

## **In Defence of a Dogmatist**

Brian Weatherson, March 20 2006  
incomplete paper, just sections 1-5 so far

The surge of attention paid to dogmatism in recent years becomes perfectly understandable once we set out sceptical reasoning carefully. For dogmatism is just the denial of what may be the weakest premise is a careful formulation of the argument from sceptical hypotheses. I'll briefly set out the sceptical argument in a way that makes the role of dogmatism explicit, then survey recent challenges to dogmatism. Many of these challenges rely, either explicitly or implicitly, on some feature of a broadly Cartesian approach to epistemology, and I'll argue that when we make those assumptions explicit the challenges look surmountable.

### *1. External World Scepticism*

A standard way to put the classic argument from sceptical hypothesis for scepticism is the following.

1. If I know I have hands, then I'm in a position to know that I'm not a handless brain in a perfectly functioning vat (BIPV)
2. I'm not in a position to know that I'm not a handless BIPV
3. So I don't know I have hands

That's basically the argument that Keith DeRose gives in [Responding to Skepticism](#), though (a) he uses variables O and H rather than a specific example, and (b) doesn't include 'position to know' in the consequent of 1 and in 2. But it's the same basic idea. DeRose seems to present the sceptic as taking 2 to be an intuitively plausible claim, and that intuition is at the heart of her argument. (See his discussion of intuitions in footnote 1 for example.) But I think we can make the sceptical case stronger if we set out an argument for premise 2. In [Scepticism, Rationalism and Externalism](#) (SRE), I suggested that the argument for 2 is the following. (In putting this the argument this way I'm following some suggestions Jim Pryor makes in [The Sceptic and the Dogmatist](#).)

1. I'm not in a position to know a priori that I'm not a handless BIPV
2. I'm not in a position to know a posteriori that I'm not a handless BIPV
3. So, I'm not in a position to know that I'm not a handless BIPV

If we take a posteriori knowledge to just be knowledge that's not a priori, this argument is valid. And I think the intuitions that the two premises of this argument are true is stronger than the

direct intuition that its conclusion is true. But I think on reflection that we can give an even more detailed argument on behalf of the sceptic, because we can argue for each of the two premises in this argument.

First the claim that I can't know a priori that I'm not a BIPV.

1. If it could have turned out that I'm a handless BIPV, then I'm not in a position to know a priori that I'm not a handless BIPV
2. It could have turned out that I'm a handless BIPV
3. I'm not in a position to know a priori that I'm not a handless BIPV

More intuitively, if I could have turned out to be a handless BIPV, then that's something I need evidence to rule out. That is, it's not something I can know to be false a priori. So I can't know that I'm not a handless BIPV a priori.

Now for the a posteriori premise.

1. It could have turned out that I'm a handless BIPV
2. If it could have turned out that I'm a handless BIPV, then it could have turned out that I'm a handless BIPV with the same empirical evidence as my actual empirical evidence
3. If I'm in a position to know a posteriori that I'm not a handless BIPV, then it could not have turned out that I'm a handless BIPV with the same empirical evidence as my actual empirical evidence
4. So I'm not in a position to know a posteriori that I'm not a handless BIPV

This argument is more complicated, but the idea is that if handless BIPVs are possible, then so are handless BIPVs that are evidentially just like me. But my evidence can't distinguish between possible ways I could be that have the same evidence. So I can't know on the basis of that evidence, i.e. a posteriori, that I'm not a BIPV.

Putting all of this together, we get the following long argument for scepticism.

- E1. If I know I have hands, then I'm in a position to know that I'm not a handless brain in a perfectly functioning vat (BIPV)
- E2. If it could have turned out that I'm a handless BIPV, then I'm not in a position to know a priori that I'm not a handless BIPV
- E3. It could have turned out that I'm a handless BIPV
- E4. If it could have turned out that I'm a handless BIPV, then it could have turned out that I'm a handless BIPV with the same empirical evidence as my actual empirical evidence
- E5. If it could have turned out that I'm a handless BIPV with the same empirical evidence as my actual evidence, then I'm not in a position to know a posteriori that I'm not a handless BIPV
- E6. So I don't know that I have hands

I've put 'E' in front of each premise because this is my form of the canonical argument for external world scepticism, and I'll be contrasting it with arguments for other kinds of scepticism below.

We could break this argument down even further. Timothy Williamson for instance thinks that the sceptic has an argument from the luminosity of evidence for premise E4. But for now I'll stop here because the argument as stated has a rather nice property. For every one of the premises, there are many philosophers who accept it, and many philosophers who reject it. And of course there are philosophers who accept the conclusion. So we have a six way breakdown of epistemologists. (The breakdown isn't *exclusive*, because someone could reject multiple premises in this argument.)

Premise E1 is rejected by philosophers such as Dretske who reject closure. Call these the anti-closure theorists.

Premise E2 is rejected by philosophers who I call *rationalists* in SRE. I'll talk a lot more about rationalism in the next section.

Premise E3 is rejected by at least some theorists who endorse the so-called Response from Semantic Externalism. The name is a little misleading, because as David Chalmers argues in [The Matrix as Metaphysics](#), there is a similar argument against the possibility of *some* sceptical hypotheses that does not have any externalist assumptions.<sup>1</sup> Because premise E3 is used in both of the sceptic's sub-arguments, theorists who deny E3 might differ amongst themselves over whether we know a priori or a posteriori that various sceptical scenarios do not obtain. It is more in the spirit of semantic externalism to say that this knowledge is a posteriori, but perhaps Chalmers' argument suggests that in some cases this knowledge is a priori.

Premise E4 is rejected by Evidence Externalists, such as Timothy Williamson. Williamson thinks that although BIPVs are possible, they have different evidence to what we do.

---

<sup>1</sup> As David Chalmers pointed out to me in correspondence, on his view it *could* have turned out that we are *recently envatted* brains in vats, so his arguments don't give us a reason to deny E3 as written.

That's because for Williamson our evidence is what we know, and we know more things than a BIPV.

