

Fools ape angels: epistemology for finite agents

(For Gardenfors e-schrift)

Adam Morton, University of Bristol
adam.morton@bristol.ac.uk

Soyez raisonnables, demandez l'impossible.

Paris 1968

Si [Dieu] ne m'a pas donné la vertu de ne point faillir, par le premier moyen que j'ai ci-dessus déclaré, qui dépend d'une claire et évidente connaissance de toutes les choses dont je puis délibérer, il a au moins laissé en ma puissance l'autre moyen, qui est de retenir fermement la résolution de ne jamais donner mon jugement sur les choses dont la vérité ne m'est pas clairement connue. Car quoique je remarque cette faiblesse en ma nature, que je ne puis attacher continuellement mon esprit à une même pensée, je puis toutefois, par une méditation attentive et souvent réitérée, me l'imprimer si fortement en la mémoire, que je ne manque jamais d'en resouvenir, toutes les fois que j'en aurai besoin, et acquérir de cette façon l'habitude de ne point faillir.

Descartes, fourth meditation

1. Epistemic virtues and ideals: Epistemology as metacognition What is the point of epistemology? Some of the reasons for asking questions about knowledge, evidence, and reasonable belief come from within philosophy. In particular there are the motives arising from scepticism, which some philosophers take very seriously. But epistemology would exist quite independently of all that. Something like epistemology will be a part of human life as long as people criticise and encourage one another's beliefs. And presumably people every where and when have said things to one another which play the role of "why on earth do you think that?", "if you believe A in circumstances C then you should believe B in C' ". Moreover this process must always have been internalized, people must always have been asking themselves critical questions of their own beliefs. In fact, if some philosophers are right, for example Davidson and Shoemaker, in order to have beliefs at all it is necessary to put some effort into keeping them tidy and presentable. If some psychologists are right every person has to learn an individual set of routines for managing their own thought, which are only with great difficulty communicated from one person to another. (This is a fashionable topic in psychology at present, labelled 'metacognition'. But there do not seem to be real theories of it!) When we do succeed in transmitting tricks of mind-management we do it in part by using the normative vocabulary of belief.

This is not to say that even without philosophical concerns we would talk in exactly the same way about beliefs and the reasons for holding them. The purpose of this paper is to argue for the opposite conclusion. There is a gap between the shape of epistemology, its vocabulary and dominant concerns, as it responds to traditional philosophical theories and the shape demanded by our need to structure, criticise, and improve our beliefs. In

particular epistemology has absorbed assumptions about human limitations and the best way to respond to them into its choice of vocabulary and problems, assumptions which would now seem very debatable, but which remain invisibly effective as disguised presuppositions of the epistemological vocabulary.

The obvious link between epistemological theory and everyday epistemic practice is the vocabulary of epistemic virtues and vices. We speak of beliefs as being sensible or off-the-wall, well-supported or conjectural, orthodox or eccentric, and we describe believers as careful or sloppy, responsible or wild, meticulous or daring. There is now a large literature on virtue epistemology, associated with such writers as Hookway and Zagzebski. (But I must admit that nothing in this literature convinces me that there *are* any specifically epistemic virtues, as opposed to virtues of intelligent activity generally.)

Virtues suggest images of ideal virtuousness, conceivable situations in which they are completely instantiated. Such ideals may or may not be attainable, and in fact there are virtues which it is hard to think of as being ideally instantiated. For example there is the virtue of coping well with ones limitations. It is hard to think of this as part of the equipment of an ideal agent. (Or not? Curious question: could one be ideally imperfect?) Let me say more about one aspect of this, as it will be important in what follows.

Suppose that we have a concept of ideal functioning, whether in epistemology, ethics, or any other intelligent activity. And suppose that some particular agent is not capable of fully satisfying that ideal. What is the best course of action for the agent? There are two general strategies that she can follow.

One, I shall call it *the A strategy*, is to approximate to the ideal, to fulfil as many of its requirements as possible. Given a long list of actions that would be performed by an ideal agent an A-strategy agent will perform as many as she can. Given a complex cognitive process that would be part of ideal functioning she will instantiate as many parts of the process as she can.

The other, which I shall call *the B strategy*, is to try to satisfy requirements which can be met, and whose complete satisfaction will give better results than partial satisfaction of the requirements of the ideal. (More specifically, better not than any old partial satisfaction but the partial satisfaction that would result from trying hard and seriously to achieve the ideal.) Given a long list of actions that would be performed by an ideal agent a B-strategy agent may perform none of them, not even those that would be easy for her, if performing acts on a quite different list would give better results. (Very often, though, the ideal list and the alternative list will overlap.) Given a complex cognitive process that would be part of ideal functioning she may instantiate a quite different process which has nothing in common with it.

