[[Forthcoming in *Synthese* – this is an uncorrected draft, so please don't quote from or circulate this version!]]

Suspending is Believing

Thomas Raleigh, traleigh@gmail.com

ABSTRACT: A good account of the agnostic attitude of Suspending Judgement should explain how it can be rendered more or less rational/justified according to the state of one's evidence – and one's relation to that evidence. I argue that the attitude of suspending judgement whether p constitutively involves having a belief; roughly, a belief that one cannot yet tell whether or not p. I show that a theory of suspending that treats it as a *sui generis* attitude, wholly distinct from belief, struggles to account for how suspension of judgement can be rendered more or less rational (or irrational) by one's evidence. I also criticise the related idea that suspension essentially requires an 'Inquiring Attitude'. I show how a belief-based theory, in contrast, neatly accounts for the rational and epistemic features of suspending and so neatly accounts for why an agnostic has a genuine neutral *opinion* concerning the question whether p, as opposed to simply having *no* opinion.

1. Introduction

When there is enough reasonable doubt as to whether a proposition is true or false then one ought to suspend judgement and remain agnostic. Just as strong evidence serves to ground our beliefs, when the evidence is equivocal or weak it can ground the agnostic attitude of suspending judgement (a.k.a. 'withholding judgement'), making it more or less justified, more or less rational. But doubt can be reasonable not only due to the state of one's evidence; it can also be reasonable due to the *cognitive relation one bears to the evidence*. For example, you might reasonably suspend judgement because you have not yet had time to review and process all the available evidence, or because you realize that you are not expert enough or not currently in a fit mental state to draw a conclusion that is properly based on the evidence.

One kind of question you might ask here is: *how much* reasonable doubt is required for suspension of judgement to be reasonable or to be rationally required? Answering this sort of question has obvious important practical/ethical consequences, e.g. in the legal setting of a jury deliberating on its verdict. I will not attempt to address this sort of question in this paper.

A different, more theoretical question, which I will try to address in this paper, is: what kind of mental state or attitude is it to suspend judgement and remain agnostic? Taking the rational, epistemic profile of suspension seriously has, I contend, implications for the psychological nature of the mental state. I will argue that the attitude of suspending judgement whether p constitutively involves having a belief – roughly, a belief about what one can currently tell based on one's evidence. For it is only by treating suspending judgement as a species of belief that we can account for how this agnostic state can be rendered epistemically rational or irrational, justified or unjustified, according to one's evidence and one's cognitive relation to the evidence. Though some philosophers have occasionally endorsed or assumed that suspension involves some kind of belief, there has so far been no extended defence of this general position[1].

[1] Rosenkranz (2007) is certainly an extended and highly sophisticated elaboration of a specific belief-based account of agnosticism. However, Rosenkranz is not concerned with *defending* the basic assumption that agnosticism essentially involves some kind of commitment or belief that we are not in a position to know either that p or that not-p. Rosenkranz simply takes this for granted in the course of

For simplicity, I will stick throughout to considering the state of suspending judgement whether a single proposition is true (or false), though in English we also speak of suspension concerning more open-ended questions – e.g. "I'm suspending judgement as to who the murderer is", "I'm agnostic over where the treasure is buried" etc[2]. In what follows, I will sometimes abbreviate suspending judgement whether p as 'Sp' – and likewise abbreviate belief that p as 'Bp' and Disbelief that p as 'DBp'[3]. A distinction is sometimes drawn between 'occurrent' and 'dispositional' beliefs. However I find the notion of an 'occurrent belief' somewhat dubious and prefer to think in the following way: a mental *act* or *event* of judging that p (generally at least) initiates a mental *state* (however short- or long-lived) of belief that p. Likewise then I will assume that an initial agnostic act or event of suspending judgement initiates an agnostic state of suspension of judgement.

The plan for this paper is as follows: in the next section, (2), I identify two constraints that I think any good theory of suspending judgement should satisfy and also consider four possible theories about the nature of suspending judgement: the non-belief view, the credence view, the belief view and the sui generis view. I suggest that we should discard the non-belief view and the credence view. In section (3) I formulate and clarify the specific kind of belief-based view that I want to defend. In section (4) I briefly defuse a line of thought that might seem to count against this belief-based view. In section (5) I present my core argument that the sui-generis view cannot meet the two basic constraints. In section (6) I criticise Jane Friedman's specific, fleshed-out version of the sui generis view and her idea that suspending essentially involves having an 'Inquiring Attitude'. In section (7) I show that the belief-based view *can* meet the two constraints – and also briefly consider some complications concerning higher-order evidence. Section (8) provides a short conclusion.

2. Two Constraints and Four Theories

I'll take for granted as a basic starting assumption that suspending judgement is an attitude that can be rendered more or less epistemically justified, more or less epistemically rational, according to one's evidence and also according to one's cognitive relation to that evidence. I take it that this assumption that suspending is subject to norms of *epistemic* rationality is widely shared insofar as suspending is standardly grouped together and discussed along with belief and disbelief as a third doxastic attitude. And just as with belief and disbelief, it seems extremely plausible and intuitive that however practically rational or irrational it would be for a subject to suspend judgement (e.g. an evil demon will lavishly reward you, or perhaps viciously punish you, for suspending judgement), there will always be a further separate question as to whether it would be epistemically rational for the subject to suspend judgement, a question that will depend (at least in large part) on the state of the subject's evidence. However, whilst I take this assumption to be extremely plausible and widely shared, it is worth noting that it is at least a logical option to insist that suspension is not subject to the demands of epistemic rationality at all and can be rendered only more or less practically rational[4].

---

explicating what he calls 'true agnosticism', which is, roughly, the view that we will continue to be ignorant as to whether p in all future states of information that we could reach using present methods. Rosenkranz's position is briefly discussed further in sections 2 and 5, below.

[2] The position I will be arguing for – that suspension essentially involves a belief about what one can currently tell– extends pretty straightforwardly, I think, to states of suspending-*wh*, where the question being suspended about is not just the truth/falsity of a single proposition, but is a question for which many propositions are potential candidate answers.

[3] I'll assume for simplicity that $DBp \leftrightarrow B\neg p$.

[4] Many thanks to an anonymous referee for this journal for pressing me to clarify this.

As a second guiding assumption, I will also take it to be uncontroversial that suspending judgement whether p is a different state of mind from simply *having no opinion whatsoever* concerning p's truth or falsity. An agnostic has a genuine opinion whether p, though it is a *neutral opinion*. Whereas someone who has never even entertained the proposition p in thought, whether consciously or unconsciously, simply has *no opinion* about p's truth or falsity and we would not count them as suspending judgement or being agnostic whether p. Likewise someone who manages to deliberately refrain from having any kind of intellectual response, explicit or implicit, to the question whether p – 'I simply refuse to think about this!' – is not naturally counted as being an agnostic on the matter. Nor is someone who starts to think about the question whether p but then just trails off and starts to think about something else without forming any opinion whatever concerning the question whether p.

In one of the very few recent papers to focus on what the mental state of suspension is exactly[5], Jane Friedman writes:

> "Think about the theistic agnostic: she is not trying to merely dodge some beliefs; she is in some state that captures **her take or opinion** about whether God exists, **about the truth of matter**." (Friedman, 2013a, 173, bold type added)
> "The agnostic about the existence of God is not someone who lacks an opinion about whether God exists, but someone who has an opinion on the matter: a (roughly) neutral one." (ibid. 178)

We have then two basic constraints that any account of suspending judgement should satisfy:

> **(I)** Suspending whether p is an attitude that can be rendered more or less epistemically justified/rational[6], at least to some extent[7], according to one's evidence and one's cognitive relation to that evidence.
> **(II)** Suspending whether p is different from simply *having no opinion* whether p.

