DOI: 10.1111/nous.12502

ARTICLE

Noûs

Indexicality, Bayesian background and self-location in fine-tuning arguments for the multiverse

Quentin Ruyant

Departamento de lógica y filosofía teórica, Facultad de filosofía, Universidad Complutense de Madrid

Correspondence
Quentin Ruyant, Departamento de lógica y filosofía teórica, Facultad de filosofía, Universidad Complutense de Madrid.
Email: qruyant@ucm.es; quentin.ruyant@ikmail.com

Funding information
H2020 European Research Council, Grant/Award Number: 101022338

Abstract

Our universe seems to be miraculously fine-tuned for life. Multiverse theories have been proposed as an explanation for this on the basis of probabilistic arguments, but various authors have objected that we should consider our total evidence that this universe in particular has life in our inference, which would block the argument. The debate thus crucially hinges on how Bayesian background and evidence are distinguished and on how indexical or demonstrative terms are analysed. The aim of this article is to take a step back and examine these various aspects of Bayesian reasoning and how they affect the arguments. The upshot is that there are reasons to resist the fine-tuning argument for the multiverse, but the "this-universe-objection" is not one of them.

Our universe seems to be miraculously fine-tuned for life: if the value of cosmological constants were slightly different, it would not have the same long lasting structure, and life would be impossible. How can we explain this? Perhaps this is because God wanted there to be life, and designed our universe in this way. But there is an alternative explanation that is apparently just as good: the multiverse hypothesis. The idea is that if there are many universes, each with different features (for example values for cosmological constants), then the fact that some of these universes are suitable for life is not miraculous. And of course, as living organisms, we happen to be in one such universe, which is not surprising either.

This fine-tuning argument for the multiverse (FTAM) can be expressed within a Bayesian framework: the probability of life in the cosmos is very high under the multiverse hypothesis, but very low under the single-universe hypothesis, which entails that the former is favoured by our evidence that there is life. But it has been objected (originally by Hacking (1987) in the context of a sequential multiverse) that this is an instance of an inverse gambler's fallacy: the kind of fallacy one commits when one enters a casino, sees someone roll dice and obtain a double-six, and infers that there must have been many dice rolls before this one for this unlikely success to happen. Although more dice rolls do imply a higher probability that a double-six will be obtained at some point, it does not imply a higher probability that the particular dice roll that one was looking at resulted in a double-six. And, the objection goes, the same could be said in the case of fine-tuning: the existence of many universes does not affect the probability that *our* universe is suitable for life.

A natural response to this objection is to recall that a selection effect is involved in FTAM. Imagine that I am only allowed to enter the casino just after a double-six has been rolled. Then I would be right to infer that there have probably been many rolls before the one that allowed me to enter, whichever it is. Similarly, we can only be the witness of a universe if it permits life, because life is required for our own existence. So, accordingly, we are right to infer that there are many universes, because it raises the probability of our own existence, in whatever universe this eventually happens. But according to White (2000), this reasoning is fallacious: in the case of the casino, any double-six is enough for me to enter the casino, but, he argues, in the case of fine-tuning, I can only be a witness of life if the constants are right in *this* universe. And in this respect, the existence of other universes is irrelevant.

This argument, which I will call the this-universe-objection (TUO), has generated various responses examining subtleties of the fine-tuning argument for the multiverse and of the inverse gambler's fallacy (Isaacs, Hawthorne, & Sanford Russell, 2022; Friederich, 2019a; Bradley, 2009; Juhl, 2007; Holder, 2002), but it remains influential (among its defenders are Goff (2019), Metcalf (2018) and K. Draper, Draper and Pust (2007)). Part of this dispute is apparently methodological: even when the opponents in this discussion all adopt a Bayesian framework, they often approach the framework differently, having in mind specific conceptions of the significance of indexicals (*our* universe), of rigidity, and of the implied division between background knowledge and evidence.

The aim of this paper is to take a step back and examine these various aspects of Bayesian reasoning and how they affect the arguments. The upshot is that in so far as we are willing to stay close to scientific practice in the way our Bayesian inferences are carried out, there are many reasons to resist FTAM, but TUO is not one of them.

1 | THE BAYESIAN FRAMEWORK

Many fine-tuned features of our universe can prompt a fine-tuning argument for the multiverse. However, I am mostly interested in this paper in the *form* that FTAM takes, and more precisely in probabilistic formulations of the argument, so the detail of how our universe is fine-tuned for life will not matter (I will just talk about the value of cosmological constants for illustration purposes). All that will matter is that the probability (in a sense to be defined) of our universe displaying these features is low. Having said that, these fine-tuned features of our universe are part of our scientific knowledge, and so, it could be thought that FTAM takes place in a scientific (or inductive, or naturalised) metaphysics methodology. A problem that I want to put forth is that the reasoning that is adopted in discussions is not always congruent with scientific reasoning (including in cosmology). My aim is to examine how the arguments and objections would transpose in a framework with more scientific rigour.

Probabilistic formulations of FTAM adopt a Bayesian framework. Although the use of Bayesianism is quite standard in analytic philosophy, this is already somehow problematic, since arguably, science is more in the business of confirming or discarding hypotheses on the basis of severe tests (Mayo & Spanos, 2006) than in the business of evaluating degrees of credence, whatever such entities could be. In my view, philosophers who adopt a naturalistic perspective should believe in the multiverse only in so far as cosmology confirms this hypothesis and discards the single-universe hypothesis and remain agnostic otherwise. Since the single universe hypothesis has not been discarded by cosmologists so far, naturalist philosophers should remain agnostic. But this might be a contentious methodological issue, both for science and for philosophy, and in any case, this Bayesian framing is shared by proponents and objectors of FTAM, so I will adopt it here for the sake of the discussion.

If we adopt the Bayesian framework, the main question is how to qualify the elements of Bayesian reasoning.

The general spirit of the Bayesian framework is the following: we have a set of probabilistic theories or hypotheses $\{T_i\}$, which, assuming that background circumstances B are present, allow us to attribute definite likelihood probabilities $Pr(O|T \wedge B)$ to all possible observations in a set $\{O_j\}$. These probabilities are simply deduced from the theory and background circumstances. They are naturally interpreted in terms of *objective (e.g. physical) likelihood probabilities* postulated by the theory. This can mean either that there are real propensities in nature, probabilistic laws governing natural phenomena, or that there are statistical patterns in the actual distribution of properties of the world, or in the initial conditions of the universe: the metaphysical interpretation will not matter much for our purpose (although we will see, at the end of this paper, one way in which it could affect the debate). In all cases, Bayes's theorem basically allows us to "reverse" these likelihood probabilities, so as to evaluate the probability $Pr(T|O \wedge B)$ of any theory $T \in \{T_i\}$ given an observation $O \in \{O_j\}$, still assuming B as a background. According to Bayesianism, if we observe O at t_1, we should update our credence in the theory according to the following formula:

$$Pr_{t_1}(T|B) = Pr_{t_0}(T|O \wedge B) = \frac{Pr(O|T \wedge B) \times Pr_{t_0}(T|B)}{\sum_i Pr(O|T_i \wedge B) \times Pr_{t_0}(T_i|B)}$$

This only requires having prior degrees of credence $Pr_{t_0}(T_i|B)$ associated with all theories in the set at t_0, and being able to derive a likelihood $Pr(O|T_i \wedge B)$ for what was actually observed from any theory in the set (the prior probability for our evidence $Pr(O|B)$ is here expressed in terms of a sum over theory predictions for this observation weighted by priors; this allows to avoid inconsistencies that could result in having probabilities higher than 1 for instance).