And I think what is distinctive about dogmatism is that it is the rejection of premise E5. The dogmatist thinks that even if our evidence is consistent with our being a BIPV, we can have genuine a posteriori knowledge on the basis of that evidence that we are not a BIPV. In particular, the dogmatist thinks that we can take our appearances at face value without having prior, and in particular without having a priori, reason to believe that in doing so we go wrong, as in a BIPV scenario. Many dogmatists say something more precise than  $\neg E5$ , but this denial is central to the view, central to its response to scepticism, and I think most of the attacks on dogmatism are basically arguments for E5.

Finally, the sceptic is the theorist who thinks that the argument has a true conclusion.

To my mind, the primary attraction of dogmatism comes from the implausibility of the sceptical conclusion, combined with the attractiveness of the other premises. Now in the case of external world scepticism, this kind of consideration isn't too compelling because the other premises are not that attractive. In particular, both E3 and E4 have a lot to be said in their favour. But external world scepticism is not the only kind of sceptical argument there is, and when we look at other sceptical arguments, the case for dogmatism gets stronger.

## 2. Inductive Scepticism

As Roger White notes in [A New Look at the Old Problem of Induction](#), we can imagine a dogmatist position about inductive justification. White notes this as part of an attack on dogmatism, but I think the point about dogmatism's potential applicability is independent of these attacks. To see this, consider the following sceptical argument. In this argument,  $p$  is some proposition about the future, for example the proposition that the next time I watch *Casablanca* Sam will twice play "As Time Goes By", and  $h_t$  is the proposition that the history of the world to  $t$  is exactly like *this*, where 'this' picks out the history of the actual world to the present moment.

- I1. If I know  $p$ , then I'm in a position to know  $\neg(h_t \wedge \neg p)$ .
- I2. If it could have turned out that  $h_t \wedge \neg p$ , then I'm not in a position to know a priori  $\neg(h_t \wedge \neg p)$ .
- I3. It could have turned out that  $h_t \wedge \neg p$ .
- I4. If it could have turned out that  $h_t \wedge \neg p$ , then it could have turned out that  $h_t \wedge \neg p$  and that I have the same empirical evidence as I actually have.
- I5. If it could have turned out that  $h_t \wedge \neg p$  and I have the same empirical evidence as my actual empirical evidence, then I'm not in a position to know a posteriori that  $\neg(h_t \wedge \neg p)$ .
- I6. So I don't know that  $p$

Although E3 and E4 are plausibly denied, I don't think that carries over to I3 and I4. There is no argument from semantic externalism against the existence of a possibility where an evil demon changes my *Casablanca* DVD so Sam plays "Take Me Out" to Rick instead of "As Time Goes By". If this is to be a world where  $h_t \wedge \neg p$ , it has to be a world where the evil demon only comes into existence after the present moment, but again there is no argument from semantic externalism against that possibility. So I think I3 is fairly secure.

That I3 is safe should not be controversial, for it has never been claimed that semantic externalism is a response to *all* sceptical arguments. That I4 is safe is a little more controversial, for it like E4 is in conflict with Timothy Williamson's position that evidence is knowledge. Now running a full argument against Williamson's theory of evidence would take us too far afield, so instead I'll just note here that the hole Williamson notes in the standard argument for E4 is not a problem in the argument for I4. Williamson thinks that the standard argument for E4, used by both sceptics and internalists about evidence, appeals to the luminosity of evidence. Since he thinks, rightly, that evidence need not be luminous, this undermines the case for E4. It isn't clear whether there is any way to fix this hole in the argument for E4, and in any case it does seem plausible that our evidence might consist of what we perceive as well as the way that we feel. But there is no similar hole in the argument for I4, and in fact we can give a quite plausible argument for I4. This argument just relies on a causal theory of evidence. It is very plausible that for *my* evidence to change, there must be some causal impact on me. Changes to the world from which I am causally isolated intuitively make no difference to my evidence. It is perhaps possible that something that makes no causal difference to me nevertheless makes a difference to the *strength* of my evidence because it changes which facts in the world my evidence is correlated with, and we have a correlationist account of evidential support. But the *identity* of my evidence does seem to depend solely on my causal history. And since my causal history is the same in the actual world as in the DVD demon world, I4 at least seems to be true.

I'm not going to give a new non-question-begging argument for either of these claims, but I think the options of either rejecting I1 or accepting I6 are not viable. I find the arguments John Hawthorne makes in [In Defence of Closure](#) to be compelling support for I1. And I6 clearly violates our standard attributive practices. I was going to illustrate inductive scepticism with an example concerning the end of *Casablanca*, but I decided against that because I thought that some readers might not have seen it, and it is more fun to see movies without knowing how they end. This is a very commonly used reason for not telling people how movies have ended on previous screenings, and it is clearly a *good* reason but if I don't know how *Casablanca* will end next time you watch it, it would be a bad reason. So I think accepting the conclusion is out of the question.

That leaves us with the rationalist option of rejecting I2, or the dogmatist option of rejecting I5. Now what is interesting about both rationalism and dogmatism is that the primary objections to them are independent of their applications. So the reasons usually given against the

dogmatist response to external world scepticism, i.e. in favour of E5, are equally strong arguments in support of I5. (We'll see this in detail when we look at these arguments in the bulk of the paper.) And the primary argument for E2, the alleged implausibility of knowing a priori that a certain possibility is not realised, is an equally strong argument for I2. So we can reasonably say that  $E2 \leftrightarrow I2$  and  $E5 \leftrightarrow I5$ . But we also know that either I2 or I5 have to be true. So we know that one of E2 and E5 are true. So if rationalism is not true, then dogmatism, both about induction and the external world, is. Since many people object strongly to rationalism, this is a reason to take dogmatism very seriously indeed.

### 3. *Dialectical Ruminations*

Despite writing this paper in defence of dogmatism, my personal inclination is not to believe it. By inclination I'm a rationalist who thinks that we have fallible, but incorrigible, a priori knowledge that many sceptical hypotheses are false. But an inclination is not an argument, and I don't have any compelling *arguments* against dogmatism. If I'm right in what follows, then nor at this stage of the development of epistemology does anyone else.