In the above-quoted paper, Friedman is primarily concerned with establishing what suspension of judgement is *not*. She argues against accounts that analyse suspending judgement whether p *not* as a positive mental state or attitude with p (or perhaps the question whether p) as part of its content, but *merely as a lack of belief* in either p or in not-p. I.e. Friedman is arguing against the following:

- **(NO-BEL)** NON-BELIEF VIEW: Suspending whether p just consists in neither believing that p nor disbelieving that p. Possible extra conditions: (i) the subject has considered the proposition p, (ii) the subject actively/deliberately refrains from believing/disbelieving that p, (iii) the subject's lack of belief/disbelief is for 'epistemic' reasons.

The second basic constraint, (II), immediately rules out the simplest version of the Non-Belief view – that suspending is just a lack of belief and disbelief. Friedman considers and rejects various possible extra necessary conditions that a (NO-BEL) theorist might try to add. So the putative requirement that the subject have considered the proposition in question is neither a

---

[5] In addition to Friedman's writings, some other recent papers which focus, in different ways, on suspension include: Booth 2012, Turri 2012, McHugh 2012 and Tang 2016.

[6] Of course, for certain purposes the notions of justification and rationality may come apart. But so far as the arguments of this paper are concerned we will not need to distinguish between them too carefully.

[7] I include this clause so as not to assume *evidentialism* about justification/rationality – i.e. that one's evidence exhaustively determines the epistemic justification/rationality of one's doxastic attitudes. I leave open here that there may be other factors, in addition to one's evidence, that may be relevant to epistemic justification/rationality.

necessary condition – since suspension could, in principle, be induced by ingenious neuro-surgery – nor a sufficient condition on suspending – since one might consider the proposition in question but then get bored or distracted (etc) and fail to form any kind of neutral opinion in response. Likewise the putative requirement that the subject actively refrain from forming a belief is not necessary – again since agnosticism could in principle be induced without any mental effort/act on the part of the subject – and not sufficient – since a subject may actively refrain from forming any opinion whatsoever, for example if she thinks that thinking about the question whether p is immoral or will bring about a panic attack, etc. And finally, the putative requirement that the subject's lack of belief be for 'epistemic reasons' does not work either – since a subject may lose or forget her epistemic reasons for refraining from belief and yet still be agnostic. The problem with these more complicated versions of NO-BEL then is that they all still fail to satisfy constraint (II), for they fail to capture how exactly agnosticism differs from simply having no opinion – i.e. from simply lacking belief and disbelief.

Notice that the problem with NO-BEL is *not* that it cannot satisfy constraint (I). Consider a simple sort of creature that can form beliefs and that has enough rational sensitivity to evidence that it generally manages to believe that p only when it has sufficiently good evidence that p and to disbelieve that p only when it has sufficiently strong evidence that not-p and otherwise it just does *nothing* – it neither believes nor disbelieves that p. This state of mere non-belief would be, for such a creature, the rational response to having weak or equivocal evidence – or at least, this non-belief would certainly be *more* rational than believing or disbelieving that p. But such a simple creature, in such a doxastically neutral state, would not have a *genuine neutral opinion* whether p – i.e. it would not be *agnostic* whether p.

Friedman's arguments against (NO-BEL) strike me as sound and so I will assume from now on that (NO-BEL) is not the correct account of suspending judgement.

Now, it might seem that a very natural way to satisfy our two constraints is to treat suspension as being a kind of credence or degree of belief – in particular a credence that falls short of either full belief or full disbelief, perhaps somewhere in the range close to 0.5.

- **(CR)** CREDENCE VIEW: Suspending whether p consists in having some 'middling' level of credence (degrees of belief) that p, which falls short of outright belief in either p or in not-p.
  See e.g. Hajek (1998), Christensen (2009).

Clearly, having some specific credence that p can be rendered more or less justified/rational according to the strength of one's evidence. And so long as one's theory of credences allows for a subject to have no credence function at all for some propositions, then the distinction between the neutral opinion of suspending and having no opinion whatever could be neatly accounted for as being the difference between having a middling credence that p and having no level of credence whatsoever about p.

But whilst the credence view of suspension can easily satisfy our two constraints, it faces other serious difficulties.

Friedman (2013b) points out that when a subject rationally suspends over each of the independent propositions p_1, p_2, p_3, p_4… she can presumably (sometimes at least) also permissibly suspend over their disjunction (p_1 v p_2 v p_3 v p_4…) or over their conjunction (p_1 & p_2 & p_3 & p_4…). But if suspending was just having a mid-level credence, once you have enough independent propositions one's credence for their conjunction or their disjunction ought to tend towards zero or towards one – i.e. you ought *not* suspend over the conjunction or the disjunction of these propositions. Lottery cases seem to provide a similar sort of difficulty for (CR) – e.g. when I have bought a lottery ticket, my credence that it will win

could be extremely low and yet it seems I can still, quite rationally, suspend judgement as to whether it is a winning ticket. And even in the case of a single proposition, as Sturgeon (forthcoming) points out, it seems psychologically very plausible that I can rationally suspend judgement as to whether p/not-p without assigning any specific subjective probability or degree of confidence to p. Decker (2012) points to the sorts of cases that are standardly used to make trouble for the Principle of Indifference to argue for the claim that in these cases, where we can apparently partition up the space of possibilities in more than way, we should rationally have gaps in our credence functions and assign no credence-value whatever to certain propositions – a state which he calls 'deep agnosticism'.

Now much more could be said about the relation between degrees of belief and suspending judgement. But so far as this paper is concerned, I will assume that (CR) does not state a necessary condition on suspension either[8]. Having ruled out (CR) and (NO-BEL), what other options are there?

- **(BEL)** BELIEF VIEW: Suspending whether p consists in having some kind of belief – e.g. about whether one ought to believe p, about whether one is in a position to know whether p.
  See e.g. Crawford (2004), Bergmann (2005), Rosenkranz (2007)

- **(SG)** SUI GENERIS VIEW: Suspending whether p consists in having a 'sui generis' mental attitude concerning p that does *not essentially involve* either full or partial belief.
  See e.g. Friedman (2013a, 2017), Sturgeon (forthcoming)

Notice that, so far, (SG) is just defined negatively, in opposition to the other three views, as being *some kind* of positive mental state with a content concerning p but which does not essentially involve any kind of full or partial belief. In section 5, below, I will discuss the only fleshed out version of the view, due to Friedman 2017, which says something more about the specific positive nature of the proposed sui generis state.

Notice also: I have formulated these views in terms of what *constitutes* the mental state of suspending judgement. But there is also the theoretical possibility of 'hybrid' positions, which would hold that suspending judgement constitutively requires meeting *more than one* of these conditions. Another possible kind of hybrid position would be to hold that whilst only one of these views is a correct account of what constitutes suspension, nevertheless one or more of the other views provide a *necessary but non-constitutive condition* – e.g. suspending judgement whether p *consists* in, say, a special sui generis attitude towards p, but it is nevertheless also a necessary condition for suspending judgement that you have, say, a middling credence that p, though this credence state is not part of what constitutes the state of suspending judgement. Yet another possibility would be to endorse a kind of disjunctive position, according to which Suspending judgement can be, e.g., constituted *either* by a having middling credence, *or* by having a certain kind of belief, though neither condition is individually necessary for suspension. Having noted that they are possible options 'in logical space', I will not consider these sorts of mixed positions any further. It will provide quite enough of a task for the present paper to compare the respective merits of the pure versions of these positions. To repeat then: according to the pure version of the Sui Generis (SG) view to be discussed below, suspending does *not* essentially involve any kind of full belief or partial belief (credence).