If we do not have an exhaustive set of theories, or if we don't want to attribute specific degrees of credence to our theories, we can instead focus on Bayes factors. Subjective Bayesianism is not really used to evaluate hypotheses in science, but Bayes factors are occasionally used (here is an example from psycholinguistics: Sagart et al. (2019)). They represent the extent to which one theory is favoured over another one by a specific piece of evidence, assuming background circumstances, and this is just the ratio of the respective likelihoods of the evidence predicted by the theories under consideration:

$$BF_{12} = \frac{Pr(O|T_1 \wedge B)}{Pr(O|T_2 \wedge B)}$$

If the factor is much higher than 1, then the first hypothesis T_1 is favoured by the evidence O. If it is close to 0, then the second one T_2 is favoured (a logarithm of this value can also be used: the diagnosis then depends on whether it is positive or negative). The advantage of Bayes factors is that we do not have to invoke arbitrary or subjective elements in their calculation.

What is generally assumed in the debate on fine-tuning is that we have two hypotheses to evaluate: multiverse (T_M) or single universe (T_U), and that the relevant evidence has to do with the fact that there is life in the universe. Fine-tuning, or the fact that only specific cosmic conditions permit life, is assumed to be part of the background circumstances B. Intuitively, the thought is that under these circumstances, T_M makes life in at least one universe (O) more likely than T_U, which means, using a Bayesian reversal, that the multiverse hypothesis is made more probable by the fact that life exists.

$$Pr(O|T_U \wedge B) \ll Pr(O|T_M \wedge B)$$

Therefore,

$$BF_{UM} \ll 1$$

But as we have seen, according to TUO, this is an instance of an inverse gambler's fallacy: the multiverse hypothesis cannot raise the probability that *this* universe in particular has life, not any more than seeing a double-six when entering a casino can raise the probability that there were many rolls before. The objection hinges on how O and B are characterised.

The requirement of total evidence plays an important role in this reasoning. This is the requirement that when doing Bayesian inference, we should take into account the strongest evidence we have, otherwise we could be led astray. White (2000) gives the following example:

> Suppose I'm wondering why I feel sick today, and someone suggests that perhaps Adam got drunk last night. I object that I have no reason to believe this hypothesis since Adam's drunkenness would not raise the probability of me feeling sick. But, the reply goes, it does raise the probability that someone in the room feels sick, and we know that this is true, since we know that you feel sick, so the fact that someone in the room feels sick is evidence that Adam got drunk.

This reasoning is fallacious, and it is the fact that only partial evidence is taken into account that leads us astray. In the same way, we should take into account, as evidence O, the strong claim that *this* universe is suitable for life, and not the weaker claim that *at least one* universe suitable for life exists. What we want, according to this reasoning, is an explanation for why there is life in *this* universe, given that this universe exists. So, the existence of this universe should be included in our background knowledge B. And since the existence of other universes than ours is irrelevant to the likelihood of anything that happens inside ours (assuming that universes are causally disconnected), the likelihood of life in *this* universe is not raised by the multiverse hypothesis (K. Draper, Draper & Pust 2007).

$$Pr(O|T_U \wedge B) = Pr(O|T_M \wedge B)$$

Therefore,

$$BF_{UM} = 1$$

This has generated various discussions regarding the soundness of the requirement of total evidence (Epstein, 2017; P. Draper, 2020; Barrett & Sober, 2020). I won't discuss it here, because I am convinced that it is a sound requirement *in principle*. My response to TUO does not rest on questioning this requirement. However, one point that can be made with respect to it, in order to anticipate the following sections, is that for practical purposes, when reasoning probabilistically, we always focus on the weak evidence that we think is most relevant, and discard almost everything else. So, maybe "requirement" is too strong a qualification: if taking into account our total evidence is *sufficient* to avoid reasoning mistakes, it is seldom *necessary*. The important question is therefore: when are we allowed to discard some strong evidence in order to focus on weaker evidence? And the answer depends on how we qualify the T, B and O that are involved in our reasoning (which crucially depends, in the case at hand, on how we analyse the semantics of demonstratives such as "this" in "this universe").

Recently, Friederich (2019b) has attempted to resist TUO by considering that the background B of our inference already contains not only the fact that this universe exists, but also the fact that it contains life, and that our evidence O is not that there is life, but that our universe is fine-tuned for life. According to him, this would correspond more closely to our epistemic situation: we knew life existed before learning that only very specific features allow for it, and therefore, we should update our credence by taking into account fine-tuning as evidence against the background that life exists rather than the other way around. So, we can see, again, that how T, B and O are characterised plays a crucial role in the debate. The remainder of this article is dedicated to examining each of these components in turn (admittedly in a rather opinionated way at times).

2 | THEORIES, HYPOTHESES AND MODELS

Let us focus on T first. There, talking in terms of hypotheses being confirmed or not is a bit too simple. It is now commonplace to assume that scientific theories are not sets of statements about the world, not linguistic entities, but rather families of models (Suppe, 1989; Winther, 2020). In this respect, cosmology should be taken to be an activity of model construction. General hypotheses are too abstract to generate directly the likelihood probabilities that we need for our Bayesian calculation, but models can. Therefore, it would be better to frame the debate in terms of cosmological models and their accuracy.

For the sake of simplicity, we can consider an abstract model M_M where many universes coexist, each one of them with their own set of cosmological constants, where the constants are determined by a probabilistic process. It will not really matter if these universes are parallel or sequential. Then we can consider an abstract model M_U that contains a single universe with definite cosmological constants. But here, we can see a problem for the framing of the discussion: should we assume that there are as many candidate M_U as possible values for the constants, each being a potentially accurate representation of our cosmos? Should we assume that cosmological constants are *model parameters* for M_U to be filled in during concrete applications? Or should we assume that according to M_U, the constants are fixed by a probabilistic process, as in M_M?

In a sense, the last option already concedes too much to the multiverse theorist. Saying that these constants could have had different values, in a physical sense of possibility, is questionable (see Colyvan, Garfield & Priest, 2005). This postulate is not at all required by contemporary cosmology. On the other hand, it is not obvious that the two first options give rise to a fine-tuning problem. Under the first option, we would have one specific single-universe model with the right constants for life M_L which would be favoured by our evidence that there is life, at least as much as the multiverse model ($BF_{LM} \geq 1$). Then a principle of parsimony (perhaps baked into our priors) could make us prefer this particular single-universe model rather than the multiverse model in the absence of positive evidence for the latter. On the second option, we would fill M_U with the right values for cosmological constants on the basis of our measurements to the same effect. In both cases, this would amount to saying that our universe just has these constant values, that's what we observe, and that explanations must stop somewhere. And this undermines FTAM (except, interestingly, Friederich (2019b)'s version of the argument).