When is it appropriate to follow the A strategy and when the B? In a way the answer is obvious: the A strategy is appropriate when partial satisfaction of the ideal has better results than complete satisfaction of any other rule. But when is this? Roughly, when you can replace "all" and

”maximum” with ”many” and ”a lot”. There are reasons to believe that it is very hard to know when this is. For many famous computationally intractable problems have the form of a demand to find the maximum or minimum of some function $f(n)$ which does not pass smoothly from n to $n+1$. For example the ”shortest telegraph network” (Steiner tree problem) problem of placing additional points among n points on a plane so that the resulting set of connections between them is as short as can be is hard precisely because the solution for n does not give you much of an idea what the solution for $n+1$ may be. Let me state this as a rough working principle: *maximization and approximation do not mix*. The relevance to A and B strategies is just this: if you are an agent trying to solve such a problem for a large n and you do your humble best by using the solution for some $m < n$ your actions are likely not to approximate to the best at all. Doing your best may produce the worst. (For a stimulating analysis of this situation see Frank Jackson and Robert Pargetter, ’Oughts, Options, and Actualism’ *Phil Rev* 95, 1986:233-255.)

In giving someone advice you certainly don’t always follow the A-strategy. Or in judging how reasonable someone’s attitude was. To decide when to give A- and when B-flavoured advice you need to know what the facts are about the task and about the agent’s capacities. The better your knowledge the better. There is a potential discipline whose aim is to apply the psychology of reasoning — especially the psychology of error — to work out the best forms of epistemic advice: what are the aims that it pays to aspire to, and what are the problems and virtues we should keep in mind? It would be a *normative naturalised epistemology*. I think this is something well worth developing — and the more imaginatively the better. It is in some ways more like the aims of the old epistemologists than what we now do under the label. (In a way Bayesianism is a set of answers to these questions. Wrong ones.)

Normative naturalised epistemology would not be a rival to most philosophical epistemology. It would be a coherent evidence-based adjunct to, and improvement on, what we might call *folk epistemology*, our everyday judgements of the tenability of beliefs and the virtues and vices we remind one another of in order to keep our beliefs up to scratch. Standard philosophical epistemology also builds on folk epistemology, and then gets involved in grand questions about scepticism and the tenability of whole systems of belief and the possibilities of Science. We don’t really know what folk epistemology looks like (any more than folk psych, and like that, whether there really is such a *unified* single thing. Cf remark about epist virtues and virtues of intelligent activity above.) My impression — its just that — is that folk epistemology makes far more distinctions than we do in philosophy. For one thing, it is permissive rather than compulsive. It says when you *may* rather than when you *must* believe something. And it applies may to other states than belief: some things are reasonable conjectures, for example, though it would be crazy under the circumstances to believe them. I think folk epistemology tends to be externalist in its account of what one may and may not believe. It often appeals to the circumstances agents find themselves in rather than the thoughts that are running through their heads. And it does

not use the philosopher's notion of justified belief. It distinguishes different ways in which a belief can be ok: or so I believe, but given my intentions for the rest of the paper you had best note this as something to think about before agreeing with me. Common sense has a variety of terms of approval and disapproval for beliefs: sensible, reasonable, acceptable, plausible, intelligent,.../crazy, unreasonable, unacceptable, implausible, unjustified, stupid,... (Curiosity: I don't think we often use the labels "reasonable" and "justified", though their negatives are frequent.) More systematically: we distinguish in every day criticism of belief between evaluation and authority. Evaluation is questions of whether p is a good thing to believe: "why I think p is true". Authority is questions of whether S is well placed to know (be trusted, relied on) that p: "why if S thinks p we should take it seriously". I don't think there is a tight common sense connection between these. We can say that someone's belief is mad though they are well placed to hold it, or vice versa.

