Chapter 5
Finite Minds

Abstract
Can finite minds encompass an infinite number of beliefs? There is a difference between being able to complete an infinite series and being able to compute its outcome; and justification is more than mere calculation. Yet the number of propositions or beliefs that are needed in order to reach a desired level of justification for the target can be determined without computing an infinite number of terms: only a finite number of reasons are required for any desired level of accuracy. This suggests a view of epistemic justification as a trade-off between the accuracy of the target and the number of reasons taken into consideration.

5.1 Ought-Implies-Can

As in the past, the idea of infinite epistemic chains is still generally regarded as being nonsensical, and often for the same reasons. Scott Aikin has divided the various objections to infinite chains into two main categories: the *ought-implies-can* arguments, which are basically pragmatic in character, and the *conceptual* arguments.[1] In this chapter we deal with the first category; the conceptual arguments we will discuss in the next chapter.

Ought-implies-can arguments in effect contain all the different versions of the notorious finite mind objection, which was already raised by Aristotle. They imply that justifying our beliefs only counts as an obligation in so far as we are capable of doing so. Given our human finitude we cannot complete an

[1] Aikin 2011, Chapter 2.

D. Atkinson, J. Peijnenburg, *Fading Foundations*, Synthese Library 383,
DOI 10.1007/978-3-319-58295-5_5

infinite series of inferential justification, hence we are not obliged to perform this task. Aikin distinguishes two kinds of ought-implies-can arguments:

> On the one hand, there are arguments that the *quantity* of beliefs (and inferences) necessary is beyond us (for various reasons). This is the argument from quantitative incapacity. On the other hand, there are arguments that the quality (or kind) of belief necessary to complete the regress appropriately is one we simply cannot have. That is, because some belief in or about the series (and necessary for the series to provide epistemic justification) will be so complex, we cannot have it. And thereby, we cannot maintain the series in a way capable of amounting to epistemic justification. This is the argument from *qualitative* incapacity.[2]

The idea is straightforward enough: because we are mortal and of restricted capacity, we are unable to handle epistemic chains that either contain an infinite number of beliefs or contain some beliefs that are too complicated for us to handle.

But straightforward as it may seem at first sight, the idea is not always clear, and it has not always been expressed in the same way. Even among the philosophers who are most pertinacious in their disapproval of infinite epistemic chains, there is no agreement on this matter. For example, Michael Bergmann, as we have seen, deems it obvious that we cannot have an infinite number of beliefs:

> ...it seems completely clear that none of us has an infinite number of actual beliefs, each of which is based on another ...[3]

Noah Lemos agrees:

> One difficulty with [the option of an infinite chain] is that it seems psychologically impossible for us to have an infinite number of beliefs. If it is psychologically impossible for us to have an infinite number of beliefs, then none of our beliefs can be supported by an infinite evidential chain.[4]

But Richard Fumerton has a different opinion:

> There is nothing absurd in the supposition that people have an infinite number of justified beliefs.[5]

---

[2] Ibid., 52. The same distinction was made by John Williams when he discriminated between an infinite number of beliefs and an infinitely complex belief (Williams 1981).

[3] Bergmann 2007, 23.

[4] Lemos 2007, 48.

[5] Fumerton 2006, 49.

> Klein is right that we do have an infinite number of beliefs.[6]

> . . . there probably is no difficulty in supposing that people can have an infinite number of beliefs.[7]

This difference of opinion should perhaps not surprise us. After all, as noted earlier, it is entirely unclear how we should count our beliefs. This observation already intimates that knock-down arguments whether we *can* or *cannot* have an infinite number of beliefs are not to be expected.

Peter Klein has defended his infinitism against the finite minds objection by arguing that the objection is based on what he calls the 'Completion Requirement'. According to this requirement, a belief can be justified for a person only if that person has actually completed the process of reasoning to the belief. Such a requirement, says Klein, is against the spirit of infinitism indeed, but it is also unrealistic in that it is too demanding:

> Of course, the infinitist cannot agree to [the Completion Requirement] because to do so would be tantamount to rejecting infinitism. More importantly, the infinitist should not agree because the Completion Argument demands more than what is required to have a justified belief.[8]

Klein regards epistemic justification as being incomplete at heart: it is essentially provisional and can always be further improved. He fleshes out this view by means of two distinctions: the distinction between propositional and doxastic justification, and that between objective and subjective availability. Propositional justification, according to Klein, depends on the objective availability of reasons in an endless chain, where objective availability means that one proposition *is* a reason for another, so that it can be said to justify even if we are not aware of it. Doxastic justification, on the other hand, is parasitic on propositional justification and hinges on an availability that is subjective: a belief *q* is doxastically justified for an epistemic agent *S* if there is, in the endless chain of reasons, a reason for *q* that *S* can "call on". Although in its entirety the chain can never be subjectively available to *S*'s finite mind, *S* can take a few steps on the endless path. How many steps *S* can take, or needs to take in order to reach doxastic justification, all depends on contextual factors:

> Infinitism is committed to an account of *propositional justification* such that a proposition, *q*, is justified for *S iff* there is an endless series of non-repeating

[6] Fumerton 2001, 7.
[7] Fumerton 1995, 140.
[8] Klein 1998, 920.