Now intuitively speaking, the fact that our universe has the right cosmological constants for life still looks like an odd coincidence. A way to save this intuition is to consider that M_U is actually a disjunction of models, or that an hypothesis in general can be made to correspond to the disjunction of the models that are compatible with it, and that we should evaluate the posterior probability of this disjunction in light of our evidence. Then the point that can be made is that our observation that life exists, given fine-tuning (or the other way around), should lower our credence in this disjunction of models, because very few of them are compatible with life. But firstly, I think that this does not fit very well with the spirit of Bayesianism, where the point is to update our credence by "reversing" a theoretical likelihood, not by reversing another degree of credence. There is a lack of homogeneity in the reasoning if constant values are attributed a subjective degree of credence under the single universe hypothesis, but an objective likelihood under the multiverse hypothesis. This is a minor point though (a convinced Bayesian would probably say that we are applying Bayes's formula to degrees of credence in all cases, it is just that they are deduced from likelihoods using Lewis's principal principle in the second case). A more problematic issue is that the whole argument now rests on an attribution of prior probabilities to models, which requires adopting a measure on the space of possible values for constants, and then applying a principle of indifference. Why we should accept this framing is questionable (although note that there is a similar "measure problem" for the multiverse hypothesis (Wenmackers, 2023)).

The problem, in a sense, is that we are trying to compare two models with different scopes: one, the multiverse model, seeks to explain where cosmological constants come from, and the other, the single universe model, takes constant values for granted and does not provide any explanation (it will be enough here to understand an explanation as something that makes the explanandum unsurprising or likely given postulated circumstances: then the connection between explanations and Bayesian reasoning is straightforward). Until now, we were implicitly taking M_U and M_M to be models of the entire cosmos, with the same target of representation. However, a target of representation for a theoretical model is normally *something to be explained* by the model. So, if really M_U and M_M must have the same target of representation, they must provide alternative explanations for the same thing, which is not the case here, since M_U does not explain the value of cosmological constants (structurally speaking, M_U could actually be contained in M_M).

A fairer comparison, and a more legitimate application of Bayesianism, would involve a model with a single universe that *explains* the value of its cosmological constants (or any fine-tuned feature), for comparison with the multiverse explanation. And the problem for FTAM is that there might be many such potential models, that is, many potential explanations for the value of cosmological constants involving a single universe, *including many yet unconceived ones*. If this is so, the single universe hypothesis is not necessarily threatened by fine-tuning in comparison with the multiverse hypothesis (although I think it would remain better to compare specific models rather than general hypotheses). It is not even clear to what extent the coincidence in cosmological constants is in need of an explanation, or to what extent it could be an artefact of contemporary theories that will disappear with new theories. Unrestricted Bayesianism implicitly assumes that we have access to the full range of potential explanations, but this is hardly the case: our

M_M and M_U (which, by postulate, integrate the notion of cosmological constant of contemporary cosmology and their physical role) are only two models among infinitely many more, and so, the putative fact that one would be strongly favoured by our evidence over the other has very limited philosophical significance in the end.

Another way to make the same point is that fine-tuning is troubling because we currently lack any good potential single-universe physical explanation for why cosmological constants have the very particular values that they have. This is why more extravagant explanations (design and multiverse) are considered. But it is not clear at all that Bayesianism is the appropriate framework to address this kind of epistemic situation, where a proper explanation is missing. Maybe its use should be restricted to the comparison of the well-defined potential explanations that we have in a specific context, instead of using it to choose between accepting to live without any explanation (the single universe hypothesis) or endorsing extravagant ones (such as the multiverse hypothesis).

Unfortunately, these issues are obscured by the fact that the discussion is framed in terms of abstract and vague hypotheses with universal scope rather than in terms of specific explanatory models within a limited scope. I am personally very sceptical that any conclusion can be reached in such abstract settings, or that the Bayesian framework is appropriate at all, and for this reason, I think that FTAM is deeply problematic. However, I will put aside these worries here, because they are apparently not shared by participants in both camps of the debate.

So, I will assume that the question is how to evaluate two abstract cosmological models, M_U and M_M, each attributing likelihood probabilities to the constant values for universes, in light of the fact that only specific values allow for life. The question is: in this context, to what extent our evidence should favour one model over the other? (TUO could also be reformulated as an objection against the idea that the multiverse explains our evidence at all, in which case it is legitimate despite the worries expressed in this section: the rest of the discussion will be relevant to assess this kind of objection.)

3 | BACKGROUND KNOWLEDGE

After theories T, the second component of Bayesianism that needs to be qualified is the background B.

A Bayesian inference is used to evaluate the probability of a theory given some available knowledge. This knowledge is divided into two parts by Bayes's formula: the part that is actually used to evaluate the theory, and the part that will be kept in the background of our reasoning. The way the division is made seems to have important consequences for the result of the fine-tuning inference: we saw that objectors to FTAM consider that "this universe exists" is part of the background, and conclude that the multiverse is not favoured by our evidence, and that others go as far as claiming that "there is life in this universe" is also part of the background, but do not reach the same conclusion. However, it is not clear who we should follow in these matters.

There is indeed a lot of confusion about how this division between background and evidence should be made. The background is sometimes supposed to correspond to old evidence (acquired before what is counted as new evidence). The thought is that the division should match our epistemic situation, or at least a fictional epistemic situation, using what is sometimes called ur-priors (Monton, 2006; Meacham, 2016). Sometimes it is claimed that if it would be logically impossible for us to be in an epistemic situation where something is not the case, then this something should be part of the background: for example, "I exist" (see K. Draper, Draper & Pust, 2007, p. 301). But nothing like this is actually required.

From a purely calculational point of view, if we adopt standard Bayesianism (not the Bayes factors: I will return to them later), then what counts as background knowledge and what counts as evidence is somehow arbitrary. In the end, Bayes's formula is a valid theorem of probability theory, and the probability that it allows to calculate is always the same: it is the probability $Pr(T|O \wedge B)$ of a theory given our total knowledge, which includes *both* background and evidence. The division between background and evidence only affects what will be taken as basic material for the calculation: which priors $Pr(T_i|B)$ and which theoretical likelihoods $Pr(O|T_i \wedge B)$ will be used. But whether the division follows the order in which information was acquired, or whatever other consideration, is not important as long as we agree on the value of these basic probabilities (priors and likelihoods): if we are careful in our calculation, the end result will always be the same, however the division is made.

What actually matters is much more practical. When dividing our total knowledge into background and evidence, we should ensure *the availability of the probabilities that we need for calculating the probability that we want* (I take this nice formulation from a blog post by Barnes (2013)). The probability that we want is the probability of the theory under evaluation given our total knowledge. The ones that we need are the prior probabilities of theories and the likelihood they assign to possible observations, *given the background*.

Focusing on likelihood probabilities, which is the main component of our calculation, this means that the background B must contain the kind of knowledge that is required in order for our theories or hypotheses to be applicable and yield well-defined likelihood probabilities (generally associated with causal processes). And our evidence O must be the kind of observation for which the theory or hypothesis yields well-defined likelihood probabilities (the result of these processes). This typically corresponds to an experimental context and the outcome of the experiment respectively. What does not fall into any of these two categories can be ignored, because it does not affect the calculation. This is how the division is done in practice, and no way else.

For example, a quantum model of the double-slit experiment does not attribute probabilities for the distance between the two slits or between the slit and the detection screen. These parameters must be fixed by the background, as a prerequisite for applying the model. Once they are fixed, the model allows one to calculate the probability of observing particles at various positions on the screen. Actual measured positions on the screen constitute our evidence that the model is accurate. All the rest can be ignored.