Much of the recent discussion of dogmatism has been connected to evaluation of Moore's famous proof of the external world. From the perspective I'm advocating here, the issues about that argument are somewhat orthogonal to the evaluation of dogmatism as such. The core dogmatist position is that it is possible for appearances to justify beliefs about the induction or the external world without a prior justified belief that there is a connection between those appearances and external facts. Hence it isn't particularly relevant to the case of G. E. Moore, who *did* have a prior justified belief that his appearances were a good guide to reality.

Put another way, dogmatism is most interesting as a theory about justification in the context of discovery, in particular in the context of discovering that various sceptical hypotheses do not obtain. Whatever Moore did do that day at the Aristotelian Society, he didn't *discover* that he had hands.

Now I know that this is all completely obvious, and the story about Moore is just meant as a dramatic device. But it's worth restating because, as I'll argue below, it is a quite misleading dramatic device. Some of the things that are suitably taken for granted in the context of looking at Moore and his hands (or fake hands) are not suitably taken for granted a priori, or at least before we discover what kind of world we live in. If some of the arguments I'm going to make below are right, dogmatism relies on their being something of a disanalogy between our initial discovery of what kind of world we live in, and our discovery of particular facts about that world, such as the fact that G. E. Moore has not lost his hands.

As I said at the top, I think many of the critics of dogmatism have taken for granted a broadly Cartesian approach to epistemology. The Cartesian approach, as I'm using the term 'Cartesian' at least, is based around these assumptions:

- Evidence supervenes on phenomenal appearances;
- Justification supervenes on evidence;
- Appearances are luminous, so that when it appears to us that  $X$ , we are justified in believing it appears to us that  $X$ , and when it does not appear to us that  $X$ , we are justified in believing it does not; and
- We are justified in believing that  $p$  only if we could justify belief in  $p$  starting from a position of radical doubt.

I'm going to concede eventually that given a Cartesian outlook, strong objections to dogmatism can be made. But the Cartesian outlook is not compulsory, and without it some of the objections to dogmatism are undermotivated.

#### *4. Bootstrapping, Dominance and Luminosity*

Several recent attacks on dogmatism seem to require an assumption that appearances are luminous in Timothy Williamson's sense. Indeed, they occasionally require that it be a priori that appearances are luminous. A dogmatist who, quite reasonably, rejects these assumptions should not feel threatened by such arguments.

The first such argument is the 'bootstrapping' argument due to Jonathan Vogel (2000) and Stewart Cohen (2002).

Imagine my 7 year old son asking me if my color-vision is reliable. I say, "Let's check it out." I set up a slide show in which the screen will change colors every few seconds. I observe, "That screen is red and I believe that it is red. Got it right that time. Now it's blue and, *look at that*, I believe it's blue. Two for two ..." I trust that no one thinks that whereas I previously did not have any evidence for the reliability of my color vision, I am now actually acquiring evidence for the reliability of my color vision. But if Reliabilism were true, that's exactly what my situation would be. (Cohen 2002)

Although Cohen only mentions reliabilism in this quote, and reliabilism is clearly Vogel's target as well, Cohen makes clear that the target is more generally any theory that allows that "S knows  $p$  prior to knowing that the cognitive source of S's knowing  $p$  is reliable." (Cohen 2005: 417). And dogmatist theories certainly fit into that category.

The first thing to note about this argument is that it isn't clear how we're supposed to get from the example to any actual evidence of reliability. Compare the following simple example. I stand in Harvard Yard, asking people who go by whether (a) they attended the Dalton school and (b) they have attended or now attend Harvard. I ignore the people who answer 'no' to the first question, or who ignore me, and note what percentage of people who answer 'yes' to the first

question answer ‘yes’ to the second. Let’s say, not surprisingly, this is 100%, and that I have a sample size of, say, 22. Do I have reason to infer that all, or even most, Dalton students go on to attend Harvard?

Surely not. For one thing, even though I could find non-Harvardians in Harvard Yard (for example I’ve occasionally been there), so my experiment isn’t guaranteed to return this result, it isn’t in any way a random sample. Nor have I said anything else about my other evidence. I might, for all I’ve told you in the example, have evidence of 100 other Dalton graduates who did not go on to Harvard. And, as will become important below, I might not have reason to trust the testimony that people give in answering my survey questions.

In explicitly applying the bootstrapping objection to dogmatism, Roger White in [Problems for Dogmatism](#) presents a version of the argument that avoids some of these pitfalls. He claims that if I do something like Cohen’s slide show, I will be able to form the following belief.

**Track Record:** This card appears to be C in color, and it is C, this one appears C’ and it is C’,..., and these are all the cards that I have viewed.

**No Errors:** My color experiences have matched the actual color of each of the many cards that I have viewed.

The dogmatist is allegedly committed the agent justifiably believing Track Record because they think that even without prior justification the agent can properly take the appearances at face value. And Track Record entails No Errors, so by closure the agent is justified in believed No Errors. (Note here that White is careful to put into Track Record the claim that I have no other relevant evidence. That avoids the problem mentioned above that I might have done many other surveys before this one, and ensures that No Errors is entailed.) And the best explanation of No Errors is that the agent is a reliable colour judger.

We could worry about the final step because there aren’t a great number of alternative explanations floated. But let that pass for the argument for Track Record involves two false assumptions.

If a card looks C in colour, and there are no reasons not to trust the appearance, then the dogmatist is committed to saying the agent is justified in believing the card is C. But it doesn’t yet follow that they are justified in believing the card looks C in colour. That only follows from the facts on the table if we assume, for no good reason, that if a card looks C then the agent is justified in believing it looks C. (Similarly in the survey example above, I might not be justified in believing the survey subjects went to Dalton, even if they say this.)

Even if we make the assumption that if a card looks C the agent is justified in believing it looks C, the crucial final conjunct of Track Record doesn’t follow. That is, even given this assumption it doesn’t follow that an agent is justified in believing that they know of all the card



appearances they have received. It is possible that they haven't noticed a card that actually looked C to them (and may or may not have been C), and there's no good reason either a priori or in dogmatism to simply rule out this possibility.

So without some fairly strong luminosity principles, the agent is not justified in believing anything about how the cards looked, and certainly isn't justified in believing any proposition that purports to be a complete record of their appearances. So dogmatism without luminosity doesn't support Track Record, and hence doesn't support No Errors, and doesn't give the agent a reason to believe in their own reliability.

