

Knowing in the Teeth of the Diallelus

How rightly not to be sceptical

J. A. Durieux¹

1 Introduction

The Münchhausen trilemma and the diallelus present a formidable obstacle to the possibility to know. In chapter 2 we look into the question whether infinite or circular reasoning can provide information. Chapter 3 is about the kind of information we might have access to, and theories of truth. Chapter 4 investigates the cases of transcendent exceptionalism and transcendent mediocrity, and concludes that only in the latter case knowledge of the external world is possible.

Chapter 2 contains a bit of probability calculus. Those who don't need to be reminded of the poignancy of the sceptical argument can skip it.

2 Grounding knowledge

The Münchhausen trilemma states that any piece of knowledge requires support of at least one of the following kinds:

1. **Foundational** – support that itself requires no further support (or support that is equal to the knowledge itself, depending on one's position on the nature of support).
2. **Infinitary** – an endless and non-repeating chain (or directed acyclic graph) of grounds, each providing support for the knowledge above it².
3. **Circular** – a graph that contains cycles, such that the transitive closure of the grounding relation in some cases reflexive.

In practice it is useful to make a more fine-grained distinction; we shall do so in separate subsections.

2.1 Immediate foundational support

In the simplest case, the fact under consideration is already part of our starting knowledge. In chapter 3 below we shall look into the question what that starting knowledge might consist of.

2.2 Finite foundational support

In the case of foundational support, there is a body of information already known to us, and finite reasoning allows us to close that body transitively with respect to implication. Of course the amount of information we already know could be infinite, or the information itself could be *about* infinite structures – such as knowledge how to compute limit values. The reasoning graph itself will be finitely deep, however³, and end in information that is certain (though, again, it may be *about* uncertainty, such as the information that a fair coin has a probability of $\frac{1}{2}$ of landing heads).

2.3 Potentially infinitary support

Infinite chains where an infinite number of nodes *also* have foundational support may be run down finitely deep to collect the information in a finite subset of the nodes. This can be compared to the following situation.

Imagine an infinite row of painters on an infinite plain, working on a huge, intricate painting of a

1 E-mail: truth@b.biep.org; orcid: [0000-0003-2582-4973](https://orcid.org/0000-0003-2582-4973); web site: <https://biep.org>.

2 It may be useful to stress that this is about doxastics, not the mere fact that *the number of propositions* may be infinite. Infinite systems of propositions that contain a finite amount of information may be solvable with finite work. For instance, the infinite system $\{a_i = 2^{-i} + a_{i+1}\}$ can be solved for any $i \in \mathbb{N}$ in finite time.

3 Given that we are finite beings, it will be finitely broad as well. Here we are only concerned with the depth – the length of the reasoning chain –, however.

perspective – a river running to the horizon. Each painter paints one kilometre of that river, and then passes on the painting to the next painter in the row. Even though I shall never be able to see the whole painting, walking arbitrarily far along the row of painters enables me to see any foreground of the painting apart from the horizon itself. In practice, there will be a hard limit, determined by my finiteness. *Achievable* information is information I can actually reach, given the circumstances I am in.

In such situations each step adds part of the picture, and depending on the nature of the sequence, in the limit the whole picture may or may not be reached. It is this limit that can be approached asymptotically.

Of the special case where this foundational support “along the way” is about probabilities, Atkinson and Peijnenburg (2017) have made a thorough analysis. In that case there is a statement that may be true or false, and at each level part of the probability spectrum is “painted in”, i.e. each node i is of the form: “of the remaining part of the probability spectrum, for some fraction τ_i the statement is true; for some fraction ϕ_i the statement is false, and for the remaining $\rho_i \triangleq \tau_i - \phi_i$ the next node in the chain will provide the information”. So at each level, an amount of probability is added, and in the limit the exact probability of the statement’s being true is known, provided $\lim_{n \rightarrow \infty} \prod_{i=0}^n \rho_i = 0$ (what they call “the usual class”)⁴.