> propositions available to S such that beginning with q, each succeeding member is a reason for the immediately preceding one. It is committed to an account of *doxastic justification* such that a belief is doxastically justified for S *iff* S has engaged in tracing the reasons in virtue of which the proposition q is justified far forward enough to satisfy the contextually determined requirements.[9]

We sympathize with Klein's view, but the previous chapters have made it clear that our position differs in two ways. On the one hand it is *weaker*: where Klein holds that justification requires the objective availability of an infinite chain, we allow that there can be justification even if the chain terminates. In those cases the foundation still exerts some justificatory influence of the target; and just how much justificatory influence it exerts depends on other characteristics of the chain, such as its length and the speed with which the series of conditional probabilities converges. On the other hand, our position is *stronger* than that of Klein: where he denies that infinite chains can be completed, we assert that they can. We only need to construe justification probabilistically and make sure that we are in what we have called 'the usual class', i.e. the domain where the probabilistic support is not too close to entailment.[10]

---

[9] Klein 2007a, 11. We have substituted q for p. *Cf.* Section 1.2.

[10] While some have taken the view that Klein's infinitism can account for propositional but not for doxastic justification, Jonathan Kvanvig has argued it fails on both counts. His argument why it fails for propositional justification goes as follows. In Klein's view, propositional justification either is relative to the total evidence available or is not so relative (where 'available' is interpreted liberally: a reason need not be present in order to be available, but may be only ready to hand). If propositional justification is *not* relative to the total evidence available, then my justification for q might depend on which book I happen to have taken from my shelves: "one source can be the start of an infinite chain of reasons for thinking [q], and the other source the start of an infinite chain for [$\neg q$]" (Kvanvig 2014, 140). If, on the other hand, propositional justification *is* relative to the total evidence available, then scepticism looms. Suppose that evidence E_1 confirms q, that E_2 confirms $\neg q$, and that $E_1 \wedge E_2$ does not confirm q. Let person S_1 have E_1 as evidence, S_2 have E_2, and the infinitist have $E_1 \wedge E_2$. Then, Kvanvig argues, "if propositional justification is relative to total information", none of these three have justification for q or for $\neg q$. Kvanvig's argument rightly points to vagueness in the term 'availability', whether interpreted liberally or strictly. However, his argument seems to presuppose an 'absolute' concept of justification and moreover to equate justification with confirmation. With the relational concept of justification that we proposed in Chapter 2, and with the assumption that confirmation is necessary but not sufficient for justification, there does not seem to be a problem.

5.2 Completion and Computation

On the basis of the previous chapters our answer to the finite mind objection will not come as a surprise. If justification is probabilistically construed, then even the 'Completion Requirement' that Klein rebuts can be met.[11] For then infinite justificatory chains *can* indeed be completed in the sense that they yield a unique and well-defined probability value for the target proposition. And if it is possible to complete infinite chains, the finite mind objection does not arise. Although this answer to the finite mind objection differs from that of Klein, who after all asserts that completion and infinitism are irreconcilable, it does enable us to account for at least two of Klein's intuitions, namely that epistemic justification gradually emerges along the chain and that contextual factors decide at which level of emergence we will decide that 'enough is enough'.[12]

However, Jeremy Gwiazda has argued that this reply to the finite mind objection does not work. As he sees it, we have not *completed* a probabilistic regress, but we have only *computed* its limit.[13] There is a great difference, according to Gwiazda, between calculating the probability value of a target proposition on the one hand and actually giving reasons for that proposition on the other. Gwiazda does not discuss in detail what the differences are, but he might be thinking of a difference in time: while we can calculate the limit of an infinite series in a finite time, we are unable to come up, in a finite time, with an infinite number of reasons. As such, the difference resembles an important distinction that Nicholas Rescher has emphasized, namely between regresses which are time-compressible and those which are not. An example of the former is generated by the Zeno-like thesis 'To *reach* a destination, you must first *reach* the halfway point to it'; an example of the latter is produced by 'To *make a journey* to a destination, you must first *make a journey* to the halfway point to it':

> The first thesis is true — and harmless: that is just how transit from point *A* to point *B* works. But the second is false and, moreover, vicious in rendering any sort of journey impossible. Zeno of Elea notwithstanding, a motion to reach or to cross endlessly many points is perfectly possible. But infinite *journeying*,

[11] This point appears to have been missed in Wright 2013.