The luminosity principle needed here is actually quite strong. The assumption is that if the agent is appeared to C-ly, then (a) she has sufficient justification to believe that she is appeared to C-ly and (b) she does believe that she is appeared to C-ly. We need (a) to say that she is justified in believing the conjuncts of Track Record recording individual cases. And we need (b) to say that she is in a position to judge that these are all of the cases. If the dogmatist has a reason to deny either (a) or (b), then they have a reason to reject this argument.

Even if this luminosity principle is actually true, it is arguable that we need an a priori luminosity principle to generate a problem for the dogmatist. When pressing the case against the dogmatist, White says that given the dogmatist's assumptions "surely I shouldn't have to waste my time performing the test. I could simply reason that no matter how the test turns out, I will be justified in believing that my color-vision has been accurate on each use, which in turn confirms that my color vision is reliable." But if the anti-dogmatist argument requires luminosity, and we can only know a posteriori that appearances (and non-appearances) are luminous, then we can't know a priori that running the test will give us a reason to believe that our colour vision is reliable. So there must be some other reason to say that the bootstrapping argument is bad. I won't press this point because it might be argued, fairly, that it is simply intuitive that bootstrapping arguments are bad, independent of whether we can explain this in terms of the uselessness of running tests.

The big problem here is not that luminosity is assumed. It isn't even that there are arguments against luminosity, as in Williamson (2000). It might be that these arguments fail (as I argue in Weatherson 2004 they do fail for some cases) and that luminosity is true. The big problem is that it is incoherent to accept luminosity and denounce bootstrapping arguments. We can see this by looking at what we can learn about the reliability of introspective judgments given luminosity. (The discussion that follows is indebted in various ways to the discussion of introspection and priming in Pryor 2004: 360ff.)

A priori I can't know whether my introspective judgments will be reliable or unreliable. But I might want to know which they are. So imagine I test myself the following way. I show myself a card of colour C, and it appears C-ly to me. Given the luminosity of appearances, I can justifiably believe that I believe it appears C-ly and that it does appear C-ly. Then I show myself a card of colour C', and it appears C'-ly to me. Again by appeal to the luminosity of appearances, I

can justifiably believe that I believe it appears C'-ly and that it does appear C'-ly. I keep doing this a lot of times until I have a wide and diverse sample. At this point I have a lot of evidence that my introspective judgments are accurate, and no evidence pointing the other way, so I infer that I'm introspectively reliable.

This argument seems to me no better or worse than the argument from dogmatism (plus luminosity) to the claim that my colour perception is reliable. At the very least, I can't see any virtue in the above argument that isn't shared in the perceptual bootstrapping argument that the dogmatist is (allegedly) committed to endorsing. So if the perceptual bootstrapping argument is bad, then the introspective bootstrapping argument is bad too, and vice versa. We'll say for short that bootstrapping is bad if both these arguments are bad, and bootstrapping is good if both these arguments are good. It turns out not to matter for anything that follows whether bootstrapping is bad or good. (For what it's worth, I can't see much wrong with the introspective bootstrapping argument, and that makes me much more sympathetic to the perceptual bootstrapping argument.) What does matter is that if luminosity is true, then the argument from the previous paragraph, i.e. the introspective bootstrapping argument, is a good argument. So if luminosity is true, then bootstrapping is good.

Now to get a bootstrapping argument against dogmatism, we need to assume both that bootstrapping is bad, and that appearances are luminous. But as we just showed, if appearances are luminous, then bootstrapping is good. So it is incoherent to assume that appearances are luminous and the bootstrapping is bad. But that's exactly what is assumed in the argument from bootstrapping against dogmatism.

We might wonder how much of a bootstrapping argument we can regenerate without luminosity. The anti-dogmatist might note that there will be *some* cases where we can justifiably believe that things appear C-ly, and they are C-ly, and *no* cases where we can justifiably believe that things appear C-ly, and they are not C-ly. This might seem to support a dubious inductive argument to the claim that we are reliable. To see that this argument fails, we just have to note White's careful formulation of the argument. To support an inductive claim that all (or even most) Fs are Gs, we don't need just a list of Fs that are Gs, but some reason to believe that (a) these are all of our sample and (b) our sample is a representative sample. And without a luminosity-type claim we can't get that, because the agent can't rule out that they have had lots of appearances about which they have no beliefs, justified or otherwise. Moreover, since it is at least possible, and even plausible, that appearances that we know we are having, clear cases of appearance types, are more reliable than appearances we don't know we are having, we have no reason to think these cases we know about are in the relevant sense a representative sample. So without luminosity there is no bootstrapping argument against dogmatism. But as we just noted, accepting luminosity is inconsistent with rejecting bootstrapping arguments.

White also discusses a more interesting argument against dogmatism to the effect that the dogmatist denies that we can justifiably believe some things that we will justifiably believe come

what may. He argues that this is implausible. I agree; indeed I use a similar argument in Scepticism, Rationalism and Externalism. The only problem is that I don't think the dogmatist is committed to this conclusion. White's argument concerns the possibility that G. E. Moore has 'super-fake hands' where a super-fake hand is a fake hand that "eliminates all evidence of super-fake hands, [and] of all other related skeptical hypotheses too." White imagines that we try and figure out whether G. E. Moore has super-fake hands by looking. He says there are three possibilities.

A: He will not appear to have hands.

B: He will appear to have hands, but we will have some reason to suspect that this appearance is deceptive.

C: He will appear to have hands, and we will have no reason to suspect that this appearance is deceptive.

He then wants to argue that in any of the three cases we'll be justified in believing that Moore does not have super-fake hands if dogmatism is true. The crucial case for us is A.

Now suppose A obtains. We will clearly be justified in believing that he is not super-fake handed, as he doesn't even appear to have a hand.

This is too quick. All that was specified in case A was that Moore doesn't appear to have hands. For us to justifiably believe that he is not super-fake handed, it would be sufficient for us to justifiably believe that he doesn't appear to have hands. But this doesn't follow from the fact that Moore doesn't appear to have hands. Or at least it doesn't follow without a fairly strong luminosity assumption, that we are justified in believing that things do not appear the way they do not in fact appear, and the dogmatist has no reason to grant that assumption.