The image there would be each successive painter i painting a horizontal strip τ_i of plains and the corresponding strip ϕ_i of sky, thus more and more constraining the possible position of the horizon on the canvas.

Now Atkinson and Peijnenburg do not make a terminological distinction between three cases:

1. The foundational support found on each level is known in advance. In this case instead of going down the chain, the problem can be recast in a finite form – computing the limit of a known series. This is their original *Barbara* example.
2. The foundational support cannot be known in advance, but is met along the way. That is the true potentially infinitary support that concerns this section.
3. *Each* bit of support needs support again – there is no foundational information on any node. They exclude that situation in some remarks⁵, though – it falls outside the scope of their analysis. We shall call that *actually infinitary support*, and it is the subject of the next section.

Atkinson and Peijnenburg give an example, where the probability a bacterium (*Barbara*) has some trait depends on both a known base probability of having mutated (the foundational knowledge), and whether her ancestor had it (the regredient knowledge). So each node in the main chain has *two* ancestors, one part of the chain (“the parent has this probability of having the trait”), and the other a leaf node with the mutation information (“*this bacterium has probability ϕ of losing the trait if it inherited it, and probability τ of gaining the trait if it didn’t inherit it, and probability ρ of inheriting whether it has it*”), and correspondingly it has both foundational support and inherited support – but only the foundational support does have any influence.

Now because in their example we are supposed already to *know* the probabilities for each bacterium, and the total of that knowledge is finite, the situation is actually foundational, and the outcome can be fully calculated in finite time⁶. If we replace the pre-known information about all generations by information that must be collected in some way individually for each generation, all that remains possible is the “walk down the row of painters”, and we have the potentially infinitary support with which this section is concerned.

2.4 Virtually infinitary support

We are finite beings, and information that takes too long to reach might as well be infinitely far away, so in practice we may not be able to gather all the information at finite depths. In that case there is a “timber line”, information below which for the sake of the analysis can be considered “at infinity”. With that adaptation, the next section will apply.

4 Because $\prod_{i=0}^n \rho_i$ is the amount of influence node n has on the outcome, and this goes to zero in the limit, they speak of “fading foundations”, i.e. the relevance for the outcome of the foundational knowledge found in deeper parts of the chain goes to zero as the subchain under consideration starts farther and farther away from the root. A slightly more mathematical version of this argument is given in the appendix, section 5 below.

5 They make this clear in e.g. their sections 4.4 and 6.4: “we are not trying to formulate an answer to the sceptic”.

6 This is comparable to the situation where there wasn’t a river, but a straight railroad in the plains. Part of the painting, plus the knowledge that the railroad was straight, would allow one to find by extrapolation the point where the painted rails will meet at the horizon.

2.5 Actually infinitary support

Here no (sufficient amount of) foundational support can be reached at any finite depth. One possibility for that is what Atkinson and Peijnenburg call the “exceptional class” (into which we shall look deeper in section 2.6 below) – but there is a more fundamental case.

Let us consider a variant of the Barbara example. Here we investigate the bacterium *Barbara celarent*, and the research question and situation are as above, except that we have no foundational support of the mutation rates – those rates are supported by research, and the research is supported by general knowledge about how to set up experiments and how to interpret the results, and those in turn are supported by probability theory, which is supported by the law of great numbers, which is supported by experiments with flipping coins, which are ... and so on, literally *ad infinitum*. And of course the claim that we *did* the research – that we didn’t just dream we did, or have a false memory – needs undergirding too, and so do the claims that the bacterium we tested was indeed *B. celarent*, that the experimental set-up was correct, that the lab journals weren’t tampered with between the research and the computations, that the underlying model was sufficiently correct, that the scientific method yields truth, that ...

In the *B. celarent* example there is no way to establish any prior probability, and as before all information from infinity tends to have zero relevance. The conjunctive nodes at any point make this the more poignant – though infinite disjunctions (alternative ways to support the conclusion) might counteract that.