So, in effect, we never use our total evidence (this would be unpractical), but we focus on relevant weak evidence by reflecting on what is predicted by our hypotheses and on the circumstances that are involved in calculating the predictions (the story could be complicated if we want to aggregate the results of various experiments involving the same models or theories; there is no place here to discuss this issue, but in any case, the debate on fine tuning can be expressed in terms of models applied to a single evidential context).

As far as likelihoods are concerned, there is no reason to seek a match between our epistemic situation and the way we separate background from evidence. We can evaluate new models in light of old evidence without any problem. It is true that what counts as background knowledge will be rigid in the inference (it remains true in all counterfactual states of affairs considered, thus potentially affecting the likelihoods by eliminating possibilities), but this does not mean that it must include all epistemic necessities: it is perfectly legitimate for me to consider implicitly that there are alternative possibilities where *I* do not exist (or equivalently that the likelihood that I exist is lower than 1), as long as the theories under consideration attribute some likelihood

$Pr(O_j|T_i \wedge B)$ to this state of affairs, and that I can acknowledge that it is not instantiated. "I am still alive, so the bomb must have been defused" is correct reasoning (see also Juhl, 2007). After all, Kripke (1980) warned us against conflating epistemic a prioricity and necessity. Under an intuitive reading, the kind of necessity involved with background rigidity has to do with the way a situation of interest is identified, and the likelihood probabilities postulated by our theories are assumed to be objective likelihoods associated with the identity or nature of the situation referred to. It might be the case that our own existence must be assumed in order to correctly identify the initial situation of interest in cases where it is identified *relative to us* (this will be examined in the next section), but this does not imply that we must continue to exist in all the possible scenarios considered in the models, because epistemic a prioricity is irrelevant (this is one reason why it makes sense to view Bayesian inference as reversing objective likelihoods in order to yield a credence update rather than as reversing prior credences).

So far, we have considered theoretical likelihoods, but not including epistemic necessities in the background might be more problematic when considering prior credences in theories $Pr(T_i|B)$, at least if these are interpreted as subjective degrees of credence. These are also probabilities that we need, after all, but what could be my degree of credence in the multiverse hypotheses prior to knowing that I exist? Does it make sense at all? No, it doesn't. But the whole business of attributing prior subjective degrees of credence to hypotheses is bankrupt anyway. No one has ever provided a non-arbitrary subjective degree of credence in anything, except perhaps (perhaps!) in tautologies and contradictions. This notion of degree of credence is a very elegant fiction, but it has nothing whatsoever to do with science or with cognition. Let us keep realistic.

If we adopt an objective version of Bayesianism instead, then priors are presumably fixed by something like a principle of indifference. Not that this is unproblematic given, for instance, Bertrand's paradox, but maybe our past experience is enough to fix an adequate measure on possibility space on the basis of natural categories or statistical data. What matters, in this context, is that the likelihood of the background information is not affected by which theory is true, otherwise assuming a principle of indifference for the $Pr(T_i|B)$ introduces a bias. We are talking about prior probabilities *given the background*, so they cannot be indifferent to the fact that background conditions are realised, otherwise it is as if one Bayesian updating were missing, so to speak.

A typical case where this matters is when a selection effect is involved. Imagine that I conduct a study in order to know whether smoking causes lung cancer. I recruit patients in a hospital, and as it happens, half of them are smokers. This fact should not be part of the background of my inference, but part of my evidence, because the hypothesis that smoking causes lung cancer raises the likelihood for a smoker to be found in a hospital, and since my study is conducted in a hospital, this introduces a selection bias. If hypotheses are attributed equal priors on the basis of a principle of indifference, this should be done *before* taking into account the fact that half of the patients in my study are smokers, otherwise, one Bayesian updating will be missing, and the connection between smoking and cancer will be underestimated. Only the general fact that all recruited patients will be found in a hospital can be part of the background. On the other hand, if 50 smokers and 50 non-smokers were selected by a randomised procedure, we could consider the fact that half of the people recruited in my study are smokers to be part of our background knowledge, because the randomised procedure is designed precisely to ensure that which objects or people are selected does not depend on which hypothesis is true.

This kind of selection effect does not depend on any deep existential consideration, or on epistemic necessities, for reasons already explained. The only question to ask is this one: does the *objective likelihood* of some piece of information depend on which theory is true? If it does, it is part of your evidence. If not, you can put it in the background. For example, if the likelihood of my existence depends on which theory is true, then my existence is useful evidence: it should not be part of the background.

In sum, from a purely calculational point of view, the division between background and evidence does not really matter. This means that there is no reason to seek any match between this division and our epistemic situation, or to take into account any existential aspect. However, the way the division is made can lead one astray in the way priors are determined, in particular when selection effects are involved. Fortunately, the criteria we end up with in the case of priors is congruent with our previous contention that the background should be something that we need in order to compute likelihoods with the theory, and not something the likelihood of which is to be calculated with the theory. If we stick to this principle, there should not be any problem.

All these considerations concern standard Bayesian reasoning. If we express the debate in terms of Bayes factors instead, that is, if we are only interested in the extent to which a piece of evidence confirms or disconfirm one theory over another one (which, I think, is our best option), then there are no priors. We do not need to attribute definite degrees of credence to all hypotheses (or to assume that we could contemplate the entire space of conceivable hypotheses for that matter). We can focus instead on the ratio between the likelihoods predicted by our two theories for what was actually observed, conditioned on the background. And here, the division between background and evidence directly affects the result, because likelihoods are calculated against the background. But the same logic applies: if the objective likelihood of something depends on which theory is true (in virtue of causal processes postulated by the theories), then this something should *not* be part of the background. If, on the other hand, we need some information in order to apply the theory and derive likelihood probabilities, including information about our selection procedure, then this information should be part of the background.

With this in mind, one could wonder what to do of the various proposals in the debate on fine tuning that differ in what they consider to be background knowledge.

Take first Friederich (2019b)'s proposal, which would supposedly match our epistemic situation, to consider as background that life exists in this universe, and as new information that life requires fine-tuning, and that therefore our universe, which contains life, is fine-tuned. A first problem is that it is not clear in what sense fine-tuning constitutes some kind of evidence rather than a component part of the cosmological theories that we are evaluating. Maybe the evidence under consideration should be the observations that confirm these theories and discard the alternatives. (Friederich (2019b) acknowledges this problem, and argues that his argument is not strictly a Bayesian updating for this reason).

A second problem is that this approach does not directly give us the probabilities that we need. The multiverse hypothesis, by itself, does not yield any well-defined theoretical likelihood for fine-tuning, because the idea that there are many universes is conceptually independent from the idea that life requires fine-tuning. Furthermore, the fact that there is life in our universe is not a prerequisite for applying the multiverse theory: it does not play the role of background circumstances in the same way as the distance between the slits is required in order to make predictions from the double-slit model. This does not mean that it is impossible to calculate a probability of fine-tuning given the multiverse and life, but in order to do so, we will need another iteration of Bayes's formula, and eventually, with this new reversal, we will fall back on the standard approach: all the work is done by the likelihood of life given the multiverse or a single universe (although it is not explicit, this move is involved in Friederich's reasoning when he claims that the probability of our universe having the right constant values is 1 given that there is life: the reason for this is that the likelihood of life if the constants were not the right ones would be 0). Therefore, if the approach differs in its results by choosing a different background, this cannot be because it starts from different theoretical likelihoods: it must be because the prior probabilities it attributes to the theories under comparison are attributed against a different background. But then the difference looks somehow artificial or subjective.