There are no real conclusions at this stage, but questions. Are there epistemic situations where the B-strategy is appropriate? Are there B-type epistemic virtues which are worth describing, that is, virtues which consist in finding and following a path which does not approximate to the one an ideal epistemic agent would follow? I believe that the answer to both of these questions is Yes, and in what follows I shall try to convince you of this. That answer may have a fairly high initial plausibility, though making the question clear and actually arguing for it are not easy. But in case you are not interested in arguments for unsurprising conclusions I shall mention another question. Does the existence of B-type epistemic virtues suggest that the language of epistemology is not well suited to serve its practical function of articulating criticism and evaluation of our beliefs? If the answer to this question is Yes then normative naturalised epistemology, or informed folk epistemology, would be quite different from philosophical epistemology. I shall suggest that the answer to this question *is* Yes. And it is certainly far from obvious that I am right about that one.

2. Cases Here are some cases in which people are and are not functioning as they should, and which it is hard to say in standard epistemological terms what is right and not right about their beliefs.

(a) *the climate modeller* A politician wants to know what the rainfall in the Sahara will be in twenty years time. (Why give development aid to a region that will be uninhabitable; why support facilities for desert life if it is going to become a garden?) She consults a team of meteorologists who run a complex global climate model on their supercomputer. They assure her that given the data available at the moment the model predicts extreme drought in the Sahara throughout the next century. So she expects drought. Climate models are tricky things though. They are constructed around simplifying assumptions.

designed to minimize false conclusions while bringing computational tractability. (All the computing power in the universe could not deduce weather from fluid dynamics; all the intelligence of all the conceivable human mathematicians could not get an all-purpose qualitative understanding of it.) Some assumptions are appropriate for some purposes, none for all: you wouldn't use a model designed to predict rainfall tomorrow to indicate whether there will be more hurricanes in a hundred years' time. And this fact is relevant to our politician's beliefs. There has been a very subtle miscommunication between her and the meteorologists, and as a result they have used a model whose assumptions are quite inappropriate to the question she is asking.

Something is wrong with the politician's belief that there will be a drought in the Sahara in twenty years time. But what? It is based on evidence in a reasonable way, using methods which she has good reason to believe are the appropriate ones. We could characterise them as reasonable or justified or warranted, but based on the false assumption, which she could not have known to be false, that the model is appropriate. Or we could characterise them as undermined, flawed, defective (in fact unjustified and unwarranted) since they are based on a method of reasoning that will not reliably give true or useable answers to the questions she is asking. (There are hard questions of model-tuning here. For my take on them see my 'Mathematical models: questions of trustworthiness', *British Journal for the Philosophy of Science*, 44 no 4, 1993, 659-674, 1993)

(Linguistic intuition, for what that's worth, suggests that her belief is at any rate not unreasonable. Simply because *she* is not unreasonable: she is going about things in a sensible way and is in full possession of her faculties.) So we can say useful folksy things but it is hard to sound philosophical.

(b) *Probability* Blaise is very innocent about probability and asks her statistician friend Pierre what the probability is that a fair coin will land heads five times in a row. Pierre answers in impressive but mysterious terms: "Since the coin is fair each toss is independent of the previous ones, so it is as if different coins were being tossed. So imagine a grid with five coins side by side and for each one its pattern of heads and tails. The number of possibilities is the area of that grid, length times width, so to get the probability of independent events you multiply their probabilities." Blaise nods her head, having absorbed the conclusion, thinks for a moment, and says "So the chance of a fair coin landing heads five times in a row is $2 \times 2 \times 2 \times 2 \times 2$, or 32."

Blaise has got the right conclusion by the right method. But the reason Pierre gave her for following that method is confused: if he were right then the answer would be not 32 but 25. (Not 2^n but n^2 .) Blaise can now solve similar problems about independent events, having been put on the right track by Pierre's nonsense. Is her belief justified? Surely we should be asking some more nuanced questions. She is a sensible person thinking in a

responsible manner, but we fear that she can be lured into error by wise-sounding pseudery.

(c) *(Over)confidence* A mathematician is working in a largely unexplored domain. He finds himself very short of intuitive holds on the material and the proofs that seem to work are too long and fiddly for him to be completely sure of them. One night as he is coming out of a movie theater a wonderful but very dubious idea occurs to him. ("Suppose that the meeting points of lines at infinity on the real projective plane have complex coordinates." Something of that kind of simple formal nuttiness.) This shows him a path through the forest. A few things become immediately clear to him, though he realises that he could have found and proved them quite independently of his path. Further things only hang together in terms of the path, and he formulates a fundamental theorem for which he manages to make a fairly convincing proof, some steps of which are clearly a bit dicey. Other mathematicians are attracted, special cases that do not need the wobbly bits of argument are proved, and a generation later a young genius proves a result that is notationally very similar to our mathematician's fundamental theorem, though some of the formalism has a significantly different interpretation. The theorem becomes known by the name of our hero.

