, P. (1982). *Probability and evidence*. Cambridge: Cambridge University Press.

Isaacs, Y., & Hawthorne, J. (2018). Fine-tuning fine-tuning. In J. Hawthorne, M. A. Benton, & D. Rabinowitz (Eds.), *Knowledge, belief, and god: New insights in religious epistemology*. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198798705.003.0008.

Leslie, J. (1989). *Universes*. London: Routledge.

Lewis, D. (1981). A subjectivist's guide to objective chance. In W. L. Harper, R. Stalnaker, & G. Pearce (Eds.), *IFS. The University of Western Ontario series in philosophy of science* (Vol. 15, pp. 267–297). Amsterdam: Springer. https://doi.org/10.1007/978-94-009-9117-0_14.

McCain, K., & Poston, T. (Eds.). (2018). *Best explanations*. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198746904.001.0001.

Nagel, T. (2012). *Mind and cosmos*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199919758.001.0001.

Parfit, D. (1998). Why anything? Why this? *London Review of Books*, *20*(2), 24–27.

Schechter, J. (2010). The reliability challenge and the epistemology of logic. *Philosophical Perspectives, 24*(1), 437–464. https://doi.org/10.1111/j.1520-8583.2010.00199.x.

Schechter, J. (2018). Is there a reliability challenge for logic? *Philosophical Issues*. https://doi.org/10.1111/phis.12128.

Schroeder, M. (2005). Cudworth and normative explanations. *Journal of Ethics and Social Philosophy, 1*(3), 1–27.

Street, S. (2008). Reply to copp: Naturalism, normativity, and the varieties of realism worth worrying about. *Philosophical Issues, 18*(1), 207–228. https://doi.org/10.1111/j.1533-6077.2008.00145.x.

Strevens, M. (2017). Notes on Bayesian confirmation theory. http://www.nyu.edu/classes/strevens/BCT/BCT.pdf. Accessed 24 Feb 2019.

White, R. (2005). Explanation as a guide to induction. *Philosophers' Imprint, 5*(2), 1–29.

White, R. (2018). The argument from cosmological fine-tuning. In G. Rosen, A. Byrne, J. Cohen, S. V. Shiffrin, & E. Harman (Eds.), *The Norton introduction to philosophy* (2nd ed., pp. 29–35). New York: Norton.