[12] This is basically the way in which Frederik Herzberg, referring to insights about probabilistic regresses, has replied to the 'new finite mind objection' that was raised by Adam Podlaskowski and Joshua Smith. See Herzberg 2013, 373-374, and Podlaskowski and Smith 2011.

[13] Gwiazda 2010. The same point was made by Matthias Steup (Steup 1989).

> with its inherent requirement for explicitly planned and acknowledged transits, is an impossibility. And the reason for this lies not in the impossibility of motion, but in the fact that making a journey to somewhere (as *distinct* from reaching or arriving there) involves deliberation and intentional goal-setting. And since man is a finite being, an infinitude of conscious mental acts is impossible for us. So while that first *structural* regress is harmless, the second regression of infinitely many consciously performed acts is an impossibility.[14]

In the same vein, it could be admitted that we, with our finite minds, are capable of calculating the probability of a target proposition (in the previous chapters we have after all done so), but are incapable of giving an infinite number of reasons for this proposition, since the latter would require an infinity of consciously performed acts. Because epistemological justification is about giving reasons, and not about making calculations, the finite mind objection applies in full force.

A similar reaction to our views has been voiced by Adam Poslaskowski and Joshua Smith.[15] They argue that, although "valuable lessons" can be drawn from our formal results, it is "entirely unclear" that these results meet a basic requirement, namely "providing an account of infinite chains of propositions *qua* reasons made available to agents".[16] Podlaskowski and Smith call this 'the availability problem':

> Given the distinctive emphasis that Peijnenburg, Atkinson, and Herzberg place on calculability, we have doubts about the extent to which (on their account) an infinite chain of propositions can serve as *reasons* that are *available* to an agent. (This is what shall be called the *availability problem* facing the distinctive brand of infinitism under consideration).[17]

> ...it is hard to see, more generally, how the emphasis on calculability yields a notion of *available reason* (or *availability*) that can serve the infinitist's purposes.[18]

Podlaskowski and Smith maintain that our analysis confuses two completely different things, namely being able to compute the probability of a target

[14] Rescher 2010, 25. Rescher uses several ways to express the distinction between time-compressible and non-time-compressible regresses; one of them is by saying that the latter need pre-conditions whereas the former only has co-conditions (ibid., 55).

[15] Podlaskowski and Smith 2014.

[16] Ibid., 212.

[17] Ibid., 214.

[18] Ibid., 215.

proposition on the one hand and having available reasons for this proposition on the other. They blame us for assuming that, "since mathematical means exist with which an agent can decide the probability of any proposition being true (even if it belongs to an infinite series), all the members of an infinite chain of reasons must thereby be *available* (as reasons) to an epistemic agent".[19] Like Gwiazda, they stress the difference between determining the probability of a target q and showing that something is a reason for q:

> deciding the probability of any given proposition ... even if there are infinite chains of propositions ... is still a far cry from showing that, as a matter of principle, each proposition in a chain of of propositions is one that can serve as a *reason* for another proposition in that chain, and do so in the right order. It appears that two dispositions have been conflated: those to make a certain sort of calculation, and those to accept any given proposition as reason for another proposition. ... [A] demonstration that finite agents can actually calculate the probability of a proposition's truth — even if it belongs to an infinite chain of reasons — does not thereby show that each reason is equally *available* to a finite agent.[20]

The observation of Gwiazda and Podlaskowski and Smith that computing and completing reflect two different dispositions is fair enough. However, as we will explain in the next section, in epistemic justification we draw on both. In this sense, justification resembles logic: there, too, we draw on an abstract, normative dimension concerning how one *ought* to reason, and a concrete, descriptive dimension concerning how one reasons *in fact*.[21] Together the two dimensions suggest a view of justification as a trade-off between the accuracy of the target proposition and the capacity of our mental housekeeping.

5.3 Probabilistic Justification as a Trade-Off

Rescher is of course right that a time-compressible regress is different from a regress that is not time-compressible. And Gwiazda and Podlaskowski and Smith are right that making a calculation is not the same as giving a proposition as reason for another proposition. The skill to compute the value of the

---

[19] Ibid., 215.
[20] Ibid., 216. Michael Rescorla's complaint that our approach falls prey to 'hyper-intellectualism' expresses a similar sentiment (Rescorla 2014).
[21] Van Benthem 2014, 2015.

target on the basis of a probabilistic epistemic chain indeed differs from the capacity to have the propositions in the chain available as reasons.