Now in practice we generally are justified in believing that our belief about whether or not someone appears to have a hand is justified. But remember that we don't particularly care about Moore. What matters at the end of the day is the sceptical argument. We certainly don't know in advance of getting any empirical evidence whatsoever that we will be justified in having beliefs about how things appear, let alone that those beliefs will be accurate. So the gap between situations like A obtaining and our justifiably believing in advance that we will justifiably believe they obtain is particularly large when we consider dogmatism as a response to scepticism.

As I argue in SRE, this *kind* of argument might go through with a somewhat weaker luminosity assumption, perhaps even a luminosity assumption that we can justify a priori. So I think the kind of argument White is running here is of some interest. And he is clearer than I am about the need for principles like the Meta-Justification Principle in this kind of dominance

argument. But there is no hope for such an argument without a luminosity assumption, and it doesn't harm dogmatists if they have to deny luminosity.

### 5. The Bayesian Objection

One of the most widespread objections to dogmatism is the following Bayesian objection. (It appears in various forms in work by John Hawthorne, Roger White and Stewart Cohen.) Let A be the proposition *it appears to me that there is a hand*, H be the proposition *there is a hand* and F be the proposition that it falsely appears to me that there is a hand, i.e.  $A \wedge \neg H$ . Then we can make the following argument, where Pr measures our *prior* probability

- |  |  |
|--|--|
| 1. $\Pr(A) < 1$                                | premise                                    |
| 2. $\Pr(F) > 0$                                | premise                                    |
| 3. $\Pr(F   A) \cdot \Pr(A) = \Pr(F \wedge A)$ | from definition of conditional probability |
| 4. $\Pr(F \wedge A) = \Pr(F)$                  | from definition of F                       |
| 5. $\Pr(F   A) \cdot \Pr(A) = \Pr(F)$          | from 3, 4                                  |
| 6. $\Pr(F   A) > \Pr(F)$                       | from 1, 2, 5                               |
| 7. $\Pr(\neg F   A) = 1 - \Pr(F   A)$          | theorem of probability                     |
| 8. $\Pr(\neg F) = 1 - \Pr(F)$                  | theorem of probability                     |
| 9. $\Pr(\neg F   A) < \Pr(\neg F)$             | from 6, 7, 8                               |

Premise 1 is rather plausible, since it isn't an a priori given that we'll see hands. And presumably we shouldn't rule out probabilistically that we'll see fake hands, so 2 is plausible. And the rest follows fairly trivially using grade school algebra. The gloss usually put on step 9 is that getting evidence A lowers the probability of  $\neg F$ , so it can't be evidence for F. But the dogmatist says we go from not being justified in believing F before we get A, to justified in believing F when we get A, which is impossible if A isn't evidence for F. I'll come back to the gloss in a second, because it seems very misleading to me, but first I want to clear up what seems to be an equivocation in the argument.

Imagine a soul whose a priori credences are given by the following table. The columns V and S represent worlds where her primary perceptual apparatus is Vision and Sonar respectively. For some reason she is certain a priori that she'll have exactly one of those two perceptual systems. And she is certain she won't have hands without seeming to, so the only possibilities she takes seriously are the six in the table, and the prior probability for each of the six is given in the corresponding cell.

	V	S
$\neg H$	0.1	0.1
$A \wedge H$	0.3	0.1
$A \wedge \neg H$	0.1	0.3

For this agent,  $\Pr(A \wedge \neg H)$  is 0.4. And  $\Pr(A \wedge \neg H \mid A)$  is 0.5, so that probability is higher, as we showed above that in general it must be. But the agent's evidence is never just  $A$ . There is no such thing as a mere hand appearance. She either gets a hand appearance through vision or a hand appearance through sonar. As it turns out she gets a visual appearance of a hand. So she should conditionalise on  $V \wedge A$ . And  $\Pr(A \wedge \neg H \mid V \wedge A) = 0.25$ , i.e. lower than her prior probability for  $A \wedge \neg H$ . In this case conditionalising on appearances led to her probability in  $\neg F$  going up.

Now the anti-dogmatist may say at this point that the serious sceptical hypothesis  $F$  is not  $A \wedge \neg H$ , but  $V \wedge A \wedge \neg H$ . And the probability of that does go up, from 0.1 to 0.25, when we conditionalise on  $V \wedge A$ . If we define  $F$  to be maximally precise about the kind of appearances we're getting, then we will be immune to the objection being made here, which basically amounts to the claim that it's possible the kind of hand appearance we're getting is one that is more probably an accurate hand appearance. But the problem is that if we define  $F$  this way, then  $\Pr(F)$  will be 0. For there are, I'd imagine, uncountably many different kinds of perceptual systems we could, logically, have had. We're all used to vision of course, but there are many logically possible alternatives. The prior probability of having any one of them is 0.

So there is a dilemma facing proponents of this argument. If  $A$  is defined generically, to just be the proposition that we have some hand appearance or other, then it is possible that the probability of  $A \wedge \neg H$  goes down when we conditionalise on the actual hand appearance we get. If  $A$  is defined precisely, to be for example a *visual* appearance of a hand, then  $\Pr(F) = 0$ , and premise 2 in the argument is false. So I'm not sure that there is a way to define the key terms to make this argument work.

But let's assume that this problem can be solved, and  $A$  defined so that the argument goes through. There is still, as always, a big step between the mathematical argument above and any philosophical conclusion. In particular, there is a big gap between the *static* claim in step 9 about the relationship between two prior probabilities, and the *dynamic* gloss on it to the effect that when evidence  $A$  is received, the probability of  $\neg F$  goes down. This is only a plausible gloss if conditionalisation is the only plausible model for belief updating, and it is not.