The *reductio argument* states that if one fact can be proven from such a groundless graph, then any statement can – including the negation of that fact. So, infinite chains cannot take us beyond the transitive closure under inference of the knowledge we start with. No knowledge can emerge from infinity.

2.6 The exceptional class and circular support

The same objection holds in the case of circular support. We can “roll out” the cycles, and reduce the cyclic graph to an infinitely-deep non-cyclic one, and apply the above reasoning. As the S values end up repeating, the product will vanish except in the degenerate case where they are 1, the exceptional class.

But what about that exceptional class? This contains cases of (asymptotic approximations of) infinite entailment rather than probabilistic reasoning. These cases are like pronomies – sentences that have more than one fixed point. Take the self-affirmation: “This sentence is true”. It can validly be true, and validly be false, and no inference can wrest out the information which of the two it *really* is. A non-circular version (say, the sequence of “Foo is bar” followed by infinitely many copies of “The previous sentence is true”) is no different in that respect – the sequence as a whole gives us no information about whether foo is or is not bar.

Suppose the trait is the ability to defend itself against a poison in its environment. Through some mechanism (let’s say gen deduplication), offspring tends to be better at defending itself. If the probability to survive an antibiotal treatment is p , then $p_{i-1} = p_i \cdot (1 + \text{rnd}(1 - p_i))$. Now if some bacterium with an infinite pedigree is given us, no amount of reasoning can help us decide whether it will or will not survive the antibiotics – all we know is that its probability to survive it will be either 0 or 1, depending on whether the infinitely-far first ancestor had none or some of the trait.

Now circular support is more restricted than repetitious support, because it enforces identity not only of the formula, but also of its value. An endless list of “The previous sentence is false” is a pronomy, whereas a sentence pointing to itself that way is an antinomy. In this case the endless list is like a two-step cycle. Sometimes, circularity can turn a pronomy into a system with a single solution.

In the probabilistic case, as long as a regression in the exceptional class forces *some* restriction on the possible probability of the starting proposition (i.e. we are not dealing with pure entailment), an endless repetition of the infinite regression would fall in the normal class. Atkinson and Peijnenburg analyse this in their chapter 8.

Having a unique fixed point doesn’t make that outcome correct, however. Loops are cases of self-reference, and we don’t know yet how to deal properly with self-reference. Treating it in the way we treat non-self-referential propositions leads to errors. If that weren’t so, the proposition “*If this statement is true, then God exists*” would be a sound argument for the existence of God. We shall return to this in section 3.1 below.

So in the cases of the exceptional class and circular support as well, there is no information emerging from the infinite or circular depth.

2.7 Pragmatically foundational support

If my wife holds up her hand, stating “Here is one hand”, I agree with her. If George Edward Moore does the same, I’ll include some unfulfilled supports: “Well, provided I can trust my visual qualia, and all they seem to imply, then yes.” In practice, while writing up the *B. celarent* report we see no need to justify *modus ponens*, give evidence for the statistical law of large numbers, or prove the existence of an external world. Belief in those is shared by all our readers (if there is an external world, and it does in fact contain readers, and ...), so we somewhat arbitrarily cut off our support graph there.

Infinetists seem to do this – they claim that support should be given till one reaches “obvious” propositions.

There is nothing wrong with that, if the assumed truths have indeed been established – which means that the world is such that those things *can* indeed be established, or at least be made probable⁷. Chapter 4 below looks into the requirements that places on reality.

3 The foundation

Since we are interested in establishing an upper bound on what can be known, we can grant without defence that we can know our qualia, including our hopes and fears, and our thoughts, including our imaginings, for which we shall use the word *dream* as a technical term. So an author devising a world with a plot has a dream, and so does a mathematician who considers the seven-point space. We can also know that we know these things⁸.