However, these two faculties are not disjunct, as the above authors seem to think.[22] Especially when it comes to epistemic justification, of which probabilistic support is an essential part, these faculties are closely and essentially connected. A justificatory regress is not just any old regress; it is a regress about *reasoning*, in this case reasoning that involves how a proposition or belief is probabilistically justified by another. This means that the actual process of 'giving probabilistic reasons' is to a certain extent subjected to the rules of the probability calculus, just as the actual process of 'giving deductive reasons' is to a certain extent subject to the rules of deductive logic. The aversion of Gwiazda and others to using calculations in the context of giving reasons might be exacerbated by the idea that this necessarily involves processing an infinite number of terms. That idea, although understandable, is however mistaken, and betrays a misconstrual of our view.

We have argued that, whenever we give a reason, A_i, for a target q, the significance of A_i as a reason depends on how much probabilistic support it gives to q. The latter in turn depends on how much A_i's support for q deviates from the 'final' support, i.e. the support that q would receive from the entire justificatory chain of which A_i is a member. And how much support q receives from the entire infinite chain depends on the chain's character, i.e. on the values of its conditional probabilities together with the value of the unconditional probability of the ground, p. While the conditional probabilities come from experiments, the unconditional probability of the ground is unknown.[23] The longer the chain, the smaller the contribution from the ground, and when the chain is infinitely long, the contribution from the ground to the target vanishes completely, leaving all the justificatory support to come from the combined conditional probabilities.

The view could be easily misunderstood. It does *not* imply that 'giving reasons' depends on 'making calculations' in the sense that we first have to calculate the limit of a probabilistic regress before we can know what our reason is worth; computing the limit is *not* necessary for weighing the quality of our actual reasons. Rather, the structure of the probabilistic justificatory

[22] Recall the claim of Podlaskowski and Smith that it is "entirely unclear" what formal calculation means for "propositions *qua* reasons made available to agents" (Podlaskowski and Smith 2014, 212).

[23] In Chapter 8, Section 8.5, we will come back to the status of the conditional probabilities. In particular, we consider the situation in which they are not given, but are themselves in need of justification. As we will explain, a network is then created with a remarkable structure that resembles a Mandelbrot fractal.

chain is such that it enables us to say *how many reasons* we need to call on in order to approach the probability of the target to a satisfactory level. To do that, we do not need to know the length of the chain; we need not even know whether it is finite or infinite. Nor do we have to know the probability of the ground. The only thing we need are the values of a certain number of conditional probabilities (sometimes more, sometimes less, depending on the speed of the convergence) that suffice to take us to within a desired level of accuracy with respect to the true, but unknown probability of the target. Once we are there, we can safely ignore the rest of the chain — such is the lesson of fading foundations.

An example might help to understand the point. Imagine I have a reason A_1 for my belief q, and know the two relevant conditional probabilities, $P(q|A_1)$ and $P(q|\neg A_1)$. Suppose I am unable or unwilling to back up A_1 by a further reason, and therefore want to cut off the chain here. We have seen that knowing the conditional probabilities is in general not enough to know the value of $P(q)$; especially with short chains like the one at hand it is indispensable that we also know the unconditional probability $P(p)$. Even if I have no clue what the value of the latter is, I do know that it cannot be greater than one and cannot be smaller than zero. I now consider these two extremal cases, i.e. where $P(p) = 1$ and where $P(p) = 0$, and I find that in the first case $P(q) = x$ and in the latter case $P(q) = y$. The condition of probabilistic support now guarantees that the real value of $P(q)$ lies in the interval between x and y, *no matter how many further* A_n we take into consideration. What is more, the condition ensures that with every reason we add, the interval will *become smaller*, making the value of $P(q)$ more precise with each step. This applies both in the uniform situation, where the conditional probabilities are all the same, and in the nonuniform case, where they are different.

As a result, I can determine how many reasons I need to have in order to approach the true probability of the target q within an error margin of, for example, 1%. If this number of reasons happens to be too large to fit into my finite mind, then I will have to relax the level, and be content with a degree of justification that is further away from the true probability of the target. But if the number of reasons is rather small, so that they all fit in my finite mind (although perhaps not in that of my four-year-old daughter), then I can always tighten up the satisfaction level, and come closer to the target's true probability. Epistemic justification thus boils down to striking a balance. In acting as responsible epistemic agents, we are instigating a trade-off between the number of reasons that we can handle and the level of accuracy that we want to reach. If we are unable or unwilling to manage a large number of reasons, we have to pay in terms of a lack of precision and hence of trustworthiness. Taking the short route thus comes at a price, but in situations where precision is not important, we can take it easy and should do so on pain of exerting ourselves unnecessarily.