[8] Notice though: even if (CR) does not state a *necessary* condition on suspending judgement, perhaps having a credence of around 0.5 can be a *sufficient* condition for *one way* of suspending whether p. See the possibility of 'disjunctive' theories discussed below.

Friedman (2013a) intriguingly suggests that the sui generis option, (SG), is the correct one, a suggestion that is further developed in Friedman (2017)[9]. I should confess that I had always intuitively thought of suspension along *something like* the lines of (BEL), though without really having considered the matter – certainly without having considered (SG) as a rival possible option. And so, having been awoken from my dogmatic slumbers, I want to say something more satisfactory in favour of (BEL) and also to say something against (SG).

3. The Meta-Cognitive Belief view

Let me now try to be a bit more precise about the nature of the belief in question.
I think that the basic thought that marks out an agnostic opinion whether p from simply having no opinion whatever is the belief that: *I cannot currently tell whether or not p*. The agnostic subject then must have a perspective on her own cognitive/epistemic position to the extent that she takes herself to be not yet able to properly form a judgement on the matter based on her evidence. Such a meta-cognitive opinion – about what one can currently tell concerning some question – plausibly requires some degree of cognitive sophistication (see section 4, below). But to be agnostic, one need not have any kind of sophisticated, philosophical notions of what counts as evidence (mental state, proposition, fact, concrete particular, etc.) or how much evidence would rationally permit judgement. The idea is just that: so far as I can currently tell, nothing yet is clearly enough indicating to me either the truth or the falsity of p. So the specific formulation we'll work with is

- **(BEL-M)** META-COGNITIVE BELIEF VIEW: Suspending whether p constitutively requires having a belief or opinion that one cannot yet tell whether or not p, based on one's evidence.

Notice: the belief in question does not have to be about whether the evidence *objectively* indicates that p, or what it would indicate to an ideally rational agent, or even what it indicates to me once I have had a bit more time to think it over. So the belief in question is not about what I am *in a position* to know (given more time to think, or more coffee, or fewer distractions etc.) It would suffice to think that: for all I can currently bring to mind and grasp, overall nothing yet seems to indicate either p or not-p clearly enough for me to form a definite judgement. So one need not have formed a *settled* neutral opinion, in order to count as suspending judgement. One might suspend judgement as to whether p having only so far considered a fraction of the available evidence that one takes to be potentially relevant. Indeed, one might have *forgotten* what exactly the available evidence is – it is enough to be agnostic that one retains a belief that the evidence, whatever it was exactly, has not yet allowed one to form a judgement either way. On the other hand, you might suspend judgement if you believe that there is, or even just may be, lots more relevant evidence out there that is as yet unavailable to you. So then even if the available evidence quite strongly indicates that p (or that not-p), given the live possibility of extra, currently unavailable evidence, you might still take your available evidence to be such that you cannot tell whether or not p.

I take the view (BEL-M) proposed here, to be in the same spirit as the following statements, both made somewhat in passing, of a meta-cognitive requirement on suspension by Crawford (2004) and Bergmann (2005):

> 'Suspension of judgement necessarily involves thoughts about one's own epistemic perspective on whether or not p, namely, that one's epistemic perspective falls short of establishing whether p and thus that one does not know whether p.' (Crawford, 2004, 226)

[9] Scott Sturgeon (forthcoming) also endorses a version of the Sui Generis option.

> '…to withhold p (in the sense I have in mind) one must, in addition, consider the prospect of one's believing p as well as the prospect of one's disbelieving p; otherwise one will not be able to resist both believing p and disbelieving p. So withholding p involves not only an attitude towards p but also attitudes towards attitudes toward p.' (Bergmann, 2005, 421)

It is also close to the following statement by Rosenkranz (2007):

> 'The most basic species of agnosticism about p asserts… that we are in a position neither to know p nor to know ¬p.' (Rosenkranz, 2007, 62)

However it should be noted that in Rosenkranz's terminology agnosticism and suspension of judgement are *different* attitudes – he uses the term 'suspension' to refer to what I am calling mere lack of belief and disbelief, i.e. mere doxastic neutrality. Also, Rosenkranz is primarily concerned to argue for the possibility of what he calls 'True Agnosticism':

> 'Roughly, True Agnosticism about a given range of statements emerges as the view that we are not at present in a position to know the truth-value of these statements and will continue to be ignorant about this matter in all subsequent states of information which we can reach by applying our present methods and working out the implications of the background knowledge we will thereby acquire.' (Rosenkranz 2007, 58)

To repeat: the sort of agnosticism I am concerned with, and which I take our everyday talk of suspending judgement to be about, does not require any such settled, strong commitment about the futility of future inquiry. Though it is important that an agnostic *may* have such a settled opinion – as we will discuss in section 6, below.

Now, as I am thinking of the agnostic attitude, neither (BEL-M) nor (SG) by themselves could be sufficient for a subject to be suspending judgement as a (conflicted/irrational) subject might satisfy either of these conditions whilst still simultaneously believing that p (or that not-p). In addition to meeting these putative conditions then, I think that one must also actually neither believe p nor disbelieve p in order to count as being agnostic towards p. So, to be clear, on the position I favour the agnostic attitude would not consist *solely* in having the required belief, as specified by (BEL-M). An agnostic must also satisfy the following condition:

- DOXASTIC NEUTRALITY: Suspending whether p also requires that the subject is in a *neutral doxastic state* with respect to p – i.e. she neither believes that p nor disbelieves that p.

In other words, doxastic neutrality is a further necessary, but not sufficient, condition for having an agnostic's genuine *neutral opinion*. (This requirement of doxastic neutrality is also assumed by both Crawford 2004 and Bergmann 2005) Notice, of course, that we will need to individuate the contents of propositional attitudes in a sufficiently fine-grained way for this constraint to be plausible. For example: it is presumably possible to both believe that "Hesperus is Venus" and suspend whether "Phosphorus is Venus".

And in fact, in addition to satisfying (BEL-M) and DOXASTIC NEUTRALITY, I think that there also needs to be the *right relation* between these two conditions. For consider a case where a subject happens to have a belief that she cannot currently tell whether or not p, but this meta-cognitive belief has *no explanatory connection* to her doxastic neutrality concerning p; and in fact her doxastic neutrality is due to some other entirely non-epistemic, non-evidential reason – such as getting bored or distracted when thinking about the question. It seems that such a subject should *not* count as an agnostic about p, but as someone who has

failed to form any opinion whatever as to whether p. Thus it seems that a genuinely agnostic subject's meta-cognitive belief should be (part of) *the reason* for her doxastic neutrality. To be clear then: the belief-based view of suspension I would want to defend holds that suspending judgement whether p is a complex state consisting of both a meta-cognitive belief (BEL-M) and the state of doxastic neutrality concerning p *and* the right sort of connection between these two factors[10].