The main hurdle for knowing is the *diallelus*: whatever we know indirectly, we know through some channel. That channel may be a reasoning, or a message bearer (e.g. an optic nerve), or a criterion for truth, or anything transporting information from a (quite possibly compound) source to us as the receiver.

The diallelus: *we can only know X on the basis of what comes to us through channel C from source S, if we already know the reliability of both S and C.*

So if I already know *modus ponens* as truth yielding⁹, and I know both p and $p \rightarrow q$, then I can know q – but can knowledge of *modus ponens* be part of our foundation? Can we *know* that *modus ponens* is true? Obviously we can *believe* it, and even find ourselves unable to disbelieve or seriously doubt it, but does that make it true?

<<Add probabilistic version: even a basis for that is problematic, so not even likely truth is achievable.>>

3.1 Coherence

If I were a solipse¹⁰, my beliefs would be the only reality out there, and therefore as true as anything can get. So the fact that I had a coherent set of beliefs would suffice for knowing.

A related position is the *coherence theory of truth*, that coherence is all that matters for truth.

This position makes sentences such as “There is a world out there independent from my thoughts” have a very different meaning from its apparent one – it no longer is about a world out there, but about coherence of that belief with other beliefs, quite independent from whether there is or is not actually a world out there in the traditional sense¹¹.

A big question is “coherence of what?” Let B be the whole of one’s beliefs.

- Internal coherence – coherence of B – is insufficient, even if there is only a single fixed point, as we saw in section 2.6 above. Otherwise “If this sentence is true, God exists” would be a valid existence proof of God, and “If this sentence is true, God does not exist” an equally valid proof of the opposite, – as long as one believes only one of those, and no other beliefs refute the one believed. On the other hand, *at least* internal coherence will be required for truth.

7 There is an amazing light-heartedness regarding scepticism. For many, a statement such as “Well, if you are going to doubt *that*, you might as well start doubting everything!” seems to be a sufficient ground for rejecting scepticism. Finding the conditions under which knowledge is possible seems a much less-taken road.

8 Arguably, if I believe I have pain, I *have* pain, and if I believe I am imagining a white sphere, I am *actually* imagining what I believe is a white sphere. I do not claim we can know that the thing we imagine is actually a white sphere, of course – external knowledge is not made true by believing it, or false by disbelieving it.

9 As Lewis Carroll (1895) showed, this is very different from knowing the fact “*modus ponens* yields truth”.

10 A *solipse* is a solipsist who is right, or a non-solipsist who is wrong. The latter position yields interesting paradoxes, which fall outside the scope of this article, however.

11 But then *that* sentence would also have another meaning, and the original idea is simply not expressible within the coherentist framework.

- Coherence of $B \cup S$, for some S - but what would that S be? It cannot be some externally-defined set, lest the truth of "That is S " be not determined by coherence but by correspondence. We shall see a satisfactory candidate in section 4.2 below.
- Coherence of a set of beliefs meeting some criterion, such as minimality or simplicity. One question is: "why *that* criterion?". Is it true that that is the correct criterion? Another is that there seems to be no guarantee that any such criterion will pick out a unique system of beliefs, that not several may be equally minimal or simple, for example.

The question of determining B itself is hard too. On pain of inconsistency the coherentist may not refer to one's actual beliefs in giving the truth condition for " B is the set of one's beliefs". Likewise the notion of coherence, and the criterion, stand in need of determination. If I believe that my beliefs are coherent, i.e. I believe "it is true that B is coherent", does that suffice? Does my belief that B is simple suffice?

3.2 Correspondence

The main contender is the *correspondence theory of truth*, which states that truth equals correspondence between belief and reality¹². But insofar our beliefs concern the external world, how can we ever decide their truth or falsity? How could we gain access to this outer world, in order to compare it to our beliefs?

One solution would be to restrict our beliefs to our inner world - our qualia, dreams, longings, beliefs, and so on - and remain agnostic about an external world, but that leaves one with a very limited domain for truth.