The point here is confidence and conjecture, pushing at the limits of what the epistemic virtue of carefulness will permit. The example could as easily have been put as another climate modelling example, except here the person would have had to have got a satisfactory result by taking a risk with a model that was not clearly appropriate. Or it could have been put in terms of a situation where an inference to the best explanation is possible, but the best explanation is rather implausible or scary, so that there is an equally strong case for cautiously refusing to make any conjecture. Again it is not clear how to deploy the language of justification and rational belief. The puzzle can also be put in terms of the relation between epistemic prudence and rational belief. We realise now that we want our thinking not only to avoid falsehoods but to produce interesting truths and we accept that we only get the latter by taking risks. But there does not seem to be any degree of epistemic caution that attaches to reasonable belief. A kind of moral luck operates: there are interesting truths that we will only discover by taking leaps in the dark, and very often people who take such leaps will end up with mistaken conclusions that others will reasonably consider to be unreasonable. (It is sometimes right to take such leaps, for example in supposing the existence of a physical world, or of other minds.)

(d) *believing the impossible* You believe all of ten propositions. Someone presents you with an utterly convincing proof that nine of them are tautologies and one of them is a contradiction. But you don't know which. All these propositions are relevant to your everyday life, so it is vital that you adopt some attitude to them. You feel that you have no choice but to continue to believe them all, putting the realisation that one is contradictory

into a little compartment in your mind for future reference. You hope that some bright person will sort the problem out later. Meanwhile life goes on. (See Richard Foley, 'Being knowingly incoherent', *Nous* 26, 1992, 181 - 203.)

The strategy you have adopted, of continuing to hold all the threatened beliefs, is not the only defensible one. You could also have abandoned them all, or adopted an attitude to all of them that fell short of belief. Or you could have abandoned that one whose loss was least destructive in practical terms. Assume that the strategy of holding on to all of them is a reasonable option, though it is not the only thing you could defensibly have done. What should we say about the resulting beliefs, in particular the belief that is the conjunction of all of the ten? It does not seem right to say that they are all justified - though *you* may be practically justified in holding on to them - since you *know* that they cannot all be true. Inasmuch as you had grounds for believing them you now know that these grounds are flawed, though you cannot identify the flaws. But it also does not seem right to say that they are *unjustified*, or at all irrational. For in holding them you are operating in a perfectly reasonable way. The situation is tolerably clear, what is not clear is how to apply the traditional terms.

3. diagnosis These four examples turn on some inescapable features of human belief. On the most superficial level they depend on our inevitable dependence on testimony, on the use of models and heuristics, and on lines of argument whose implications we do not fully grasp. There are more basic reasons why these are inevitable.

One is the impossibility of the epistemically isolated agent. Most of what every person believes is derived from the beliefs of others, and nearly all of the social acquisition of belief works with a norm of trust, without which we would be lost, but which can be counted on to betray us regularly. (These ideas are pushed in the writings of Welbourne, Craig, and Plantinga.) Even in mathematics proof is a many-person activity: one person produces a draft proof and others check it, fill gaps, and express worries. To a large extent each one could do what each other is doing, but no one mortal could do it all.

Another basic reason is the stratification of reasoning. We sometimes start with assumptions and then think our way directly to a conclusion. We also often think about the kinds of reasoning that are likely to give good results in the circumstances, and then adopt one or another. We would be unable to use mathematical models, inferential short-cuts, and statistical cookbooks if we did not reason in this way. And sometimes, as in the philosophy of science, we operate at a third level, thinking about what kinds of reason-guiding thinking are going to select first order reasoning well.

(These two reasons, non-isolation and stratification interact with each other. We learn ways of thinking from one another. That's a function of epistemology. And we think about how to fit testimony and common belief into our patterns of reasoning.)

A third factor behind examples such as those in the last section is holism. In considering whether to add, subtract, or maintain a proposition as belief a person will take explicit account of some other beliefs. But many other beliefs, and in fact the overall pattern of the person's beliefs, form an essential background. Were that presupposed pattern different, the beliefs that are explicitly considered would not have the force that they do. They might seem ludicrously irrelevant. This is a familiar point of post-Quinian epistemology. It too generates situations in which it is hard to think of beliefs as simply reasonable (justified, supported by evidence) or not. Holism can be found in the examples above, but its relevance would emerge more clearly from different examples, in which, for example, people take for granted background beliefs which in fact are open to challenge though challenging them is not relevant to the enquiries they are pursuing. (Assumptions underlying racism or sexism are beginning to look wobbly, but you are a conventional, if reasonable, member of your conservative culture: that sort of thing.)