Can we say something more specific about the nature of this connection? Well, I suspect that there will be some tricky issues lurking here – witness the large literature that exists trying to formulate exactly what the 'basing relation' is, between a reason and the belief for which it is a reason[11]. And so I will not attempt to say fully and precisely what this 'right sort of connection' consists in. But I think it is clear that we want the meta-cognitive belief to be the subject's *motivating reason* for the doxastic neutrality and not just a merely causal-explanatory reason[12]. The insight here is due originally to Donald Davidson's discussion of reasons for action, but the same point also applies to the basing relation in epistemology[13]. A mental state can be both a reason to phi and also be part of the causal explanation for why the subject phi's and yet the mental state is *not* the subject's motivating reason for phi-ing – or as it is sometimes put, the state is not the reason *for which* the subject phi-d. Here is Davidson's classic example:

> 'A climber might want to rid himself of the weight and danger of holding another man on a rope, and he might know that by loosening his hold on the rope he could rid himself of the weight and danger. The belief and want might so unnerve him as to cause him to loosen his hold, and yet it might be the case that he never *chose* to loosen his hold, nor did he do it intentionally.' (Davidson, 1980, 79)

The climber's belief+desire are *a* reason to loosen his grip and were in fact (part of) the cause of his loosening his grip, but they were not the climber's motivating reason for which he loosened his grip. So now likewise, consider a situation where the subject has the relevant meta-cognitive belief (that she cannot currently tell whether p or not-p), but this meta-cognitive belief causes the subject to be doxastically neutral in some roundabout, non-motivating-reason way – e.g. the belief causes an intensely painful headache which causes the subject to forget about the question whether p entirely. Here again it seems that we would not want to count the subject as genuinely agnostic – she simply has no opinion whatever on the question whether p. Of course this still leaves us with the difficult question of how to determine whether a causal chain is of the right kind to allow a mental state to be a motivating reason, or whether it is of the wrong, 'roundabout' or 'deviant' kind. I will not attempt to say anything in answer to this here; I just take for granted that there is this distinction between a

---

[10] Notice that this complex state, consisting of both the meta-cognitive belief and the doxastic neutrality where the former is the subject's reason for the latter, is *different* from a version of the NON-BELIEF view, discussed in section 2, which Friedman (2013a) argued against. That NON-BELIEF view was that suspending judgement consisted *simply* in the state of doxastic neutrality, but with an added *necessary pre-condition* that this neutral state is maintained for 'epistemic reasons'. Friedman objected, correctly, that a subject may lose or forget her reasons for suspending and yet still be suspending. But this is not a problem for the complex state being proposed here. For it is certainly possible for the subject to lose her reasons for being in this overall complex state, whilst still being in the complex state. All that is required is that *within* the complex state of suspending judgement, one constituent – the meta-cognitive belief – is the subject's reason for the other constituent – the doxastic neutrality.

[11] See, for example, Lehrer (1971, 1990), Swain (1979), Audi (1986), Alston (1989), Korcz (1997), Evans (2013). A recent volume of essays devoted to the basing relation is Carter & Bondy (eds.) (forthcoming).

[12] Again, I am very grateful to an anonymous referee for this journal for pressing me to clarify this point.

[13] See, for example, Neta (forthcoming).

motivating reason and merely causal-explanatory reason. (We will return briefly to this requirement on the connection between (BEL-M) and DOXASTIC NEUTRALITY in section 6, below.)

Now, some may doubt the further requirement of DOXASTIC NEUTRALITY – indeed Freidman (2017) explicitly rejects it. Compare: some philosophers hold that it is possible for a subject to simultaneously believe p and disbelieve p. E.g. perhaps there simply need to be two mental representations with contradictory contents, both with the right sort of causal-functional roles to qualify as beliefs. Of course this would be an irrational combination, but it would qualify the subject as both believing and disbelieving that p. Likewise then you might think that a subject could (albeit irrationally) both believe that p and suspend whether p at the same time – again, it could just be a case of having rationally conflicting inner representational states.

Let me just flatly state that this strikes me as wrong. If a subject can truly be said to believe that p, whatever your preferred metaphysics for beliefs, then she is *not suspending judgement whether p* – she has precisely *failed* to suspend judgement concerning p! Suspending judgement on some question and remaining agnostic is a difficult cognitive achievement that we humans, with our thirst to form beliefs and rush to judgement, often fail at. If a subject unconsciously believes that p, no matter how sincerely they avow that they are suspending judgement, or 'keeping an open mind', they cannot count as genuinely agnostic[14].

Another issue to consider concerns what you might call 'higher-order' suspension of judgement[15]. Assuming that we sometimes have imperfect self-knowledge, it seems fairly straightforward that one might both S(p) and S(Sp). The higher-order attitude, S(Sp), would evince introspective caution or uncertainty, and if one did in fact lack suitable introspective access to one's own 1st-order agnostic attitude it seems that this could be a perfectly rational combination of attitudes to have. However, consider instead the following combination of attitudes:

S[p] and S[I cannot currently tell whether p]

Given the assumption that DOXASTIC NEUTRALITY is a requirement on suspension, this combination of attitudes will be ruled out as impossible. For, such a combination would require the subject to:

B[I cannot currently tell whether p] and S[I cannot currently tell whether p]

But according to DOXASTIC NEUTRALITY, suspending judgement whether p requires the *absence* of belief or disbelief that p.

Another obvious consequence of any belief-based account of suspension is that it will be impossible to *exclusively* suspend judgement about every proposition one entertains in thought and to believe/disbelieve nothing at all – for suspending is held to constitutively involve a belief. Any belief-based account of suspension then will be committed to rejecting the following two claims, which we might label as "Pyrhhonian"[16] theses:

---

[14] Again, I am assuming that the contents of these attitudes are sufficiently fine-grained to rule out Frege-style cases where a subject rationally believes one claim and suspends about another claim that is, in fact, equivalent.

[15] See Sorensen (2009), who uses the label 'meta-agnosticism', for further discussion.

[16] I cannot vouch for the historical accuracy of applying the label 'Pyrrhonian' to these theses! I only mean to suggest that, rightly or wrongly, these kinds of theses will naturally invoke the idea of Pyrhonnian *epoché* in many modern readers.

(i) It can be rational to suspend judgement about every proposition one entertains.
(ii) It is possible to exclusively suspend judgement without believing anything at all.

Bear in mind, though, constraint (II) above: suspending judgement is different from having no opinion whatever. And so, of course, rejecting these two theses is consistent with allowing that it could be possible to simply have *no opinion whatever* – neither belief, disbelief nor suspension of judgement – about anything.

Notice: in stating (BEL-M) I used the phrase 'Suspending whether p requires having a belief **or opinion**…' – this was in order to leave open the possibility that *full belief* about what one can currently tell might not be required. Perhaps one could have some fairly high level of credence, which falls short of outright belief, that one cannot currently tell whether or not p and yet still count as agnostic whether p. But what about the following sort of attitude?

Cr[I cannot tell whether or not p] ≈ 0.5

According to belief-based accounts of suspension, having a credence of around 0.5 that one's evidence is inconclusive is not the same as suspending judgement whether one's evidence is inconclusive. I don't want to take a stand here on the relation between full/outright belief and degrees of belief, so I leave open whether it is even possible to both B(p) and have a Cr(p) ≈ 0.5. If it is impossible, then S(p) and Cr[I cannot tell whether or not p] ≈ 0.5 would also be an impossible combination of attitudes. On the other hand, if it is merely irrational but still possible to both B(p) and have a Cr(p) ≈ 0.5, then likewise S(p) and Cr[I cannot tell whether or not p] ≈ 0.5 will be a possible but irrational combination.