Which leads us to the third problem with this approach: the fact that the background that is selected (that this universe exists and that there is life in it) supposedly matches our epistemic situation is irrelevant, and this background is not obviously independent from the theories under consideration. It rather looks like the kind of thing to which these theories would attribute different likelihoods. But this last point is debatable, since it could depend on how "this universe" is analysed.

Now turn to the proposal, from defenders of TUO, to assume "*this* universe exists" as background knowledge. Whether it is legitimate depends on the answer to this question: do our hypotheses affect the likelihood of this universe existing? Or is it an assumption that is required in order to yield probabilities for various possible observations, in particular the observation that there is life in this universe? If the former, the existence of this universe counts as useful evidence, if the latter, it counts as background knowledge (if none, it is irrelevant).

One motivation given for considering that "this universe exists" is part of our background knowledge is the thought that if it were considered evidence, then fine-tuning for life would play no more role in our inference: the multiverse would essentially make the existence of our universe more likely independently of the fact that it contains life, since the likelihood that it has life given that it exists is the same as with a single universe hypothesis (see K. Draper, Draper & Pust, 2007, p. 296; White, 2000). But this argument is fallacious. First, if really our hypotheses affect the likelihood that our universe exists, then its existence should not be part of our background knowledge, otherwise we introduce a bias. And second, the worry expressed here is unfounded. If more lottery tickets are produced, then the likelihood of at least one of them winning the lottery is boosted. This is true even though creating more tickets only raises the likelihood of existence of these tickets, and not the likelihood for any of them to win, and yet, the fact that the likelihood of winning is very low does play an essential role in the inference (in this case, a selection effect is involved: our observations concern winning tickets, whichever they happen to be, rather than specific tickets). I am not claiming that this lottery case is perfectly analogous to the fine-tuning case, but at least, we can see that the idea that the multiverse hypothesis indirectly raises the likelihood for there being life by directly raising the probability for our universe (and others) to exist is not necessarily problematic.

So, whether "this universe exists" should be part of the background or part of our evidence remains to be examined, and again, this will depend on how "this universe" is analysed. Which leads us to the third component of Bayesianism: *evidence*, and in particular the use of indexical or demonstrative terms in our total evidence, to which we now turn.

4 | EVIDENCE AND INDEXICALITY

Arguably, most scientific models are indexical. By this, I do not mean anything particularly deep. An indexical sentence, such as "I am hungry", can be uttered in various contexts to express any state of affairs of the appropriate kind (a state of hunger) from the perspective of any speaker. Similarly, most scientific models do not refer absolutely to one particular object in the world, but they can be used in various contexts to represent any object of the appropriate kind from any perspective. Take the quantum model of the hydrogen atom for instance: it does not represent one specific hydrogen atom in the universe, but can be used to represent any hydrogen atom (or collection thereof) in any context where there is one. The content of such a model can be analysed as a *function from context to content*, following how indexicality is analysed by philosophers of language: to any context where there is an object of the right kind, it assigns specific predictions (see Ruyant, 2021).

What happens, roughly speaking, when a model is tested is that measurements are made on an object or on a sample of objects of the relevant kind in particular circumstances of interest. This is the context. Then the outcomes of these measurements are compared to what the model predicts for objects of this kind in these circumstances. For example, scientists observe the emission spectrum of a bunch of hydrogen atoms. If a model (of the hydrogen atom in this example) gives a very low probability for what is observed, and if we are confident that the experiment was conducted correctly, then the model can be discarded. Otherwise it is confirmed.

In practice, the experimental context can be represented by a model of the experiment that specifies what objects or characteristics are observed exactly. The theoretical model is then embedded in this more encompassing model, so as to yield specific contextual predictions. In cases where a selection bias cannot be avoided (as in the hospital illustration from the previous section), the selection procedure is incorporated in this model of the experiment. Then this bias in the procedure becomes part of the background of our inference, and affects the likelihood of our observations (including, in our illustration, the percentage of recruited smokers). This is how the bias is corrected. Yet the model still works as a function from context to prediction.

It could be thought that cosmological models are different, because they do not really represent a *kind* of object. They represent a single entity: the cosmos. However, they do not function very differently. Cosmological models do not come equipped with a "you are here" sign. They do not represent the cosmos from one specific context or perspective. On the contrary, they are generally homogeneous, postulating that there is no privileged place in the universe, that the same physical laws hold anywhere and at anytime (something called the cosmological principle), and even when a singularity such as the Big Bang is postulated, there is no sense in asking "where did the Big Bang occur?". In consequence, if we want to assess the model, then we have to locate our own context within the model. If the model is homogeneous, it does not matter where we locate ourselves. If the model is not, or not entirely (for example, if the universe expands with time), then it will affect the predictions, so it is important that we self-locate correctly. In any case, a cosmological model is still indexical in the sense described above: it is a function from context to predictions for the context.

This feature is important in the context of fine-tuning if we want to resist the idea that any event whatsoever would confirm the multiverse hypothesis. For example, imagine I just rolled four dice and got "5, 2, 4, 2". The probability for this to happen is only $(\frac{1}{6})^4$, as for any sequence. However, one could reason that if there are many universes with copies of me rolling dice, the probability that this happens in at least one universe is close to 1. So, the fact that this random sequence, or any other, occurred in front of my eyes would confirm that there are many universes.

As White rightly complains (see Isaacs, Hawthorne, and Sanford Russell (2022) for an illuminating discussion), this rationale, if sound, would make fine-tuning for life completely irrelevant to assess the multiverse hypothesis: life only adds one tiny improbability on top of the many ones we witness every time a die is rolled, every time a book is written, or for any contingency whatsoever. Any contingency ultimately favours the multiverse hypothesis, even without life. This looks like a reductio ad absurdum of FTAM. But this reasoning is fallacious if we accept this idea that scientific models, including cosmological ones, are indexical. They do not make unrestricted existential predictions of the form "there is at least one die roll with this sequence in the cosmos" (note, in this respect, that only the toy multiverse models presented in philosophical discussions actually specify how many universes there are exactly). If they make existential predictions, it is conditioned on a context, plausibly in this case the fact that four dice were rolled. And the multiverse hypothesis does not make it more likely for any such context to result in a specific outcome than the single universe hypothesis does.

This is the kind of reasoning that TUO puts forth in order to counter FTAM, and in this sense, our analysis bolsters the objection. But other aspects of TUO are undermined by this picture of model confirmation, and notably the way it analyses "this universe".

Paradoxically, the indexical nature of scientific models means that contextual (or centred) information, expressed by means of indexicals or demonstratives, is, in some sense, irrelevant to assess their adequacy. The fact that it is *us* who are observing, or that it is *this* object that we observe, does not matter. This is because scientific models are designed to be adequate *whatever the context*, at least within a given range. The point of scientific knowledge, or of understanding or explanations in general, is not to collect disparate data, but to make them fit in a general framework. At some point, experimenters must refer to particular objects *from their perspective* in order to confront a model with experimental data, but they will assume that these objects are *representative of their kind*: perhaps they were picked by a random procedure, or perhaps an unavoidable bias will be modelled in the model of the experiment, but at any rate, the fact that these objects in particular are related to these experimenters in particular at a specific place and time should not matter. This is why scientists need not write down the precise time and location of their experiments or name the objects involved in their report: by postulate, this contextual information is not supposed to affect their observations, and since this postulate is built into the theories or models under test, these models are completely indifferent to such contextual information.