There is a bad argument around here that the dogmatist might make. It might be argued that since the Bayesian approach (including conditionalisation) involves so much idealisation it could not be applicable to real agents. That's a bad argument because the Bayesian approach might provide us with a good model for real agents, and models can be useful without being scale models. As long as the Bayesian model is the most appropriate model in the circumstances, then

we can draw conclusions for the real world from facts about the model. The problem arises if there are alternative models which seem to fit just as well, but in which principles like *getting evidence A always lowers the probability of  $\neg F$*  are not true. Call that the anti-dogmatist principle. If there are alternative models that seem better suited (or at least just as well suited) to modelling the situation of initial evidence acquisition, and those models do not make the anti-dogmatist principle true, then we might think the derivation of that principle in the Bayesian model is a mere consequence of the arbitrary choice of model. In the rest of this section I will develop just such a model. I won't argue that it is the best model, let alone the only alternative to the Bayesian model. But I will argue that it is as good for these purposes as the Bayesian model, and it does not have the anti-dogmatist dynamic consequences of the Bayesian model.

The traditional Bayesian model starts with the following two principles.

- At any moment, the agent's credal states are represented by a probability function.
- From moment to moment, the agent's credal states are updated by conditionalisation on the evidence received.

Over recent decades many philosophers have been interested in models that relax those assumptions. One particular model that has got a lot of attention (from e.g. Isaac Levi, Richard Jeffrey, Bas van Fraassen, Alan Hájek and many others) is what I'll call the *static Keynesian model*. This model has the following features.

- At any moment, the agent's credal states are represented by a set of probability functions, called their representor.
- The agent holds that  $p$  is more probable than  $q$  iff the probability of  $p$  is greater than the probability of  $q$  according to all probability functions in their representor. The agent holds that  $p$  and  $q$  are equally probable iff the probability of  $p$  is equal to the probability of  $q$  according to all probability functions in their representor.
- From moment to moment, the agent's credal states are updated by conditionalising each of the functions in the representor on the evidence received.

The second point is the big attraction. It allows that the agent need not hold that  $p$  is more probable than  $q$ , or  $q$  more probable than  $p$ , or that  $p$  and  $q$  are equally probable, for arbitrary  $p$  and  $q$ . And that's good because it isn't a rationality requirement that agents make pairwise probability judgments about all pairs of propositions. Largely because of this feature, I argued in [Keynes, Uncertainty and Interest Rates](#) that this model could be use to formalise the key philosophical ideas in Keynes's *Treatise on Probability*. That's the reason I call this a 'Keynesian' model.

The modifier ‘static’ might seem a little strange, because the agent’s representor does change when she receives new evidence. But the change is always of a certain kind. Her ‘hypothetical priors’ do not change. If at  $t_1$  her evidence is  $E_1$  and her representor  $R_1$ , and at  $t_2$  her evidence is  $E_2$  and her representor  $R_2$ , then there is a ‘prior’ representor  $R_0$  such that the following two claims are true for all probability functions  $Pr$ .

- $Pr \in R_1 \leftrightarrow [\exists Pr_0 \in R_0: \forall p (Pr(p) = Pr_0(p | E_1))]$
- $Pr \in R_2 \leftrightarrow [\exists Pr_0 \in R_0: \forall p (Pr(p) = Pr_0(p | E_2))]$

That is, there is a set of probability functions such that the agent’s representor at any time is the result of conditionalising each of those functions on her evidence. I’ll call any model with this property a static model, so the model described above is the static Keynesian model.

Now there is a lot to like about the static Keynesian model, and I have made extensive use of it previous work. It is a particularly useful model to use when we need to distinguish between risk and uncertainty in the sense that these terms are used in Keynes (1937).<sup>2</sup> The traditional Bayesian model assumes that all propositions are risky, but in real life some propositions are uncertain as well, and in positions of radical doubt, where we have little or no empirical evidence, presumably most propositions are extremely uncertain. And using the static Keynesian model does not mean we have to abandon the great work done in Bayesian epistemology and philosophy of science. Since a Bayesian model is a (degenerate) static Keynesian model, we can say that in many circumstances (namely circumstances where *uncertainty* can be properly ignored) the Bayesian model will be appropriate. Indeed, these days it is something like a consensus among probabilists or Bayesians that the static Keynesian model is a useful generalisation of the Bayesian model. For example in Christensen 2005 it is noted, almost as an afterthought, that the static Keynesian model will be more realistic, and hence

---

<sup>2</sup> The clearest statement of the distinction that I know is from Keynes’s 1937 paper “The General Theory of Unemployment”

By ‘uncertain’ knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn. Or, again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth owners in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know. Nevertheless, the necessity for action and decision compels us as practical men to do our best to overlook this awkward fact and to behave exactly as we should if we had behind us a good Benthamite calculation of a series of prospective advantages and disadvantages, each multiplied by its appropriate probability, waiting to be summed. (114-5)

potentially more useful, than the traditional Bayesian model. Christensen doesn't appear to take this as any kind of *objection* to Bayesianism, and I think this is just the right attitude.

But just as the static Keynesian is more general than the Bayesian model, there are bound to be interesting models that are more general than the static Keynesian model. One such model is what I call the *dynamic* Keynesian model. This model has been used by Seth Yalcin to explicate some interesting semantic theories, but to the best of my knowledge it has not been used for epistemological purposes before. That should change. The model is like the static Keynesian model in its use of representors, but it changes the way updating is modelled. When an agent with representor  $R$  receives evidence  $E$ , she should update her representor by a two step process.

- Replace  $R$  with  $U(R, E)$
- Conditionalise  $U(R, E)$ , i.e. replace it with  $\{Pr(\cdot \mid E) : Pr \text{ is in } U(S, E)\}$

In this story,  $U$  is a function that takes two inputs: a representor and a piece of evidence, and returns a representor that is a subset of the original representor. Intuitively, this models the effect of learning, via getting evidence  $E$ , what evidential relationships obtain. In the static Keynesian model, it is assumed that before the agent receives evidence  $E$ , she could already say which propositions would receive probabilistic support from  $E$ . All of the relations of evidential support were encoded in her conditional probabilities. There is no place in the model for *learning* about *fundamental* evidential relationships. In the dynamic Keynesian model, this is possible. When the agent receives evidence  $E$ , she might learn that certain functions that were previously in her representor misrepresented the relationship between evidence and hypotheses, particularly between evidence  $E$  and other hypotheses. In those cases,  $U(R, E)$  will be her old representor  $R$ , minus the functions that  $E$  teaches her misrepresent these evidential relationships.