One final possible complication for any candidate account of suspension: could it be possible to suspend judgement whether p *without* thereby also suspending whether not-p? In terms of the (BEL-M) account this would require having a belief that one cannot yet tell whether or not p, whilst also having *no belief* as to whether one can tell whether or not not-p. Prima facie, this seems bizarre to the point of impossibility – surely to suspend judgement about p's truth is also to suspend judgement about its falsity? Friedman (2017) appears to endorse this idea: 'In general, suspension seems to be closed under negation – one suspends about p if and only if one suspends about ¬p' (p3) [17]. However, I think we might be cautious here as in various non-classical contexts, this sort of bi-conditional may become less plausible. For example in a context where truth is taken to be equivalent to provability, a subject might think that the evidence definitely does not allow for a proof of p, but have no view at all as to whether the evidence allows for a proof of not-p. Thus she might disbelieve that p but suspend judgement whether not-p.

No doubt further possible clarifications and issues concerning the specific nature of the belief required by (BEL-M) might be made or might be needed for some purposes. But I hope that for present purposes we now have a specific-enough version of the belief-based account of suspension to be getting on with. And though it is worth having a tolerably precise thesis 'on the table', in fact the specific, fine-grained details of the belief in question won't matter all that much to the argument I'll present. Much the same line of thought, mutatis mutandis, could be used to support a more normatively-couched version of the belief view – according to which suspending whether p requires a belief that (given what one can presently tell based on the evidence) one *ought* neither believe nor disbelieve p.

[17] I say 'appears' as the phrase 'in general' is ambiguous between '*most* of the time' and 'in total exception-less generality'.

4. Simplicity and Sophistication

Before presenting my case against (SG) and in favour of (BEL-M), I want to briefly defuse a line of thought that might seem to favour (SG) over (BEL-M).

It is clear enough that one can form a belief that p despite having no opinion whatsoever about one's evidence for p, nor about one's cognitive relation to the evidence, nor about whether one ought/ought not believe that p, etc. I think it is also pretty plausible, unless you are some sort of fairly extreme internalist, that one could form a belief that p *as a rational response* to one's evidence, without having any explicit further belief *about* that evidence. And so you might think that we should favour an account that treats suspending whether p along the same lines – which would then seem to favour (SG) over (BEL-M).

But why should we accept that it is a theoretical desideratum for belief and suspension to be treated the same way in this respect? Perhaps the idea would be that we should treat suspension along the same theoretical lines as belief because we should favour theoretical simplicity or unity – and so we should then count it as a defect of (BEL-M) that it requires the two attitudes to be treated differently in this specific respect. But of course there are all sorts of differences between the attitudes of belief and suspension, so it is not clear whether this sort of appeal to theoretical simplicity or unity should carry much dialectical force. A couple of such differences that seem particularly relevant are the following:

An obvious and important motivation for *why* we do not want (rational) belief that p to require having a belief about one's evidence for p is that it would *over-intellectualize* belief. We are generally happy to ascribe simple beliefs to young children or to animals for whom it would, I take it, be quite implausible to ascribe more sophisticated beliefs about their evidence being conclusive or inconclusive. Moreover if forming one (rational) belief always required forming a further belief about the evidence, that would seem to set us on the path to a *regress* of endlessly many beliefs about evidence. But these motivations do not seem to carry over to suspending whether p. Firstly, it is not at all clear that we are prepared to ascribe the attitude of suspending judgement to simple believers such as young children or animals – I might describe my toddler or my dog as hesitant or uncertain or just ignorant whether p, but it would sound pretty strange to describe them as agnostic whether p. Agnosticism is plausibly a *more sophisticated*, intellectually demanding attitude than belief – something that (BEL-M) neatly explains since suspending, unlike belief, requires taking a view about one's own cognitive/evidential position. Secondly, the claim that suspending judgement constitutively requires a belief does not raise any parallel worry about an infinite regress of beliefs.

Finally, the champion of (BEL-M) might want to protest that on the contrary it is (BEL-M) that has the virtue of simplicity or theoretical unity, as suspension is treated as being fundamentally a species of belief. And so, on (BEL-M)'s story, we need posit fewer fundamental mental categories and so, perhaps, it can give a simpler/more unified story etc. Now, I am not suggesting that this sort of appeal to simplicity/unity *in favour* of (BEL-M) should carry much dialectical weight either. I am merely pointing out that it has as least as much plausibility as the appeal to simplicity/unity *against* (BEL-M).

5. An argument against (SG)

In this section I present the first, core argument against the sui generis view. I will argue that appealing to a sui-generis kind of propositional attitude seems unable to account for how suspension satisfies constraint (I): S(p) can be rendered more or less rational/justified depending on the state of one's evidence and one's cognitive relation to that evidence.

According to (SG), suspending judgement does not constitutively involve any kind of belief, neither full belief nor partial belief (credence). So the alleged sui generis state would embody neither a committment to, nor a degree of confidence about, the truth or falsity of *anything* – it would be an *alethically uncommitted* state. Whilst the sui generis attitude may have a propositional content, p, it does not involve any commitment or degree of confidence that p is actually the case, *nor that anything else is or is not the case*. Likewise: a desire that p or imagining that p has a propositional content about the world being some way, but the content is not 'asserted', as it were - these mental states do not constitutively involve any commitment about how the world actually is (or is not).

But then why should one's *evidence* have any bearing on the epistemic rationality or justification of such an alethically uncommitted state? Whatever one thinks about the correct ontological category of evidence – mental state, proposition, fact, object etc – evidence *just is* something which indicates what is/is not the case. Other alethically uncommitted propositional attitudes do not appear to have an epistemically rational/justificatory status that is directly connected to the subject's evidence in this way. E.g. Hoping, desiring, imagining etc. are not rendered more/less epistemically rational or justified simply according to one's evidence. In contrast: beliefs, judgements, credences, guesses – these are all attitudes that take some kind of stance on how the world actually is and so it is readily intelligible why evidence bearing on whether something is/is not the case will thereby also bear on the rational/justificatory status of such alethically committed attitudes[18].

Appealing to a sui-generis attitude that is alethically uncommitted thus leaves it wholly mysterious and unexplained *why* suspending judgement can be rendered more/less epistemcially rational/justified by one's evidence.

Here's a first stab at regimenting this into an argument against (SG). Let's assume for the moment, in line with Friedman (2017), that the (SG) theory does *not* require doxastic neutrality as a condition for suspending judgement. The argument would then run as follows:

(1) Sp is an attitude that is rendered more/less epistemically rational/justified according to one's evidence and one's cognitive relation to the evidence.
(2) IF (SG) is the correct account of Sp, THEN Sp is an alethically uncommitted attitude.
(3) IF an attitude is alethically uncommitted, THEN it is *not* rendered more/less epistemically rational/justified according to one's evidence.
(4) SO: (SG) is not the correct account of Sp.

Premise 1 is an assumption which I take to be uncontroversial. Premise (2) seems to just fall out of the core committment of (SG) – that the sui generis state does not require/involve any full or partial belief means that it would not essentially involve any alethic commitment nor any partial alethic commitment (commitment to some degree of confidence). It looks then like the main action comes with premise (3).