There are good reasons for this. Science is an empirical activity, but it is doubtful that we have any empirical access to metaphysical essences or to haecceities. We refer to and keep track of objects on the basis of their manifestations; pointing at them from our perspective and uttering "*this*" does not give us any kind of metaphysical super-power. Whether an essential indexical (Lewis, 1979; Perry, 1979) associated with our subjective point of view *here and now* is ultimately needed on top of our non-centred knowledge in order to account for how knowledge informs our actions is a contentious issue (see Cappelen and Dever (2013) for a book-length argument against this idea). I personally think that it is helpful to think of experimental contexts, from which observations are derived, in these terms. But in any case, the point remains that centred information should not matter for evaluating models and theories in science, because science is a collective enterprise that seeks general knowledge.

This means that even if we accept the requirement of total evidence mentioned previously, indexicals, demonstratives, proper names or any other referential device have no place in our total evidence. They can (and should!) be replaced with no loss by purely qualitative information. To illustrate this thesis, we can modify the scenario reproduced in a previous section:

> Suppose I'm wondering why I feel sick today, and someone suggests that perhaps Adam got drunk last night. I object that I have no reason to believe this hypothesis since Adam's drunkenness would not raise the probability of me feeling sick. I should consider my total evidence: not only the fact that someone is feeling sick in the room,

> but also the fact that it is *me*. But, the reply goes, how is it any kind of evidence? Maybe you and Adam are the same person.

I believe that this last reply is perfectly sound, if not in practice (because we know how to differentiate people), at least in principle. What matters for indexicals, demonstratives or proper names to have any empirical import is to explain how we identify objects, keep track of their identities, and distinguish them from other objects. And in this respect, saying "*this* atom that *I* was pointing at *yesterday* turned out to behave in such a way" bears no evidential weight on top of saying "*one* atom that happens to be connected at some time in such a way to an experimenter turns out to behave in such a way". This does not contradict the requirement of total evidence: it is just that centred information is irrelevant in order to assess non-centred hypotheses, because non-centred hypotheses have no centred consequences.

This point has non-trivial implications with regards to rigidity and background specification, and it is very important to take these implications into account if we want to avoid all selection biases. But before explaining why, it can be helpful to present in more detail how the division between background and evidence that was examined in the previous section is related to these considerations.

Remember that a model is a function from context to predictions. The input of this function corresponds to the background of our Bayesian inference, and the output must correspond to our evidence if the model is adequate. As just explained, our complete knowledge can be described in purely qualitative terms, say $(\exists x, y, z)Q(x, y, z)$, since the fact that it is *us* making the observation does not really matter. However, we might need referential devices such as indexicals or proper names, on top of this qualitative description, in order to make our division between background and evidence. The reason is that we need to keep track of the objects involved in our context: objects mentioned in the background must be recognisable in evidential claims. So, for example, our total qualitative knowledge $(\exists x, y, z)B(x, y, z) \wedge E(x, y, z)$ would become $B(a, b, c)$ and $E(a, b, c)$, with a, b, c proper names or indexicals. These terms will be mapped with component parts of our model (we will "locate" our context in the model) on the basis of their associated qualitative features B. This can be done by constructing a model of the experiment that embeds the theoretical model and describes a, b, c having property B in connection with it (some uncertainty about the connection can be involved here). From this, predictions concerning our objects can be derived and compared with E. In this case, we have indeed some kind of rigidity associated with proper names. We assume that specific objects characterised by $B(a, b, c)$ are present in all counterfactual states of affairs. But this rigidity derives entirely from our division between background and evidence (which should follow the recommendations from the previous section).

Note that existential claims for particular objects will not necessarily fall in the background. This depends on the way the experiment is conducted: if particular objects are first selected randomly, and then observed, then all existential claims will be part of the background, but if whether some objects are observed is part of the predictions of our theories, as in the case of a selection effect, then some existential claims can be part of the evidence. Take again our example where patients are recruited in a hospital for a study. In this case, the fact that all recruited patients are in the hospital is background knowledge ($(\forall x)Rx \rightarrow Hx$). However, which patients exactly are recruited in the study is not: it should be part of our evidence, because the propensity for a patient to be found in the hospital, and therefore recruited, depends on whether the patient smokes, and on the likelihood of lung cancer for smokers, which is among the predictions of our hypotheses.

We could even consider the possibility that no patient is recruited at all because the hospital is empty. This means that our observations will not be reported in terms of predicates applied to proper names that were introduced in the background. They will take the form of new existential claims ($(\exists x)Rx \wedge Ox$) for patients that are eventually recruited, if any. However, in general, these existential claims will still be "anchored" to an object of the background context associated with the experimenters (here, the hospital where observations are made), because models are indexical.

This is exactly where taking demonstratives or proper names too seriously, and assuming that they systematically generate rigidity by means of necessity of origin or similar devices (which is what White does with "this universe"), can lead one astray.

Imagine that during my study, I am wondering what percentage of the population has lung cancer. I go to the hospital, to the service where lung cancer patients are, and meet Adam. I could reason as follows: I know that Adam exists, and his existence does not depend on my hypothesis, so I can consider it background knowledge. I also know, as background, that people with lung cancer are likely to go to the hospital, and that all people that I will meet at the hospital today have lung cancer; assuming this, the fact that I just met Adam confirms the hypothesis that most people in the population have lung cancer, because the more people with this condition, the more likely it is for Adam in particular to be among them, and then the more likely it is for him to go to the hospital and to meet me. If I meet another person, Bea say, the hypothesis is boosted again for the same reasons.

Of course, this reasoning is fallacious. Observing people at the hospital, in the service where patients with lung cancer are, cannot inform me about which percentage of the population has lung cancer, because the sample is biased (all people there have lung cancer). This shows that modelling the selection procedure, in this case that every one that I will meet today has lung cancer, is not sufficient to avoid a bias. It is also crucial that we analyse any referential device correctly. In the case at hand, I should analyse Adam as referring to the first person I met today, *not* as a direct reference to a particular person, and so, "Adam exists" should be analysed as "I met at least one person today, named Adam". And this cannot be part of the background of my inference if it depends on my hypotheses: I cannot refer rigidly to Adam in the context of my reasoning (things would be different if I had met Adam before the study: the referential link should still be analysed in evidential terms, but since it would not depend on the hypotheses, I would be able to refer to Adam rigidly).

The same goes, mutadis mutandis, for the universe in the case of FTAM. "This universe" cannot be simply understood as direct reference to a particular universe without further analysis (pace P. Draper (2020)). A qualitative description of our referential link to the universe should be explicated as part of our evidence, for example "the universe we happen to find ourselves in" (the indexical "we" could be further analysed). Our selection procedure should also be modelled somehow, and be part of our background knowledge ("the universe we find ourselves in, if there is one, contains life"), in the same way that the fact that recruited patients were found in an hospital was part of the background in our previous illustration. But then "This universe exists" should be replaced by "We find ourselves in a universe" (just as for Adam in the previous case), and we cannot consider this fact to be part of the background of our inference, otherwise the fact that there is life in our universe would already be presupposed in the background, even though the likelihood of this could depend on our theories. It must be part of our evidence.