4 Bootstrapping

But the situation is worse, for how can we get at that prior in the first place?

So the only option for us to know anything is for it to be grounded by a finite, acyclic graph, the leaf nodes of which are foundational. The leaves must be material that is already available to our minds, and reasoning can then reach a transitive closure of that material under probabilistic reasoning. But what is that material?

The answer depends on the way the world is - is *transcendent exceptionalism*¹³ true, or *transcendent mediocrity*? We shall investigate both cases in separate subsections.

4.1 Transcendent exceptionalism

According to *transcendent exceptionalism* the world is immanent - we may be dreaming other worlds, and people in it, but no-one is dreaming our world, with us in it.

There is me, at least a bundle of knowledge. There are qualia, dreams, propositions, and possibly other elements in that bundle, and it is reasonable to assume that I can have knowledge of those, their nature and shape. So those are possible leaf nodes of inference.

Knowing a proposition does not imply knowing its truth value, of course. I may have strong convictions that an external world exists, but that doesn't make it knowledge. The fact *that* I have those convictions may constitute knowledge, though.

I can have knowledge of my dreams, of the laws that hold in them, and of any bit of them - I am the one dreaming them, after all. Whether those laws hold in the external world, if there is one, is another matter. There is no *a priori* reason why any of my beliefs about an external world are correct - and that includes beliefs concerning logic and probability theory.

For instance, I might try to reason from repeatable experiments ("if I will to close my eyes I consistently get the experience of reddish twilight") to a lawful external world, but that requires induction, and there is no way to get to induction without using it already. From "induction worked in the past" to "induction will probably work again" requires an inductive step. Induction presupposes order, and there is no way to discover order¹⁴.

Likewise logic: there is no way to infer *modus ponens* or *universal instantiation* without applying those - or applying some rule that in the end will require those for its grounding (Finn 2019). The fact that logic

12 This is independent of the precise natures of both the truth bearers and of the truth conditions, about which we shall remain agnostic in this paper.

13 See my [Transcendent Mediocrity is the Neutral Position](#) for an explanation of the notions of transcendent exceptionalism and mediocrity, and of the technical sense in which I use *dream*, *dreaming*, and so on.

14 This is explained in my [The World's Haecceity is the Dual of My Thrownness](#).

reasoning feels coherent does not prove anything either – maybe it is the very error-riddenness of our reasoning that makes us conclude its coherence.

So we seem limited to reasoning about our dream worlds, where we are in control of the facts, including any logical and physical laws that may hold there.

One interesting conclusion is that if transcendent exceptionalism is true, we have no way of knowing it – for all I know I might be a solipse, or part of someone’s dream world. Believing or stating that transcendent exceptionalism holds is a performative contradiction.

4.2 Transcendent mediocrity

Under *transcendent mediocrity*, the situation is very different. Now we have a transcendent mind thinking this world, including us, and including our thoughts. Let us call this mind *God*. Assuming for simplicity that God is a solipse¹⁵, he knows his dream – our world – and can convey to us the transitive closure of this knowledge¹⁶, and the knowledge that what he conveys to us is reliable. After all, as the dreamer of our world, he also dreams the rules holding here.

There are some conditions that come with this, given the nature of our world, but we shall not specify those here¹⁷. One fact about this world is the existence of evil, which includes disbelief in the truth. We can deny facts we know, bringing ourselves to disbelief, or lull ourselves into believing facts we *want* to be true. This is so for facts about the world – say, the gambler convincing herself that she will win next time and so cover her losses – and for the very knowledge of the existence and nature of God¹⁸.

4.2.1 Truth

The truth about our world being precisely what God believes about it, if *S* is the whole of God’s beliefs about our world and *B* is what we believe about it, then the coherence of $B \cup S$ is a correct criterion for the truth of *B*, for the correspondence of *B* with the actual (state of the) world.