The dynamic Keynesian model seems well suited to the dogmatist, indeed to any epistemological theory that allows for fundamental evidential relationships to be only knowable a posteriori. In what follows I want to argue for five things. First, the dynamic Keynesian model is *coherent*. Second, there are reasons to want to model fundamental uncertainty using the dynamic Keynesian model rather than the static Keynesian model, let alone the traditional Bayesian model. Third, the dynamic Keynesian model does *not* validate the anti-dogmatist principle. Fourth, there are reasons to think the dynamic Keynesian model is particularly appropriate to cases where earlier models suggested the anti-dogmatist principle could be derived. Finally, the dynamic Keynesian model itself could well be generalised by later models.

It is sometimes argued that any alternative to the Bayesian model is incoherent. The most common argument for this is the Dutch Book argument due to David Lewis and Paul Teller. (What I have to say here generalises I think to other arguments for the unique coherence of the traditional Bayesian model.) This argument purports to show that any agent who cannot be modelled by the traditional Bayesian model will positively value each member of a set of bets



(placed at different times) that combined leads to sure loss. This kind of reason cannot tell against the dynamic Keynesian model. To see this, let  $Pr$  be some probability function in  $U(R, E)$ , and consider any series of bets made before and after updating on  $E$ . Since  $Pr$  is in  $U(R, E)$ , and  $U(R, E)$  is a subset of  $R$ ,  $Pr$  is in  $R$ . So if the agent positively values each bet in the series, each bet must have a positive expected value according to  $Pr$ . (I'm assuming here, as is standard in Keynesian models, that a bet is positively valued iff it has positive value according to all members of the representor.) So if this agent positively values each bet, then an agent whose credences are represented by  $Pr$ , and who updates by conditionalisation (i.e. an agent who can be modelled by the traditional Bayesian model) will positively value each such bet. But it can be, and usually is, shown that there is no Dutch Book against such agents. So similarly there is no Dutch Book to be made against agents in the dynamic Keynesian model. (This point about the impotence of Dutch Book arguments as anti-Keynesian arguments has been made by Alan Hájek, though he focussed on the synchronic Dutch Book case.)

One really nice consequence of the dynamic Keynesian approach is that it lets us say what the representor of an agent with no empirical information should be. Say a proposition is *a priori certain* iff it is a priori that all rational agents assign credence 1 to that proposition. Then the representor of the agent with no empirical evidence is  $\{Pr: \forall p: \text{If } p \text{ is a priori certain, then } Pr(p) = 1\}$ . Apart from assigning probability 1 to the a priori certainties, the representor is silent. Hence it treats all propositions that are not a priori certain in exactly the same way. This kind of symmetric treatment of propositions is not possible on the traditional Bayesian conception for logical reasons. (The reasons are set out in the various discussions of the paradoxes of indifference, going back to Bertrand 1888.) Such a representor is consistent with the static Keynesian conception, but it yields implausible results, since conditionalising on  $E$  has no effect on the distribution of values of  $Pr(p)$  among functions in the representor for any  $p$  not made a priori certain by  $E$ . (We'll say  $p$  is made a priori certain by  $E$  iff  $E \supset p$  is a priori certain.) So if this is our starting representor, we can't even get probabilistic evidence for things that are not made certain by our evidence.<sup>3</sup> So on the static Keynesian model, this attractively symmetric prior representor is not available.

I think one of the motivations of anti-dogmatist thinking is the thought that we *should* be able to tell a priori what is evidence for what. If it looking like there is a cow in front of us is a reason to think there is a cow in front of us, that should be knowable a priori. I think the motivation for this kind of position shrinks a little when we realise that an a priori prior that represented all the connections between evidence and hypotheses would have to give us a lot of

---

<sup>3</sup> The argument in the text goes by a little quickly, because I've defined representors in terms on unconditional probabilities and this leads to complications to do with conditionalising on propositions of zero probability. A better thing to do, as suggested by Hájek 2003, is to take conditional probability as primitive. If we do this we'll define representors as sets of conditional probability functions, and the a priori representor will be  $\{Pr: \text{If } p \supset q \text{ is a priori certain, then } Pr(q | p) = 1\}$ . Then the claim in the text will follow.

guidance as to what to do (epistemically speaking) in worlds quite unlike our own. Moreover, there is no reason we should have lots that information. So consider, for a minute, a soul in a world with no spatial dimensions and three temporal dimensions, where the primary source of evidence for souls is empathic connection with other souls from which they get a (fallible) guide to those souls mental states. When such a soul conditionalises on the evidence “A soul seems to be jealous of me” (that’s the kind of evidence they get) what should their posterior probability be that there is indeed a soul that is jealous of them? What if the souls have a very alien mental life, so they instantiate mental concepts very unlike our own, and souls get fallible evidence of these alien concepts being instantiated through empathy? I think it’s pretty clear we *don’t* know the answers to these questions. And those souls are presumably just as ignorant about the epistemologically appropriate reaction to the kinds of evidence we get, like seeing a cow or hearing a doorbell, as we are when thinking about their evidence. The dynamic Keynesian model can allow for this, especially if we use the very weak prior representor described above. When we get the kind of evidence we actually get, the effect of  $U$  is to shrink our representors to sets of probability functions which are broadly speaking epistemically appropriate for the kind of world we are in. Before we got that evidence, we didn’t know how we should respond to it, just like the spaceless empathic souls don’t know how to respond to it, just like we don’t know how to respond to their evidence.

The Bayesian argument against dogmatism was based around the principle that getting evidence  $A$  lowers the probability of  $\neg F$ . To see whether this is true on the dynamic Keynesian model, we need to say what it is to *lower* the probability of some proposition. Since representors map propositions onto intervals rather than numbers, we can’t simply talk about one ‘probability’ being a smaller number than another.<sup>4</sup> On the static Keynesian model, the most natural move is to say that conditionalisation on  $E$  *lowers* the probability of  $p$  iff for all  $Pr$  in the representor,  $Pr(p) > Pr(p | E)$ . This implies that if every function in the representor says that  $E$  is negatively relevant to  $p$ , then conditionalising on  $E$  makes  $p$  less probable. Importantly, it allows this even if the values that  $Pr(p)$  takes across the representor before and after conditionalisation overlap. So what should we say on the dynamic Keynesian model? The weakest approach that seems viable, and not coincidentally the most plausible approach, is to say that updating on  $E$  lowers the probability of  $p$  iff the following conditions are met:

- For all  $Pr$  in  $U(R, E)$ ,  $Pr(p | E) < Pr(p)$
- For all  $Pr$  in  $R$  but not in  $U(R, E)$ , there is a  $Pr'$  in  $U(R, E)$  such that  $Pr'(p | E) < Pr(p)$

---

<sup>4</sup> Strictly speaking, the story I’ve told so far does not guarantee that for any proposition  $p$ , the values that  $Pr(p)$  takes (for  $Pr$  in the representor) form an interval. But it is usual in more detailed presentations of the model to put constraints on the representor to guarantee that happens, and I’ll assume we’ve done that.