Now, a counter-example to the conditional claim in premise (3) comes from alethically uncommitted states that have a content that is *about* an attitude that *is* alethically committed. E.g. intending not to believe that p. Such an intention in itself does not constitutively involve any kind of belief/credence. But you might think that such an intention would be rendered *epistemically* irrational by evidence that very strongly indicates that p; and would be rendered *epistemically* rational by evidence that is inconclusive or that indicates that not-p. And likewise for: desiring not to believe that p, or, hoping not to believe that p, or, mentally

[18] I am assuming, as I take to be widely accepted, that there is *some* important, constitutive link between the notions of truth and justification. E.g. Earl Connee states: 'the existence of some special intimate connection between epistemic justification and truth seems to be beyond reasonable doubt' (Connee, 1992, 657). However, such a link has been denied by Pollock and Cruz (2004).

striving not to believe that p, or, dissatisfaction at the prospect of believing that p etc – you might think that these alethically uncommitted attitudes could be rendered epistemically rational/irrational according to the strength of your evidence for/against p.

However, I think it is clear that these are cases of, at best, derivative or parasitic rationality/irrationality. Intending to Phi is not in general the kind of attitude/state that is rendered epistemically rational/irrational (justified/unjustified) according to one's evidence. It is only when the Phi in question is: holding an attitude that is evaluable as epistemically rational/irrational, that it makes sense to evaluate the intention to Phi in the same way. The intention inherits a 'second-hand' epistemic rationality/irrationality in virtue of being a means to the end of having an attitude that is the real locus of epistemic rationality/irrationality.

And I think it should seem, prima facie, quite implausible that suspending judgement is an attitude that has mere inherited/derivative epistemic rationality in virtue of being an intention or means to bringing about some other attitude, which is the real/original locus of epistemic rationality. When one's evidence is inconclusive one ought to suspend judgement, and the epistemic rationality of suspending in this way does not at all seem to be a matter of whatever further attitudes suspending may later bring about – suspending is already *in itself* the epistemically rational response here.

One way of fixing the previous argument to take account of this:

(1*) Sp is an attitude that is *non-derivatively* rendered more/less epistemically rational/justified according to one's evidence.
(2*) IF (SG) is the correct account of Sp, THEN Sp is an alethically uncommitted attitude.
(3*) IF an attitude is alethically uncommitted, THEN it is *not non-derivatively* rendered more/less epistemically rational/justified according to one's evidence.
(4*) SO: (SG) is not the correct account of Sp.

'Non-derivatively' here means – the rationality/irrationality of the attitude is not in virtue of its aiming or functioning to produce *some other* distinct attitude that is rational/irrational according to evidence. However, contra the non-derivative clause in (1*), Friedman (2017) offers an account of suspending judgement which *does* treat it as essentially involving a kind of intention or aim – the intention or aim of answering a question (and so gaining a true belief). I will discuss and respond to Friedman's account in section 5 below.

Now if, unlike Friedman, (SG) theorists were to accept that suspending judgement *does* constitutively require DOXASTIC NEUTRALITY (in addition to the proposed sui generis attitude) then they may have the materials to satisfy constraint (I). For as we have already seen in the discussion of (NO-BEL), back in section 1, a mere lack of belief/disbelief as a response to evidence could be evaluated as epistemically rational/irrational depending on the strength of that evidence.

However, as discussed in section 3, it is plausible that if one insists on both a belief component *and* DOXASTIC NEUTRALITY as a requirements on suspension, one should also require that these two factors are *connected* in the right way – the meta-cognitive belief should be the *motivating reason* for the doxastic neutrality. Likewise then it seems that if the (SG) theorist were to endorse a view on which suspending required *both* a special, sui generis attitude *and* doxastic neutrality, they should also require that these two factors be connected in the right way. It is not enough that the sui generis attitude is a mere cause of the doxastic neutrality, it would have to be the subject's motivating reason for which she is doxastically neutral. But now the question arises – how could an alethically uncommitted state be an epistemic reason for being doxastically neutral? How could such an alethically uncommitted state be the reason *for which* the subject remains doxastically neutral, such that the subject can thereby be evaluated as epistemically rational or irrational? Or to put this slightly

differently, how could an alethically uncommitted state make the difference between having a genuinely agnostic neutral *opinion* whether p and just deliberately having no opinion whatever whether p?

These are effectively worries about how (SG) can satisfy constraint (II). Of course (SG) does say *something* about how Sp and mere neutral lack of opinion are supposed to differ – the claim is that the former but not the latter essentially involves this special sui generis state (distinct from any belief or credence). But that just raises the question of *why/how* exactly this alleged sui generis state is different from just having no opinion (e.g. how is it different from simply refusing to have any opinion whether p). If the alleged attitude does not constitutively involve any kind of alethic commitments at all (not even about one's evidence or the know-ability of p etc.) then it remains murky what the difference amounts to. Whence the epistemic rational difference between the sui generis state and just deliberately having no opinion whatever? At least until some more details/explanation is supplied, (SG) leaves unexplained what this epistemic difference consists in.

At this point then we should turn from considering the simple/minimal characterization of (SG), given above in section 2, to considering Friedman's (2017) theory, which is the only account thus far in the literature supplying such further details about the nature of the alleged sui generis state.

6. Inquiring Attitudes

In this section I criticise Friedman's (2017) specific version of (SG). Friedman fleshes-out the theory, and tries to account for the rational norms that govern it, by appeal to the idea that suspending judgement necessarily involves an 'inquiring attitude', (or 'IA'), where such an attitude embodies an aim or intention to answer a question. The idea that suspending is fundamentally an attitude to a question rather than a proposition is a very interesting proposal, and Friedman's rich paper deserves more detailed discussion than I can provide here. However, I will argue that appealing to the essential or constitutive involvement of an IA as part of suspending judgement does not help the (SG) theorist to satisfy our two constraints, (I) and (II).

Friedman defends the following biconditional:

> **(S=IA)** One is inquiring into some matter iff one Suspends judgement on the matter. (Friedman, 2017, 302)

Notice first, it is far from obvious that either treating suspension as fundamentally question-directed, or as necessarily connected to having an inquiring attitude, should be incompatible with holding that it constitutively involves some kind of belief. Indeed, Freidman herself seems at points to allow that a kind of meta-cognitive belief may somehow be involved.

Friedman glosses the idea of having an inquiring attitude towards some question, by saying not only that the subject aims/intends to answer the question but also that it "involves a kind of openness and sensitivity to certain kinds of information". A defender of a belief-based account of suspending will naturally wonder whether having an inquiring attitude towards a question necessarily involves some kind of belief – for example: a belief that one's evidence does not currently allow one to answer the question, or perhaps that one does not currently know the answer, or that one ought not yet commit to an answer etc. If it does, then Friedman's account of suspending as tied to inquiry, whatever its other virtues, will effectively be a species of the belief-based view (or at least a hybrid view, one of whose elements is a belief) rather than being a (pure) sui-generis account. And in fact at one point,

Friedman comes close to allowing that suspending judgement does involve a kind of belief or judgement:

> '…we first consider and decide whether we already know [the answer to the question] Q, and if we decide that we don't but somehow want to, then we suspend and open the question up for genuine investigation.' (Friedman, 2017, 314)

Although Friedman here talks of 'deciding' that we don't already know the answer to the question, this is presumably a kind of judgement or belief that we cannot tell or don't know the answer[19].

When it comes to the rational/irrational status of this alleged Suspending-as-Inquiring-Attitude, Friedman states the following:

> "**Ignorance Norm for the IAs**: Necessarily, if one knows Q at t, then one ought not have an IA towards Q at t." (ibid., 311)

Notice again: if we were to assume that a genuinely inquiring attitude to Q *does* constitutively involve a belief – such as "I cannot currently tell the answer", or the 'the evidence does not currently allow me to answer Q" – then inquiring into Q whilst knowing the answer would amount to a kind of Moorean combination of beliefs: Bp & B[I don't know whether p], or, Bp & B[my evidence does not yet allow me to tell whether p]. Thus we could have an explanation for the 'normative pressure' to avoid having an I.A. towards a question to which one knows the answer[20]. More generally, much of what Friedman claims about the normative profile of Inquiring Attitudes is very naturally explained, I suggest, by the assumption that such IAs require some kind of *belief* whose content is that one is not currently able to answer the question. But of course this would no longer be a (pure) Sui Generis account.