Holism, non-isolation, and stratification all interact. An individual person has a pretty tenuous hold on much of the background relevant to changes in their beliefs. Or, to put it differently, the overall pattern relevant to the epistemic status of one person's beliefs is largely found in the minds of others. (The conjunction of epistemological holism and semantic externalism is largely unexplored: but it won't be for long. The writings of Keith DeRose are the first of many.) The beliefs of other people are part of the background which determines whether the reasoning you are using now is appropriate.)

There is a third, more fundamental, explanation of the inevitable effect of holism, non-isolation, and stratification, and thus of the force of examples such as those of the last section: human finiteness. The difficulty of talking about finiteness lies in imagining its absence, so consider some possible idealisations of our situation.

A *Cartesian agent* can handle only finite amounts of data, but applies to it an unlimited reasoning power and allows itself an unlimited time to get through its reasoning. Cartesian agents start with the same kinds of evidence as non-idealised humans, but they are incomparably smarter and very very patient.

A *Humean agent* can handle only finite amounts of data, applies an unlimited reasoning power to it, but allows itself only a limited amount of time to complete its reasoning. Humean agents are as smart as Cartesian agents, but in more of a hurry.

Neither Cartesian nor Humean agents would find themselves in the situations I have been describing. Or, to put it more carefully, such situations would not be inevitable for them. Any one Cartesian or Humean agent could reproduce the beliefs of their whole community within their own mind, and as evidence accumulates this whole body of belief could evolve as that of one person. (Would this then evolve just as the beliefs of a community might? Hard question.) Cartesian or Humean agents would not need cognitive shortcuts: given a theory they could deduce its consequences directly without mediating models: given a problem in logic or statistics they would

think it out from first principles, using no rules of thumb. And though such agents might change their beliefs in terms of patterns of global coherence - Cartesian doesn't have to mean foundationalist here - they would not need to manage this in terms of foreground and background. In considering a large body of beliefs such an agent could think of it as a number of individual beliefs, all to be considered as such. The web of belief would be a scatter of individual dots.

Cartesian and Humean agents seem to us now impossibly remote from our situation. They don't seem like the appropriate idealisations for measuring ourselves against. We think of ourselves as what I'll call *Simonian agents*: limited reasoning power, little time to spend on any task, and limited data. We have been pushed from Cartesian to Simonian agenthood as our conception of ourselves by some profound and general realisations about the world and our relation to it. Three realisations seem to me particularly important.

We have learned that simple quasi-deductive reasoning will not do justice to the patterns of evidence needed to discriminate true hypotheses. Neither Humean induction nor Mill's methods nor anything like them will do. We need real statistics, and real statistics are inevitably really complicated.

We have learned that no belief is so secure that more evidence cannot overturn it. And this applies to true beliefs too, in fact to beliefs that are knowledge. You can know that p and then more evidence or further thinking can force you, if you are rational, to abandon what was in fact knowledge. Reason gives and reason takes away, even knowledge.

We have learned that reasoning about reasoning is often harder than simple reasoning. This is clearest in the case of deductive logic, where following deductions from premises is one thing, but knowing whether a given consequence follows from given premises is much more demanding. So there are crucial problems of fitting inference to purposes - "is following this line of reasoning going to answer this question?" - which are not going to be solved by anything as simple as churning out deductions. There's no alternative to short-cuts, rules of thumb, inspired guesses about which inferential strategy is likely to pay off. (These issues surface clearly in some writings of Timothy Williamson, see for example 'Inexact Knowledge', *Mind* 101, no 402, 1992, 217-242.)

The point is Simonianhood rather than the discoveries that have forced us to it. But in the light of them it is clear that getting into pickles like those of the people in the cases of the previous section is not a sign of stupidity, carelessness, or defective functioning. They are situations that anything like a human being is going to encounter regularly. Or to put it differently, since the days of Descartes we have not come to think that people are stupider or less resolute. If anything we have come to appreciate quite how complex the thinking we can produce and follow is: the problem is that the world is so complex that we *need* this complicated thinking, so complicated that while we can produce it we cannot master it. And we should not take "thinking" here

to be just inference. It includes the complexity of, for example, the social organization of science too.