What immanently is a correspondence theory of truth, transcendentally is a coherence theory: coherence with the maximal set of beliefs that *do* correspond with the world – with what God believes. This does not meet the requirements of the coherence theory of truth, as it still points to a set of beliefs in a way that theory doesn’t allow us to.

However, a modified coherence theory might allow access to one’s own beliefs – and since those beliefs contain pointers to God and his beliefs, at the very least the criterion can be *stated* coherently. *Evaluating* beliefs for truth is a different matter. From above (i.e. by God) that can be done within the framework, as my beliefs are coded in his beliefs about my beliefs. From below (by us) that can be done precisely to the point that our access through the knowledge *a priori* allows us. What is it that God gives us in that respect?

4.2.2 Knowledge *a priori*

This knowledge is immediate – my thoughts *are* thoughts of the transcendent mind, and are in a way akin to qualia. I can *know*, but I cannot *prove* that I have pain, or long for a cold drink – not to others, and not even to myself. Likewise, I can *know* that God exists, that *modus ponens* is true, that there is an external world, but I cannot *prove* it. The very feature that makes such knowledge immune to the *diallelus* also prevents it from being grounded by reason. If I stub my toe, and have pain, all I can hope for is *resonance* in others – their ability to understand my pain by being sufficiently like me and being willing to imagine undergoing the same, and to consider the qualia that would go with that. If someone were flatly to deny understanding what I was talking about, or being aware of a quale that would fit my attempts to describe my pain, I have no means to change that – other than pointing to others who *do* admit understanding.

15 If it is not, we shall need a second-order transcendent mind to dream the first-order transcendent mind, and so on, for the knowledge condition in this section to be fulfilled. Let the mind in question be the solipse at the root of the tree described in [Transcendent Mediocrity is the Neutral Position](#).

16 This is one way of putting it, but there is no reason why this would involve an infinite regress. For the purpose, he might decide there be such a thing as “self-evident truth”, with all the relevant properties and attributes. Such are the joys of being a dreamer (and a solipse at that).

17 See my [Fundamentals of Philosophy](#) for a very basic statement of them.

18 This corresponds to the claim made by many theologians and philosophers of religion, that we all know that God exists, but repress that knowledge. And that is not only true for atheists, or people believing in false religions, but for all of us. At best, some of us allow a modicum of that truth to enlighten their thinking.

Likewise, I must now switch to a description that I merely hope will resonate with some readers, and may bring some of the others to a searching introspection. Maybe the fact that the alternative – having no foundation for any beliefs beyond self-beliefs – is so unappealing may bring readers to it, and hopefully some will discover deep down, at the very foundation of their thinking and believing, the things I am about to describe.

Let me then start by stating that what I find there is not propositional knowledge. The English language is unfortunately poor here, but there is body, soul, and mind knowledge (German *können*, *kennen*, *wissen*; French *pouvoir*¹⁹, *connaître*, *savoir*). If we appropriate the triad *can*, *ken*, *wit* for that distinction, then what I find is not some “witting”, but something encompassing “canning” en “kenning”.

I may be able to speak my mother tongue grammatically faultlessly, while finding myself unable to formulate correct grammar rules for it. Yet, if I hear a sentence formed by my or others’ rules that is ungrammatical, I will know so. Likewise, I find an ability to reason, which also gives me a familiarity with reason – and which helps me to recognise the rules of logic as basically correct.

Analogously I find such a can/ken foundation for understanding my sensory qualia, and from these derives my witting that an external world exists. All our applications of those foundations are fallible, but in this case the foundation itself let me also ken the fallibility of sensory perception.

The can/ken foundation for the existence of other minds is even more extreme in this respect: it lets me ken an extreme potential fallibility – and tends to compensate for that by tentatively attributing mind to anything to which it might even vaguely apply.