On the other hand if having a genuinely inquiring attitude does *not* require any such belief about one's current epistemic position, then the proposal will count as a (pure) sui generis theory. But, as I will now argue, it then becomes much less clear why having such an inquiring attitude should be rendered irrational/rational by one's evidence, or by what one knows – i.e. it is not clear how it can satisfy (I). It is also simply implausible, I will argue, that such a belief-free I.A. is an essential/constitutive requirement on being agnostic – i.e. it also seems unsuited to satisfying (II)[21]. Let's proceed by considering both directions of the bi-conditional (S=IA)

Taking the left-to-right direction (i.e. IA→S) first: there seem to be possible cases in which a subject does not suspend judgement but rather fully believes that p, indeed knows that p, and yet she might still *rationally* go about gathering evidence and making further inquiries. E.g. because she wants to make the strongest case possible in order to persuade others, or perhaps because she remains interested in proving the case for p in lots of different ways.

Friedman is aware of this kind of potential counter-example. In response she suggests that if a *rational* subject continues to act and behave as if involved in genuine inquiry into Q, even though she already knows the answer to Q, 'what she's doing looks like a sort of ersatz

---

[19] Note - Friedman wants to say that this 'decision' as to whether we already know the answer to Q occurs as part of a *pre*-suspension/*pre*-inquiry stage in the process. But it is not clear why considering whether we already know the answer to Q should not count as genuine inquiry, but only as pre-inquiry.

[20] It should be noted that some recent theorists – e.g. Lasonen-Aarnio (2014), Elga (2013), Sliwa & Horowitz (forthcoming) – have been prepared to deny that these Moorean combinations are always irrational.

[21] See also Archer (2018) for criticisms of Friedman's proposals that overlap, to some extent, with the criticisms I make here.

inquiry into Q at best, its not genuine inquiry'. I.e. inquiring into Q when one already knows the answer must either be irrational or not 'genuine' inquiry. Now, were we to stipulate that a 'genuine' inquiring attitude essentially requires some kind of *belief* that one does not currently know the answer, this might make sense – for perhaps a rational subject could not both B(p) and B(I cannot tell whether p), so such a rational subject could only be making ersatz as opposed to genuine inquiries. But in the absence of such a belief-based requirement on 'genuine inquiry', why should a rational subject who genuinely aims and acts to gather more information relevant to answering question Q, who is genuinely open and sensitive to information relevant to Q, not count as *genuinely inquiring*, even if she already knows the answer to Q? Friedman, as far as I can see, supplies no reason.

Nor is it obvious why Inquiring Attitudes, if they do not essentially involve any meta-cognitive belief, should be rendered irrational by what one knows (or one's evidence). Recall the comparison with other alethically uncommitted attitudes in the previous section – there is no inherent epistemic/rational clash between believing that p and imagining that p, or believing that p and hoping that not-p, etc. So why should there be any such clash between an *alethically uncommitted* inquiring attitude concerning p – involving, say intentions to gather evidence relevant to p, etc. – and a belief that p?

In support of her ignorance norm, Friedman appeals to the rational inappropriateness of incessantly double-checking what one already knows in the absence of any new reason for doubt:

> "While double-checking might seem fine sometimes, the incessant double-checker is very far indeed from the epistemic ideal.' (Friedman, 2017, 312)

However, whether this is really a matter of *epistemic* rationality depends, I suggest, on whether the double-checking inquiry involves a genuine *doubt* – i.e. a *belief* or *judgement* that perhaps one does not know after all, or that the evidence is not after all so strong, or at least a reduction in *credence*. If there is such a doxastic component to the double-checking inquiry, then we can readily make sense of the idea that it would be *epistemically* deficient – for there is, by hypothesis, no new epistemic reason to doubt one's knowledge or to lower one's credence. But if the double-checking inquiry does *not* essentially involve any such belief, and is rather just a pure intention or aim to check again, then it is not at all clear why there should be any kind of *epistemic* failing. Any inappropriateness would seem to be a matter only of *practical rationality*. For consider, if we massively increase the downside for being wrong, then any amount of double-checking can come to seem entirely appropriate, given the high stakes and risks involved. Or compare: is there anything *epistemically* wrong with endlessly working through a certain mathematical proof again and again? Of course this may well be *practically* irrational and pointless (or maybe not if one gets endless aesthetic pleasure from it) – but why should an inference that is *epistemically* rational to draw on the basis of some premises on the first occasion become *epistemically* irrational when the subject draws the same conclusion on the same basis for the nth time?

Thus the appeal to Inquiring Attitudes – once we are clear that they do not essentially involve any kind of belief – do not seem to be able to account for constraint (I) on suspension. Moreover, since one can apparently have a such an alethically uncommitted inquiring attitude without being agnostic, IAs do not seem to help with constraint (II) either – for they do not seem to mark out the distinctive difference between a genuine agnostic opinion whether p and just having no opinion whatever whether p.

Turning now to the right-to-left direction (i.e. S→ I.A.) there seems to be a whole range of cases in which a subject *does* suspend judgement about some question, Q, despite having no aim or intention whatever to answer Q. For example: one might simply not care in the slightest about the Q – e.g. is the total number of hairs in my left nostril odd or even? Or one

might believe that there simply is no available evidence, or indeed that there *couldn't b*e any evidence, that could answer Q – e.g. did the universe pop into existence 10 minutes ago, complete with fake fossils and fake memories etc.? I suspect that many people who self-identify as agnostics about the existence of God would fall into this category – they have no expectation that there is any humanly possible route to definitively answering the question.

Then there are also somewhat more exotic cases of *anti-expertise* that are, I think, even worse for the right-to-left direction of (S=IA). A subject, S, is an anti-expert concerning some proposition, p, when the following condition holds: S believes that p iff not-p. For example:

> "I'm a neurologist, and know there's a device that has been shown to induce the following state in people: they believe that their brains are in state S iff their brains are not in state S. I watch many trials with the device, and become extremely confident that it's extremely reliable. I'm also confident that my brain is not in state S. Then the device is placed on my head and switched on. My confidence that my brain is not in state S ..... well, it's not clear here what should happen here." (Christensen, 2010, based on Conee, 1982)

Conee has argued that in such a case, where S knows or at least reasonably believes she is an anti-expert about p, then the rational best (or least worst) option is to Sp. Whether or not Conee is ultimately right about this[22], it seems pretty clear that faced with this kind of tricky situation I *might* end up suspending judgement (whether rationally or not), and yet I would be convinced that I *could not* successfully answer the question: 'whether p?', precisely because I'm convinced that I am an anti-expert about p.

In all of these cases, the subject can apparently suspend judgement despite lacking any kind of aim or intention to answer the question being suspended over, and so lacking any expectation of ceasing to suspend Judgement at some later date. Indeed, recall that Rosenkranz (2007) considers this kind of *settled suspension* to be 'True Agnosticism' (see section 3, above). Once more then, given that we can suspend without having an IA it seems that IAs are not what explains the difference between genuine agnosticism and mere lack of any opinion – i.e. they do not help with constraint (II).