4. Virtues of limitation Cartesian or Humean agenthood is not an appropriate ideal because aspiring to it does not bring good results. That is, in many epistemic situations what in section 1 I called the B-strategy, of performing actions which are appropriate for an imperfect agent, rather than approximating to perfection, works best. That is not to say that all situations are like that. Sometimes the best advice is "try as hard as you can, ignore your limitations". Which is which?

A fairly broad-brush approach to the question would distinguish a number of wide types of situation in which we assess a person's beliefs. It is not at all obvious what the best such taxonomy is, but here is a not very profound suggestion. Distinguish between the following questions about a person's belief:

Evidence: Is the belief ok with respect to other beliefs that could reasonably be counted as providing evidence?

Acquisition: Is this a belief that it is ok to acquire in the person's circumstances?

Change: Is it ok to hold on to this belief given that specific other beliefs have changed?

Globality: Is it ok to hold on to this belief given the (long-term) changes in the person's overall system of beliefs that it would be reasonable to make? ("Ok", "reasonable" are deliberately vague: note that they don't indicate epistemic obligation but epistemic permission, though.)

These questions differ along several dimensions. Acquisition versus maintenance versus defence, short term versus long term, local versus global. There is no reason to believe that the kinds of capacities and virtues of finite agents that are relevant to one of them are relevant to others. (Every reason to believe the opposite.) In particular, one can see in a rough general way how they differ in how suitable they are to the A or the B strategy.

Acquisition for example is a B-strategy topic: Simonian agents do not have time or resources to make global changes to their beliefs with every passing perception. Some jobs have to be deferred. And *Globality* is also a B-strategy topic: the problems of maximizing global coherence among a large number of beliefs are so complex that the only feasible approach is piece-by-piece, accepting temporary conflicts and anomalies. The other topics are harder to assess.

There is a form of *Evidence* that is arguably an A-strategy topic. Given specific items of evidence and a specific target belief, how strong is it supported by that evidence alone? This is an extremely artificial question, but inasmuch as we can make sense of it, the only admissible answers seem to be in terms of an objective relation of evidential support. Nothing in this paper is inconsistent with there being such a relation, or indeed with an orthodox Bayesian construal of it. (That assumption does nothing to defuse the cases

in section 2.) And if you are asked to say how strongly E_1, \dots, E_n supports H then what you have to do is to figure it out, trying as hard as you can to approximate to the correct answer that a less limited agent would find.

A very subtle question now arises. Usually when we consider how well evidenced a belief is we are not relating it to a fixed set of evidence propositions. We are relating it to The Evidence, as available to the believer. Then the believer's task, and questions we can ask about her accomplishment, are more complex/different: has she found the relevant evidence? how much of it? has she summarised it in a useful form? has she isolated the individual items that should be considered individually? And then, building on all of these: has she used a suitable method of thinking out the relation of the belief in question to this summary and these individually salient items?

Here's the subtlety: The suitable method for a given person to relate a belief to summarised and underlined evidence is rarely one that an ideal agent would have used, so the choice of such a method is not an A-strategy task. But once the method is chosen carrying it out is often best done with no short-cuts or approximations: so that's a B-strategy task. You set up a job fit for a Simonian agent, and then carry it out with Cartesian industriousness and patience.

This is a pattern that occurs all over epistemology, I think: A-strategies nested inside B-strategies. The pattern would be explained if a general principle were true, which is independently plausible (but very general).

Belief-forming tasks that involve global relations between indefinite numbers of beliefs are B-type tasks, best evaluated by norms that accommodate human limitations. Tasks that involve relations between small numbers of particular beliefs are A-type tasks, best evaluated by norms of perfect performance.

Why should we think that this principle is true? Well, it fits a number of particular cases. It also fits with the principle formulated in section one above, that maximization and approximation do not mix, since questions of global coherence are nearly always questions about how to maximize logical and explanatory coherence. Suppose that it is true, though, and we have a problem for the traditional vocabulary. Different tasks, nested within one another, are subject to norms of very different character. We cannot ask "is this person justified in believing this?" Instead we have to focus on particular tasks - collecting evidence, immediate reaction to it, summarising it, sorting out a whole system of beliefs - and ask whether the right procedures for those tasks have been found, and then how they have been applied. There are many evaluative questions we can ask of the different components of this process, and their answers rarely line up in any simple way.

5. If you've got this far:

HAPPY BIRTHDAY PETER