I shall not try to be exhaustive. Time, order, good and evil, and so on all have their can/ken foundation. All those foundations also let themselves be known as reliable (within their own limits), and from God. God himself also has his can/ken foundation – I may be wrong in many ways *about* God, i.e. my witting about him may be very wrong, but I ken him and can interact with him. And in a way that foundation is the foundation of those foundations too, but again I find that I ken it as such, but find it hard to put that in wit terms – in communicable, propositional language²⁰.

4.2.3 Fallibility

Our knowledge is fallible, but this does not translate to fallibility of the can/ken foundations. Only if they were to present something as *infallible* whereas in fact it were wrong would that be the case. But on the contrary, the foundations make me aware of my fallibility at all times.

If e.g. my reasoning foundation would present my reasoning tools as infallible, I would never come to know that they weren’t – any contradiction derived would merely drive me to beliefs that everything is true and false at the same time. It is only because I already have the knowledge that my reasoning is *potentially* wrong that I can recognise when it is *actually* wrong.

It is precisely because the can/ken foundations present much of their knowledge as fallible that we can deny unwanted elements of it. Taking ethical knowledge as an example, we can harden ourselves against the “voice” of our foundation for morality, and convince ourselves that certain behaviour is not immoral – and end up truly believing that, on a higher level. We can “dehumanise” someone, or some group, i.e. deny a human being. Dehumanising perceived enemies, or torture victims, are well-known examples. Likewise we can deny a soul in animals, and torture them – fishing for sport, cooking lobsters alive, and lots of ancient folk games, many of which have been banned nowadays.

4.2.4 Denial

One can also deny the existence of, or access too, the can/ken foundations. Possibly some people actually don’t have access to some of them – psychopaths to the ethics foundation, for instance, or atheists to the God foundation. This presents no epistemic problem; it simply means that such people have less knowledge they can use as a foundation.

It does become an epistemic problem if someone on the one hand claims not to have some foundation, but on the other hand uses knowledge from it as well-founded. Someone who trusts logic while denying the can/ken foundation for reasoning is making an epistemic error. Others – who do recognise the foundation – may have reason to trust his reasoning, but he himself hasn’t²¹.

19 Unfortunately, *pouvoir* is otherwise overloaded, and one says « Je sais faire cela » for skills on any of the three levels.

20 This may be one reason among many others why people tend to share their experience with God partly through song, dance, ritual, and other means of communication that are closer to canning and kenning. Reducing knowledge about God to witting would amount to *canning God* it in the common sense of “canning”.

21 This seems to be where presuppositional apologetics starts.

5 Appendix

It has become traditional to shirk the responsibility to define “justification”. We shall follow that tradition, and accept the approach of Atkinson and Peijnenburg (2017) of only requiring that “ A_i is grounded in A_j ” implies that $P(A_i|A_j) > P(A_i|\neg A_j)$. This implies no loss of generality, since if $P(A_i|A_j) = P(A_i|\neg A_j)$ there is no support from A_j for A_i , and in the case that $P(A_i|A_j) < P(A_i|\neg A_j)$ replacing A_j by $\neg A_j$ will produce the desired inequality. So we assume that $S_i \triangleq P(A_i|A_{i+1}) - P(A_i|\neg A_{i+1})$ is positive.

In order to see why true infinitary support is impossible, it is useful to distinguish the prior probability P and the posterior probability P^+ . P^- includes all achievable foundational information, that is the maximum information the reasoner is able to collect by going finitely far down the chain. Then we can rewrite the Rule of Total Probability so as to separate out the prior probability and the influence $\Delta \triangleq P^+ - P^-$ of the support. So $P^+ = P^- + \Delta$, and support is possible if $\Delta \neq 0$, i.e. if adding the support makes any difference.

Let us assume a linear potentially infinitary support of the kind $A_0 \leftarrow A_1 \leftarrow A_2 \leftarrow A_3 \leftarrow A_4 \leftarrow A_5 \leftarrow \dots$, where the arrow indicates “grounds”.