It isn't too hard to show that for some models, updating on A does not lower the probability of  $\neg F$ , if lowering is understood this way. The following is an extreme example, but it suffices to make the logical point. Let R be the a priori representor, i.e. the set of all probability functions that assign probability 1 to a priori certainties. And let  $U(R, A)$  be the singleton of the probability function given in the table above (when discussing the vision/sonar case). Then the probability of  $\neg F$  after updating is 0.5. (More precisely, according to all  $Pr$  in  $U(R, A)$ ,  $Pr(\neg F) = 0.5$ .) Since before updating there were  $Pr$  in R such that  $Pr(\neg F) < 0.5$ , in fact there were  $Pr$  in R such that  $Pr(\neg F) = 0$ , updating on A did not *lower* the probability of  $\neg F$ . So the dynamic Keynesian model does not, in general, have as a consequence that updating on A lowers the probability of  $\neg F$ . This suggests that it is possible that evidence A support our knowledge that  $\neg F$ .

It might be objected that if evidence A supports our knowledge that  $\neg F$ , then updating on A should *raise* the probability of  $\neg F$ . And if we define probability raising the same way we just defined probability lowering, updating on A *never* raises the probability of  $\neg F$ . From a Keynesian perspective, we should simply deny that evidence has to raise the probability of the propositions known on the basis of that evidence. It might be sufficient that getting this evidence removes the uncertainty associated with those propositions. Even on the static Keynesian model, it is possible for evidence to remove uncertainty related to propositions without raising the probability of that proposition. A little informally, we might note that whether an agent with representor R is sufficiently confident in  $p$  to know that  $p$  depends on the lowest value that  $Pr(p)$  takes for  $Pr \in R$ , and updating can raise the value of this 'lower bound' without raising the value of  $Pr(p)$  according to all functions in R, and hence without strictly speaking *raising* the probability of  $p$ .

The above illustration is obviously unrealistic, in part because U could not behave that way. It's tempting at this stage to ask just how U does behave so we can work out if there are more realistic examples. Indeed, it's tempting to try to attempt to provide a formal description of U. This temptation should be resisted. The whole point of the model is that we can only learn which hypotheses are supported by certain evidence by actually getting that evidence. If we could say just what U is, we would be able to know what was supported by any kind of evidence without getting that evidence. The best we can do with respect to U is to discover some of its contours with respect to evidence much like our own. And the way to make those discoveries will be to do epistemology; it isn't obvious that, say, looking for nice formal properties of U will help at all.

Although the dynamic Keynesian model is different to the traditional Bayesian model, it is not strictly speaking a competitor to the Bayesian model. Static Keynesian models are special cases of dynamic Keynesian models, and traditional Bayesian models are special cases of static Keynesian models. So even if in general we should model learning using a dynamic Keynesian model, that does not suggest that it is wrong to ever use traditional Bayesian models. It just has to be the case that (a) there are some conditions such that when they are satisfied, we don't need to

extra generality the dynamic Keynesian model provides, and (b) those conditions are satisfied in the particular application we are looking at.

In general, we can use a static Keynesian model rather than a dynamic Keynesian model if the situation being modelled does not involve epistemological learning. That is, if the situation being modelled is not such that the new evidence agents get provides them with justification for epistemological theses they were previously not justified in believing. For what *U* represents is just this kind of learning; learning what kinds of hypotheses are justified by what evidence. And we can use a traditional Bayesian model rather than a static Keynesian model if the distinction between risk and uncertainty is not relevant to the case. Now it is very plausible that in the cases where Bayesian modelling has proven most useful, such as in explaining away Hempel's paradox of confirmation, those two conditions are satisfied. But we should not expect that the conditions will *always* be satisfied.

In particular, in the case at issue here, *neither* condition is plausibly satisfied. Unless we assume that facts about which evidence justifies belief in which propositions is a priori, we should think this *is* a situation where getting new evidence involves epistemological, not just empirical, learning. And if there is any situation where the risk/uncertainty distinction matters, it is the situation the Bayesian is trying to model with their hypothetical priors, i.e. the situation of an agent with no empirical evidence. So there is strong reason to think that, even if we are sympathetic to a broadly Bayesian approach to issues in philosophy of science, that we need a more general approach here. But if we use more general models, such as dynamic Keynesian models, the anti-dogmatist argument does not go through. So I conclude that this anti-dogmatist argument relies on improperly applying models that are appropriate to everyday cases, i.e. Bayesian models, to a situation not at all like the everyday, i.e. the situation of receiving for the first time empirical evidence about the external world.

Nothing I've said here should be taken to imply that the dynamic Keynesian model is the *most* general model we might want to use. Although it is much more general than the traditional Bayesian model, there are two ways in which we might for some purposes want a more general model. First, in cases where we are trying to model logical ignorance, we might want a model that didn't assign all tautologies probability 1. Second, in cases where we think the distinction between empirical learning and epistemological learning cannot be made precise, we might not want to assume that we can factorise the effect of evidence on credence into a contraction of the representor (representing the epistemological bit of learning) and conditionalisation of the remaining functions (representing the empirical bit of learning). It is an open question I believe just what a model that relaxes these assumptions should look like, but it is easy to imagine that such a model is desirable. But the desirability of such a model should not imply the uselessness of the somewhat less general dynamic Keynesian model. And it certainly should not make us think, as the Bayesian argument against dogmatism requires, that we should instead look to a *much* less general model, the traditional Bayesian model, when doing foundational work in epistemology.