Let us spell out this idea of a trade-off more fully and more formally. Assume a finite chain to consist of five propositions, the target proposition q, the intermediate propositions A_1 to A_3, and the ground A_4:

$$P(q) = \beta_0 + \gamma_0\beta_1 + \gamma_0\gamma_1\beta_2 + \gamma_0\gamma_1\gamma_2\beta_3 + \gamma_0\gamma_1\gamma_2\gamma_3 P(A_4)\,. \tag{5.1}$$

As we explained in Sections 3.5 and 3.6, the right-hand side consists of two terms. The first term is the sum of the conditional probabilities,

$$\beta_0 + \gamma_0\beta_1 + \gamma_0\gamma_1\beta_2 + \gamma_0\gamma_1\gamma_2\beta_3\,,$$

and the second is the remainder term,

$$\gamma_0\gamma_1\gamma_2\gamma_3 P(A_4)\,.$$

This remainder term is a product of two factors, $\gamma_0\gamma_1\gamma_2\gamma_3$ and $P(A_4)$. Since we suppose the conditional probabilities to be known, there is only one probability that we need to know in order to compute $P(q)$. This is $P(A_4)$, i.e. the unconditional probability of the ground. If we did know $P(A_4)$, then we would know $P(q)$.

However, suppose we have no clue as to the value of $P(A_4)$. What to do? Because of the condition of probabilistic support, (2.1), all the γ_n are positive, which means that every term in (5.1) is positive too. Therefore the smallest value that $P(q)$ could have, given the conditional probabilities, is obtained by giving $P(A_4)$ the minimum value that it could have, which is zero, leaving only

$$\beta_0 + \gamma_0\beta_1 + \gamma_0\gamma_1\beta_2 + \gamma_0\gamma_1\gamma_2\beta_3\,. \tag{5.2}$$

On the other hand, the largest value that $P(q)$ could have is obtained by giving $P(A_4)$ the maximum value that it could have, which is one, yielding

$$\beta_0 + \gamma_0\beta_1 + \gamma_0\gamma_1\beta_2 + \gamma_0\gamma_1\gamma_2\beta_3 + \gamma_0\gamma_1\gamma_2\gamma_3\,. \tag{5.3}$$

We know that the value of $P(q)$ must lie somewhere between the two extremes (5.2) and (5.3). If we were to assume the value of $P(q)$ to be one extreme, for example (5.2), then we would be sure that our error could not be larger than the difference between the maximum, (5.3), and the minimum, (5.2), namely $\gamma_0\gamma_1\gamma_2\gamma_3$.

Now imagine that the error term $\gamma_0\gamma_1\gamma_2\gamma_3$ turns out to be, for example, only 1% of the minimum value (5.2). And suppose further that we proclaim ourselves satisfied with a value that deviates by no more than 1% from the true value of $P(q)$. Then we need go no further in inquiring as to any support that the ground, A_4, might have from some other proposition. This is because any extension of the chain, obtained by adding a proposition, A_5, that supports the erstwhile ground A_4 would only *increase* the minimum (5.2) to

$$\beta_0 + \gamma_0\beta_1 + \gamma_0\gamma_1\beta_2 + \gamma_0\gamma_1\gamma_2\beta_3 + \gamma_0\gamma_1\gamma_2\gamma_3\beta_4\,,$$

and *decrease* the error to $\gamma_0\gamma_1\gamma_2\gamma_3\gamma_4$ (this *is* smaller, because the extra factor, γ_4, is less than one). This is precisely what fading foundations imply. So in this case we know exactly how many reasons we need in order to approach the true value of the target to a level that satisfies us. If we are content with a value that deviates no more than 1% from the true value of $P(q)$, then we require no more than four reasons for q, namely A_1 to A_4. And if our mind is big enough to store these four reasons, then we have accomplished our task: we have justified q to a satisfactory level, staying neatly within the limitations of our finite mind. Note that we have performed our task without knowing the true value of $P(q)$ or that of $P(A_4)$.