The Rule of Total Probability then states that

$$\begin{aligned} P^+(A_i) &= P(A_i|A_{i+1}) \cdot P^+(A_{i+1}) + P(A_i|\neg A_{i+1}) \cdot P^+(\neg A_{i+1}) \\ P^-(A_i) &= P(A_i|A_{i+1}) \cdot P^-(A_{i+1}) + P(A_i|\neg A_{i+1}) \cdot P^-(\neg A_{i+1}) \end{aligned}$$

Now since

$$\begin{aligned} P^+(A_{i+1}) &= P^-(A_{i+1}) + \Delta_{i+1} \\ P^+(\neg A_{i+1}) &= P^-(\neg A_{i+1}) - \Delta_{i+1} \end{aligned}$$

we get

$$\begin{aligned} \Delta_i &= P^+(A_i) - P^-(A_i) \\ &= P(A_i|A_{i+1}) \cdot P^+(A_{i+1}) + P(A_i|\neg A_{i+1}) \cdot P^+(\neg A_{i+1}) - P^-(A_i) \\ &= P(A_i|A_{i+1}) \cdot (P^-(A_{i+1}) + \Delta_{i+1}) + P(A_i|\neg A_{i+1}) \cdot (P^-(\neg A_{i+1}) - \Delta_{i+1}) - P^-(A_i) \\ &= P(A_i|A_{i+1}) \cdot P^-(A_{i+1}) + P(A_i|A_{i+1}) \cdot \Delta_{i+1} + P(A_i|\neg A_{i+1}) \cdot P^-(\neg A_{i+1}) + P(A_i|\neg A_{i+1}) \cdot (-\Delta_{i+1}) - P^-(A_i) \\ &= P(A_i|A_{i+1}) \cdot \Delta_{i+1} - P(A_i|\neg A_{i+1}) \cdot \Delta_{i+1} + P(A_i|A_{i+1}) \cdot P^-(A_{i+1}) + P(A_i|\neg A_{i+1}) \cdot P^-(\neg A_{i+1}) - P^-(A_i) \\ &= S_i \cdot \Delta_{i+1} + P(A_i|A_{i+1}) \cdot P^-(A_{i+1}) + P(A_i|\neg A_{i+1}) \cdot P^-(\neg A_{i+1}) - (P(A_i|A_{i+1}) \cdot P^-(A_{i+1}) + P(A_i|\neg A_{i+1}) \cdot P^-(\neg A_{i+1})) \\ &= S_i \cdot \Delta_{i+1} \end{aligned}$$

This means, that for any n , $\Delta_0 = \Delta_n \cdot \prod_{i=0}^n S_i$.

For the support to be relevant, Δ_0 must be different from 0, which it only will be in what Atkinson and Peijnenburg call the *exceptional class* - where inference equals or asymptotically approaches entailment, and which we consider in section 2.6 above. For cases from the *usual class* Δ_0 will be zero, so infinitary support chains cannot change probabilities away from the prior. This is what Atkinson and Peijnenburg call the *fading foundation*. And as they remark, it is roughly the opposite of Bayesian washing out - here it is the *evidence* (at infinity) that washes out, and the *prior probability* that in the limit is the sole determinant of the posterior probability.

5.1 Virtually infinitary support

If there is a maximum depth we can reach, P^- includes precisely the information to be found up to that maximum depth. All other information is considered “infinitely far away”, as if we were in the exceptional class, and the same result as found in section 2.6 above obtains.

6 References

Atkinson, David and Jeanne Peijnenburg (2017). *Fading Foundations - Probability and the Regress Problem*. Synthese Library #383.

Carroll, Lewis (Dodgson, Charles Lutwidge) (1895). *What the Tortoise Said to Achilles*. Mind 104.14, pp. 278-280.

Finn, Suki (2019). *The Adoption Problem and Anti-Exceptionalism about Logic*. The Australian Journal of Logic.