In light of this, the reasoning adopted by TUO is fallacious. It would make sense if our evidence was "We selected a universe randomly, and against all expectations, it turned out to be life-permitting". Then "this universe exists" could be part of the background, because our referential link would not depend on the fact that it contains life. But this, of course, has nothing to do with our actual evidential situation (in this I completely agree with Epstein (2017), even if my diagnosis does not hinge on a criticism of the requirement of total evidence).

Note in passing that issues of rigidity cannot always be inferred directly from grammatical form, as P. Draper (2020) seems to think. Definite descriptions can be used rigidly or not depending on the context, for example: “The tallest man on earth could have been another person (if another very tall person had not died too soon)” versus “The tallest man on earth could have failed to be the tallest man on earth (if he had eaten less in his childhood)”. The same goes for demonstratives: although less commonly, they can be non-rigid as well, as in “This car that I am driving could have been a different one (if I have had more money)”. The point made above could have been expressed as the idea that the semantics of “this” in “this universe that we find ourselves in” should be read exactly as in this last example. However, it is important to keep in mind that the core of the issue is not semantic, but epistemic. Nothing prevents us from rigidifying anything we want in principle, but what primarily matters is avoiding reasoning mistakes when assessing hypotheses. This includes the kind of mistakes that are induced by the grammar of natural languages.

Also note that the thought experiments devised by White (2000) in order to rebut this idea that a selection effect is relevant to fine-tuning are not actual counterexamples to our analysis. These thought experiments (being tied to a specific table where two dice are rolled, and awakened only if a double six comes out) are precisely designed to remove any effect of the theories under consideration (how many such tables there are) on our referential context (the particular table to which we are tied). White is right that the fact that we could not observe any other result than the one we observe is in principle irrelevant. The point of our analysis is *not* that we should put all epistemic necessities, such as “I observe something”, in the background. The point is that we should analyse properly how the reference of indexicals and demonstratives is fixed, and a probabilistic dependence between our hypotheses and the fact that we are able to refer to a universe at all *is* relevant to our inferences, because then this referential relation should not be rigidified into the background. This is where a selection effect is involved.

As for the idea that all observers would be metaphysically tied to the universe where they are found (because of necessity of origin for instance), so that we could not find ourselves in any other universe than this one, even if true, this assumption is not enough to put back “this universe exists” in the background or to change the way demonstratives are analysed. But in any case, this assumption is not a characterization of our evidence, but a metaphysical posit. At best, it could be built into our models so as to yield specific self location rules. I am not denying that this could constitute a way of defending the intuitions behind TUO, and undermine the multiverse explanation for fine-tuning, but it will not be as simple as the original objection pretends. Before to conclude, let us examine these aspects.

5 SELF-LOCATION UNCERTAINTY

So far, I have argued for a particular way of dividing background and evidence in our total knowledge: the background should be the part of our knowledge that is required in order to derive theoretical likelihoods from the hypotheses at stake, while evidence should be the part of our knowledge for which our hypotheses predict likelihoods. I have also argued that the background should typically refer to a *context of application* that will be qualified by the evidence, because scientific hypotheses do not make existential claims. Indexicality or direct reference are only relevant in so far as they allow us to keep track of objects from background to evidence, but our total knowledge could in principle be expressed in purely qualitative terms without loss of evidential support. Finally, I have argued that it is very important to make explicit how reference is fixed if we want to avoid selection biases. We should not leave demonstratives and indexicals unanalysed, but explain how they are connected to the observer. If this connection does not depend on which hypothesis is true, then the existence of the object referred to, together with its connection to the observer, are part of the background. Otherwise they are part of the evidence.

Following all these prescriptions, we reach two conclusions regarding TUO. On the one hand, this objection to FTAM superficially fits the mould just presented, in that it defines a background context (this universe exists) and then considers the likelihood of our observations against this background (that this universe contains life). Acknowledging indexicality in this way avoids having the multiverse hypothesis favoured by any contingency whatsoever, which, I guess, is one important motivation for this move. But on the other hand, the crucial reliance of TUO on centred information interpreted in terms of rigid reference to a specific universe is problematic, and TUO ultimately fails to account for a selection effect for this reason.

At this point, it seems that we find ourselves in an impasse. The problem is that it is impossible to satisfy all our prescriptions in the case of fine-tuning: we cannot refer rigidly to a proper background context for our inferences, so as to respect the indexicality prescription, in a way that is not already biased, because the likelihood of our very existence, hence of any referential connection between us, observers, and any potential context, presumably depends on our hypotheses. This is a problem that affects cosmological topics in general, where deep existential questions are involved (Why is there something rather than nothing? Why do *we* exist? etc.), for scientific models have universal aspirations, but existential questions, as their name suggests, cannot be framed in universal terms. This problem undermines TUO, but also FTAM.

Is there a way out of the impasse? Perhaps the wisest attitude is to remain quietist on the matter. However, if we insist in finding a solution, then we can relax our requirements, our willingness to stick to best scientific practice, to only confirm universal claims by induction, and start to speculate. We are talking metaphysics, and we cannot expect a confirmation of our hypotheses that is as robust as what normal science achieves, but maybe we can still achieve *something*. However, it should not surprise us if the right way of doing Bayesian inference in this metaphysical context becomes itself radically underdetermined by evidence.

This is what the literature on self-location uncertainty is all about: the aim is to provide well-defined probabilities for our own location in a probabilistic model, as if we were to locate ourselves on a map, so as to be able to infer centred propositions from uncentred ones. This is more or less our problem here: we have a multiverse model and a single universe model, and we want to assign probabilities to our possible locations in these models, so as to know whether the multiverse model increases the likelihood of our own existence (with the difference that the models considered in the literature on self-location are not purely generic: they make existential claims). However, there is no consensus on how to do so.

Isaacs, Hawthorne and Sanford Russell (2022) have proved that all standard approaches give the same result, that M_M is still favoured by our evidence, and that it must be the case for all reasonable self-location rules. It makes sense: after all, even if a perfectly rigorous approach remains silent on the likelihood of our very existence (not doppelgangers, but *us*), it remains undeniable that if a model makes our existence impossible, it should be discarded. Making life very unlikely is as close as we can get to making it impossible, and this is what the single universe hypothesis does in this framing of the discussion. So, all else being equal, raising the likelihood of having at least one location corresponding exactly to our situation (prior to anything that happens or not in this situation) should make a model more credible, and invoking the requirement of total evidence to counter this inference is useless, because purely centred information bears no evidential weight, so it cannot trump this result.

Here is a rationale to bolster this idea. Consider a static "map" of a multiverse with ten universes, such that only one of the ten, call it α, contains life. In this case, it seems straightforward that *if* this map is an accurate representation of our cosmos, *then* our universe is α. This is as straightforward as locating oneself on a city map, assuming it is accurate, when only one location corresponds to ours. Furthermore, there does not seem to be any reason to lower our prior credence in the accuracy of this map on the sole basis that α might not be *our* universe (as TUO would have it). This would be like doubting the accuracy of a city map while being in front of a library on the ground that all the buildings on the map that are not libraries might secretly be *our* building (the one in front of us): this does not make any sense; our map is not about the haecceities of buildings (and this is another illustration of why demonstratives must be properly analysed). If now there are two universes with life on the map (or two libraries), we can attribute equal probability weight to the proposition that we are located in one or the other, without changing our prior credence in the accuracy of the map, because questions of accuracy should exclusively be based on objective information, not on centred information.