What to do when the error term $\gamma_0\gamma_1\gamma_2\gamma_3$ turns out to be very big, for example 90% of the minimum value (5.2)? How should we proceed now? If our level of required accuracy is still 1%, then there is not much that we can do in this case. We might sadly conclude is that there is much uncertainty, due to the fact that the justificatory influence of the unknown $P(A_4)$ on $P(q)$ is very great, but that is as far as we can get. For the four reasons that we can avail ourselves of, A_1 to A_4, are of little help: jointly they bring us to a point where the deviation from the true value of $P(q)$ may be as great as 90% .

However, let us now make the finite chain considerably longer. Rather than assuming that there are four reasons for q, let us suppose that there are one hundred:

$$P(q) = \beta_0 + \gamma_0\beta_1 + \gamma_0\gamma_1\beta_2 + \ldots + \gamma_0\gamma_1\ldots\gamma_{m-1}\beta_m + \gamma_0\gamma_1\ldots\gamma_m P(A_{m+1})\,, \tag{5.4}$$

where $m = 99$. It is unlikely that I can store all these reasons in my finite mind, so I decide to cut off chain (5.4) at number seven, making a provisional stop at proposition A_6. So I get:

$$P(q) = \beta_0 + \gamma_0\beta_1 + \gamma_0\gamma_1\beta_2 + \ldots + \gamma_0\gamma_1\gamma_2\gamma_3\gamma_4\beta_5 + \gamma_0\gamma_1\gamma_2\gamma_3\gamma_4\gamma_5 P(A_6)\,. \tag{5.5}$$

In formula (5.5) I can only compute $P(q)$ if I know $P(A_6)$. Since I have no idea as to the value of the latter, I apply the same reasoning as above. That is, I first recall that the value of $P(q)$ must lie between two extremes. The one extreme is obtained by putting the unknown $P(A_6)$ equal to zero. The other extreme is obtained by putting it equal to one. Suppose that I adopt the first extreme, $P(A_6) = 0$, so my estimation of $P(q)$ is that it has its minimum value. I know that my error in making this estimation cannot be larger than the difference between this minimum value and the maximum value of $P(q)$, obtained by putting $P(A_6) = 1$. The difference itself is given by our error term, which in this case is $\gamma_0\gamma_1\gamma_2\gamma_3\gamma_4\gamma_5$.

Now suppose that the error term $\gamma_0\gamma_1\gamma_2\gamma_3\gamma_4\gamma_5$ is only 1% of the minimum value of $P(q)$. And suppose again that I am satisfied with an accuracy that deviates no more than 1% from the true value of $P(q)$. If I am capable of storing six reasons in my head, then I am done. In particular, I do not have to go on and find a justification for A_6. A we have seen, the reason for this lies in the fact that, as the chain lengthens, the minimum value of $P(q)$ increases and the maximum value decreases — which is a direct consequence of the condition of probabilistic support. This condition implies that any extension of the chain would only make the minimum value of $P(q)$ greater, and thus would make the error term itself smaller.[24] Consequently, adding a proposition to chain (5.5), for example proposition A_7, would bring us closer to the true value of $P(q)$; and since we are already satisfied with our level of approximation, there is no need to engage in this project. We have in fact reached the point where 'enough is enough', and this expression now has a very precise meaning. For any justificatory chain, I can first define the level of accuracy within which I want to approach the true value of $P(q)$, and I can then determine how many reasons I need to reach this level. In order to perform these tasks, I need not know the value of $P(q)$, nor that of $P(A_6)$, nor that of any other ground. More importantly, I can blissfully neglect the rest of the chain. For not only is it so that I am within the desired 1% of the true probability value, it is also the case that calling on any further reason will only bring me closer to that true value. As the chain gets longer, the remainder term gets smaller (in accordance with fading foundations) and the sum of the conditional probabilities gets larger (in accordance with the condition of probabilistic support). So as m gets bigger, the value of the sum of the conditional probabilities increases monotonically, whereas the remainder term decreases monotonically. Therefore, if we are already satisfied with 1%, any extension of the chain will bring us still closer to the true value of $P(q)$; there is thus no need to call on more reasons than the six reasons that we have (subjectively) available.

But now suppose that the error term $\gamma_0\gamma_1\gamma_2\gamma_3\gamma_4\gamma_5$ in formula (5.5) still greatly differs from the minimum value of $P(q)$; let us now say by 80%. The

---

[24] See Appendix A.2 for the proof.

situation is not the same as it was in the case (5.1). Since Eq.(5.5) is part of a larger chain we have the option to go on, and to look for the justification of A_6 in terms of A_7; after that, we can go further and justify A_7 in terms of A_8, and so on. The more propositions we add, the more we lengthen our chain, and the smaller will be the difference between (5.4) and the minimum value of $P(q)$. We are now able to reduce the error to less than 80% of the true value of $P(q)$. However, there is a price to pay. In getting closer and closer to the real value of $P(q)$, we are calling on more and more reasons, and our finite minds have to accommodate each and every extra reason that we call on. It could happen that our minds lack the capacity to take in all the reasons that our level of accuracy requires. In that case the only option left open for us is to relax the accuracy level to a degree where it corresponds to a number of reasons that *can* be housed in our heads. We are committed to a trade-off: we simply cannot have our cake and eat it too.[25]

What we have said above is of course not restricted to finite chains such as (5.1), (5.4), and (5.5). The reasoning about error terms works just as well with an infinite chain as with a finite chain. In both cases we can work out, in a finite number of steps, how many terms we need to reach a particular, pragmatically determined level of accuracy. If it turns out that our level of accuracy requires more reasons than we can accommodate, then we are living

[25] Whether a particular number of reasons can or cannot be housed in our heads might depend not just on size or on capacities, but also on other factors. Linda Zagzebski has distinguished between two kinds of epistemic reasons for believing a proposition q: theoretical reasons, which are third personal and "connect facts about the world with the truth of [q]", and deliberative reasons, which are first personal and "connect me to getting the truth of [q]" (Zagzebski 2014, 244). Even if, impossibly, we were able to complete our search for theoretical reasons, that would still leave us with the second problem that what we call 'reasons' may not indicate the truth: "We would still need trust that there is *any* connection between what we think are the theoretical reasons and the truth" (ibid., 250). Zagzebski argues that this second problem can only be solved by calling on a deliberative reason with a special status, viz. epistemic self-trust, which ends our urge to search for further theoretical or deliberative reasons. It is not excluded that Zagzebski's epistemic self-trust might be a factor in the process of trading-off. Other possible factors might perhaps be the localist considerations of Adam Leite, or the "plausibility considerations" that Ted Poston mentions in support of his claim that "there is more to epistemic justification than can be expressed in any reasoning session" (Leite 2005; Poston 2014, 182-183). We expect that our trade-off can also be combined with Andrew Norman's "dialectical equilibrium" (Norman 1997, 487) and Michael Rescorla's "dialectical egalitarianism" (Rescorla 2009), although we are not sure if the authors themselves would agree.

beyond our means. We then should either work harder and try to create more space in our finite minds, or become more modest and lower our desire for accuracy. All this can be done without having to call on, or even to calculate, all the terms in a (finite or an infinite) series.

The two tables below illustrate the idea. In the first, the conditional probabilities, α and β, have the values 0.99 and 0.04 respectively; in the second, they are 0.95 and 0.45. 'Maximum $P(q)$' and 'Minimum $P(q)$' refer to the values that $P(q)$ has when $P(p)$ is one or zero, respectively.

Table 5.1 Extremal values of $P(q)$ when $\alpha = 0.99$ and $\beta = 0.04$.

Number of A_n	1	2	5	10	15	25	50	100	∞
Minimum $P(q)$	.078	.114	.212	.345	.448	.589	.742	.796	.8
Maximum $P(q)$	.981	.971	.947	.914	.888	.853	.815	.801	.8

Table 5.2 Extremal values of $P(q)$ when $\alpha = 0.95$ and $\beta = 0.45$.

Number of A_n	1	2	3	4	5	6	8	10	∞
Minimum $P(q)$	.675	.788	.844	.872	.886	.893	.8982	.8996	.9
Maximum $P(q)$	.925	.913	.906	.904	.902	.901	.9002	.9000	.9

In the first table one needs more than fifty intermediate reasons A_n to ensure that the difference between the maximum and the minimum of $P(q)$ is relatively small, whereas in the second table a similar uncertainty is already reached after a mere three reasons A_n. There the situation is much more amenable. Justification as a form of trade-off sheds light on the difference between propositional and doxastic justification that we discussed in 4.2. Some scholars appear to be of the opinion that propositional and doxastic justification can never be combined, since the former is abstract and infinite, while the latter is concrete and finite by definition.

Others have however argued that doxastic justification is parasitic on propositional justification, and that the context determines when exactly it comes to an end. Our considerations in this section clarify the latter position, and they make clear how this contextualism can be interpreted.

5.4 Carl the Calculator

When commenting on our approach, Podlaskowski and Smith write that "care must be taken when assessing the significance of these formal results".[26] Of course we agree, and it can be added that the same applies to assessing results that are not formal: whenever we informally discuss reasoning, or justification, or probability, we must take care what we say. For example, as we have seen, it is incorrect to say that an infinite probabilistic regress yields zero for the target, or that knowing the value of the target requires knowing the value of a basic belief. Intuitive as these claims might be, they are incorrect as they stand.