This approach is called compartmentalised conditionalisation, and I think it makes perfect sense in the case of maps. In a sense, a multiverse model is a collection of such maps with a likelihood probability attributed to each of them. If we accept to apply the same rationale to each of them separately, we arrive at the conclusion that the multiverse is favoured by our own existence to the extent that it allows that there exist at least one universe suitable for life. The likelihood for this is higher than in the single universe case.

Perhaps this conclusion could be resisted. If compartmentalised conditionalisation seems to make sense for epistemic possibilities (and I am quite convinced that it should be adopted by anyone who thinks that the universe is deterministic and that all probabilities are ultimately epistemic), it is less intuitive when natural possibilities such as propensities or causal relations are considered. This is a way in which our metaphysical interpretation of likelihoods could matter in the debate.

As far as I know, standard approaches towards self-location uncertainty never introduce any differential treatment between epistemic and natural possibilities: all possibilities implied by all epistemically possible hypotheses are treated as possible worlds on a par with all the others. Interestingly, a particular differential treatment (self indication for natural possibilities within models, compartmentalised in epistemically possible models) provides a way to resist the conclusion of Isaacs et al, and to reach the same conclusion as TUO, that the multiverse is not favoured by our evidence (and even disfavoured if we do not assume trans-world identity for universes). This can be shown easily using a simple two-universe model versus a single universe one. The idea, which fits well with necessity of origin, is that *if* a given model is correct (comparmentalised conditionalisation), *then* we (or better, our origins) are located at exactly one place in its possibility tree, without any counterparts in other branches (self-indication).

So yes, strictly speaking, all options remain on the table. But defenders of this solution cannot make as if they had a knockdown argument against the multiverse on the basis of muddled indexical semantics. They defend a particular metaphysical view on self-location that no one should feel obliged to endorse without convincing arguments.

ACKNOWLEDGMENTS

I would like to thank Mark O'Brien for his very helpful comments. This work has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 101022338.

ORCID

Quentin Ruyant https://orcid.org/0000-0002-4040-9711

REFERENCES

Barnes, L. (2013). Bayes' Theorem: What Is This "Background Information"? *Letters to Nature* (blog). https://letterstonature.wordpress.com/2013/11/17/bayes-theorem-what-is-this-background-information/

Barrett, M. & Sober, E. (2020). The Requirement of Total Evidence: A Reply to Epstein's Critique. *Philosophy of Science*, *87*, 191–203. https://doi.org/10.1086/706086

Bradley, D. (2009). Multiple Universes and Observation Selection Effects. *American Philosophical Quarterly*, *46*, 61–72. https://www.jstor.org/stable/20464437

Cappelen, H & Dever, J. (2013). *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. Oxford University Press. 10.1093/acprof:oso/9780199686742.001.0001

Colyvan, M., Garfield, J. & Priest, G. (2005). Problems With the Argument From Fine Tuning. *Synthese*, *145*, 325–338. 10.1007/s11229-005-6195-0

Draper, K., Draper, P. & Pust, J. (2007). Probabilistic Arguments for Multiple Universes. *Pacific Philosophical Quarterly*, *88*, 288–307. 10.1111/j.1468-0114.2007.00293.x

Draper, P. (2020). In Defense of the Requirement of Total Evidence. *Philosophy of Science*, *87*, 179–190. 10.1086/706084

Epstein, P. (2017). The Fine-Tuning Argument and the Requirement of Total Evidence. *Philosophy of Science*, *84*, 639–658. 10.1086/693465

Friederich, S. (2019a). Reconsidering the Inverse Gambler's Fallacy Charge Against the Fine-Tuning Argument for the Multiverse. *Journal for General Philosophy of Science*, *50*, 29–41. 10.1007/s10838-018-9422-3

Friederich, S. (2019b). A New Fine-Tuning Argument for the Multiverse. *Foundations of Physics*, *49*, 1011–1021. https://doi.org/10.1007/s10701-019-00246-2

Goff, P. (2019). Did the Universe Design Itself? *International Journal for Philosophy of Religion*, *85*, 99–122. 10.1007/s11153-018-9692-z

Hacking, I. (1987). The Inverse Gambler's Fallacy: The Argument from Design. The Anthropic Principle Applied to Wheeler Universes. *Mind*, *XCVI*, 331–340. 10.1093/mind/XCVI.383.331

Holder, R. (2002). Fine-Tuning, Multiple Universes and Theism. *Noûs*, *36*, 295–312. 10.1111/1468-0068.00372

Isaacs, Y., Hawthorne, J. & Sanford Russell, J. (2022). Multiple Universes and Self-Locating Evidence. *The Philosophical Review*, *131*, 241–294. 10.1215/00318108-9743809

Juhl, C. (2007). Fine-Tuning and Old Evidence. *Noûs*, *41*, 550–558. 10.1111/j.1468-0068.2007.00661.x

Kripke, S. (1980). *Naming and Necessity*. Harvard University Press.

Lewis, D. (1979). Attitudes De Dicto and De Se. *The Philosophical Review*, *88*, 513–543. 10.2307/2184843

Mayo, D. & Spanos, A. (2006). Severe Testing as a Basic Concept in a Neyman–Pearson Philosophy of Induction. *The British Journal for the Philosophy of Science*, *57*, 323–357. 10.1093/bjps/axl003

Meacham, C. (2016). Ur-Priors, Conditionalization, and Ur-Prior Conditionalization. *Ergo, an Open Access Journal of Philosophy*, 3. 10.3998/ergo.12405314.0003.017

Metcalf, T. (2018). Fine-Tuning the Multiverse. *Faith and Philosophy*, *35*, 3–32. 10.5840/faithphil20181992

Monton, B. (2006). God, Fine-Tuning, and the Problem of Old Evidence. *The British Journal for the Philosophy of Science*, *57*, 405–24. 10.1093/bjps/axl008

Perry, J. (1979). 'The Problem of the Essential Indexical'. *Noûs*, *13*, 3–21. 10.2307/2214792

Ruyant, Q. (2021). True Griceanism: Filling the Gaps in Callender and Cohen's Account of Scientific Representation. *Philosophy of Science*, *88*, 533–553. 10.1086/712882

Sagart, L., Jacques, G., Lai, Y., Ryder, R., Thouzeau, V., Greenhill, S. & List, J. (2019). Dated Language Phylogenies Shed Light on the Ancestry of Sino-Tibetan. *Proceedings of the National Academy of Sciences*, *116*, 10317–10322. 10.1073/pnas.1817972116

Suppe, F. (1989). *The Semantic Conception of Theories and Scientific Realism*. University of Illinois Press.
Wenmackers, S. (2023). Uniform Probability in Cosmology. *Studies in History and Philosophy of Science*, *101*, 48–60. 10.1016/j.shpsa.2023.08.009
White, R. (2000). Fine-Tuning and Multiple Universes. *Noûs*, *34*, 260–276. 10.1111/0029-4624.00210
Winther, R. (2020). The Structure of Scientific Theories. In E. Zalta (ed.), *Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.

How to cite this article: Ruyant, Q. (2024). Indexicality, Bayesian background and self-location in fine-tuning arguments for the multiverse. *Noûs*, 1–20. https://doi.org/10.1111/nous.12502