The difference of opinion between Podlaskowski and Smith and us, if there is one, concerns the relation between the ability to calculate and the ability to give reasons. As we explained in the previous section, we believe that epistemic justification involves both. Podlaskowski and Smith seem however to interpret us differently, thinking that for us having the mathematical ability to calculate is sufficient for having justification. This is for instance the message from their instructive example about Carl, who is a real pundit when it comes to calculating probabilities, but who cannot understand the meaning of reasons:

> [I]magine Carl, whose impressive talent in calculating conditional probabilities is strangely at odds with his ability to grasp various concepts. Carl has no problem solving all manner of complex equations, including those involving conditional probabilities (such as Peijnenburg, Atkinson, and Herzberg provide). Yet, there are various concepts which he is entirely incapable of grasping, some of which might feature in reasons whose probabilities of being true are conditional on other reasons. Suppose that Carl is given two lists, an infinite list of conditional probability assignments and an infinite list of reasons. Unbeknownst to Carl, the two lists correspond perfectly: the list of probabilities is meant to capture the probability of each reason being true, conditional on its predecessor. Moreover, some of the members of the list of reasons are comprised of those concepts that Carl is incapable of grasping. Even if Carl were capable of working through some infinite list of reasons, at some point on the list at hand, Carl would fail to comprehend the concepts deployed. But he would have *no problem* doing the corresponding calculations. Does merely calculating the probability of the chain make Carl justified in holding any of those beliefs, when Carl is *incapable* of understanding the concepts on which those beliefs depend? Surely not. If an agent *cannot* understand some of the

[26] Podlaskowski and Smith 2014, 212.

reasons in the infinite chain, it is difficult to see how those reasons can do any justificatory work for him.[27]

Podlaskowski and Smith suggest that, according to us, Carl has justified his beliefs. This is however not so: for us, as for Podlaskowski and Smith, Carl fails to justify. Our view is not that calculation implies justification, but that justification implies a certain amount of 'calculation'. Of course we realize that people often put forward probabilistic reasons for their beliefs without knowing anything about the probability calculus. As epistemologists who want to take the concept of probabilistic reasoning seriously, however, we believe that a minimum of adjustment to the probability calculus seems to be required, even if it is only in a rational reconstruction.

Podlaskowski and Smith seem to have anticipated this response when they write:

> One might . . . suspect that we have crafted the Carl case too narrowly, and that it misses some important aspect of what mathematical analyses of probabilistic regresses are supposed to be doing.[28]

However, they then suggest that our response requires a new notion of 'available reason' which cannot be developed within our approach:

> Perhaps there is a notion of *available reason* that can supplement the project of Peijnenburg et al. that avoids the problems raised by the Carl case. The problem with successfully developing such a response, however, is that it is entirely unclear what sort of notion they could use, given their emphasis on calculability. To see this, consider the spectrum of possible views. On one end, the notion of *availability* drops out. This end of the spectrum has the unfortunate consequence that the view collapses into maintaining that a belief is justified for a person when there merely exists an infinite, non-repeating chain of reasons that makes the belief probable. . . . On the other end of the spectrum, one might hold a very strong notion of *availability*, according to which it is required that one *actually believe* a reason for it to be available. But this is far too strong, as it runs face-first into the original finite minds objection to infinitism. . . . One lesson to draw from the Carl case is that moving a brand of infinitism beyond Klein's middle ground on the notion of *availability* proves seriously problematic . . . .[29]

Here Podlaskowski and Smith write as if there are only two possibilities: either we merely calculate, and then no reason qua reason is available, or

[27] Ibid., 216.
[28] Ibid., 217.
[29] Ibid.

we hold on to a strong notion of availablity, but then we run into the finite mind objection. Our remarks in the previous section provide us with a notion of availability that avoids the two extremes that Podlaskowski and Smith present. Often it is enough that only a few reasons are available in order to draw conclusions that go far beyond what is implied by these available reasons themselves. If the reasons in question bring us close enough to the true value of the target, then the phenomenon of fading foundations tells us that we can ignore the rest of the chain. If, on the other hand, the reasons do not bring us within a desired level of accuracy, then we will have to achieve a balance between the number of reasons that we can handle and the degree to which we can approach the final value of the target. Thanks to the condition of probabilistic support we can determine how many reasons we need in order to conclude that the rest of the chain is irrelevant. In this chapter we have explained the idea in quantitative terms, but it can quite easily be grasped in an intuitive and qualitative way.