Now again, Friedman is aware of at least some of these kinds of apparent counter-examples to the right-to-left direction of (S=IA). And she offers two possible strategies in response.
Either: (i) claim that the subject lacking an I.A. is not really suspending any more but is merely doxastically neutral and so lacking any genuine opinion on the matter.
Or: (ii) claim that even though the subject is not doing anything much in the way of actively inquiring she *is* still open and sensitive to the possibility of new information, and thus can still count as suspending.

But whilst response (i) may be plausible for cases where the subject simply doesn't care about the question ('is the number of hairs in my left nostril odd or even?'), it is not at all plausible that we can dismiss cases such as the settled, weary agnostic about 'does God exist?' or the philosophical agnostic about 'did the earth pop into existence 10 minutes ago?', or the Rosenkranz-style 'True Agnostic', as subjects who are not really suspending judgement and who have merely 'shrugged off' the question. These subjects think that the question is genuine and has an answer but that it is impossible for them ever to reach that answer. Again, they suspend *not* with the intention of later achieving some other epistemically superior state; rather, they suspend precisely because they think suspension is the best that can be hoped for.

In reply to (ii): the problem with lowering the bar for being in an inquiring state of mind in this way is that we would then all presumably count as being in an inquiring state of mind

[22] See Sorensen (1987) for criticism of Conee's proposal.

about virtually any/every proposition we consider, including all/many of those we currently know or outright believe/disbelieve – which would bring us all into constant conflict with Friedman's purported ignorance norm. For it is very plausible that in some minimal sense of 'being sensitive to new evidence', we are all still sensitive to possible new evidence whether p, even when we Kp – and this seems entirely rational. After all, think of all those much-discussed cases where an initial state of Kp is *defeated* by new misleading evidence – presumably we think that subjects in such cases *ought* to be open and sensitive to such new defeating evidence. Or consider Kripke's 'dogmatism paradox'[23] – the conclusion of which is that once I Kp, I can then rationally disregard any possible future evidence that seems to indicate not-p. The reason that most of us find this conclusion paradoxical/implausible is presumably that we intuitively take Kp to be *rationally consistent* with continuing sensitivity/open-ness to possible future evidence that not-p. But, if we interpret what it takes to be in an inquiring attitude in a sufficiently liberal way, requiring only a minimal sensitivity to possible future evidence, the proposed ignorance norm would implausibly condemn us all for even this minimal sensitivity to future possible evidence.

In summary then, we have the following dilemma for any (SG) theorist hoping to appeal to Inquiring Attitudes in order to explain the two constraints on suspension:

- *If* having an inquiring state of mind essentially involves some kind of belief that the evidence is inconclusive or that the question is not yet answered etc, then the account becomes a belief-based view[24].
- But: *If* having an inquiring state of mind does *not* essentially involve any kind of belief, then it does not help with either constraint, (I) or (II). It is quite unclear why merely having such a pure intention or aim to continue gathering evidence – i.e. a *non*-alethic commitment – should *epistemically* conflict with Kp, or should be rendered more or less *epistemically* rational depending on the state of one's evidence. Nor is it plausible that such an alethically uncommitted IA is either necessary or sufficient for Sp.

The fundamental problem with giving an account of suspension in terms of having a certain *intention* or *aim* to inquire further and eventually answer a question, comes back to the point made at the end of section 5: suspending whether p, like believing or disbelieving p, is an attitude that can *in itself* be rendered rational/irrational according to one's evidence (and one's relation to the evidence). It is not rational/irrational only by virtue of its having an aim or intention to bring about some *other* future attitude that is the real locus or source of rationality. It can be rational simply because the evidence demands (or permits) it. At the end of her paper Friedman writes:

> "We suspend judgment in order to get to the truth or to know (and the like). If we want to move ourselves from a state of ignorance to a state of knowing by inquiring, then at least two things need to happen: a question needs to be opened or asked and then it needs to be closed in the right sort of way or answered. The point or purpose or aim of opening is typically the very same point or purpose or aim of closing when the circumstances are right: to improve our epistemic standing on some matter – to settle a question and to come to know." (ibid., 321-322)

But it is just not plausible that the fundamental point or aim of suspending is in order to get to the truth – for we also suspend so as to *avoid falling into false belief.* And as is illustrated most clearly with the case of anti-expertise, above, we might sometimes suspend judgement

[23] Though credited to Kripke, the paradox is first reported by Harman (1973, 148)

[24] Though, to repeat, that would certainly *not* mean that Friedman's rich account of Inquiring Attitudes would then be of no further interest!

without any further purpose or aim of ever improving our epistemic standing on the matter, but rather in the full expectation that suspending will forever remain the best epistemic standing that we can achieve.

7. In Favour of (BEL-M)

Having presented my case against (SG) – both in it's generic/minimal form, and in the specific form that Friedman (2017) advocates – in this section I will now briefly state how (BEL-M), unlike (SG), is able to meet the two constraints.

Let's start with constraint (II). According to (BEL-M), the state of suspending judgement *does* still essentially involve some kind of alethic commitment. For whilst suspending whether p does not, of course, involve any commitment about the truth or falsity of p, it does, according to (BEL-M), constitutively require a belief that one cannot currently tell whether or not p. (BEL-M) thus gives a very straightforward explanation for how/why suspending judgement differs from mere, non-opinionated, doxastic neutrality – S(p) involves taking a stance about what one can currently tell given one's evidence.

What about constraint (I)? Well, given the requirement that DOXASTIC NEUTRALITY is still part of the complex state of suspending judgement together with the meta-cognitive belief (and the further requirement that the neutrality and the belief are connected in the right way), such doxastic neutrality about p can very plausibly be rendered more rational (at least to some degree) by weak, inconclusive evidence and be rendered less rational (at least to some degree) by evidence that overall points strongly either to p or to not-p. [For such neutrality will be the result of a sensitivity to the evidence – given that it is connected to the belief about what on can currently tell based on the evidence.]

But of course the belief component of the (BEL-M) is also something that can be rendered more or less rational/justified by one's *total* evidence – that is *including* any higher-order evidence about one's first-order evidence concerning p and any evidence one has about one's current cognitive state and relation to the evidence. For it is clear that such evidence of one's evidence about p, and evidence about one's current cognitive ability to tell what this evidence indicates about p, will rationalize/justify a belief that: one cannot currently tell whether p.

Thus, the overall complex state of suspending judgement whether p – consisting of both the meta-cognitive belief and the connected doxastic neutrality – is a state that can clearly be rendered more or less rational/justified, at least to some extent, according to one's evidence. And it can also do justice to the fact, mentioned right at the start, that doubt can be reasonable not only due to the state of one's evidence but also due to one's current cognitive relation to that evidence.

These are, I think, very significant advantages for (BEL-M) over (SG).

Now, as soon as one allows that a doxastic attitude (whether belief or suspension) can be rationalised/justified to some extent *both* by one's 1st-order evidence *and* by one's higher-order evidence of that evidence, the possibility opens up that these two factors might, as it were, be pulling in opposite directions. For example, one could have (misleading) higher-order evidence that one's 1st-order evidence whether p does not indicate what it in fact does indicate. Or one could have higher-order evidence that one has taken some kind of mind-altering drug, which affects one's ability to evaluate 1st-order evidence. Or perhaps one could just lack introspective/reflective access to one's 1st-order evidence. Thus we can have two kinds of cases:

- 1st order evidence strongly indicates that p and so, in itself, points towards Bp, *but* H.O. evidence (i.e. evidence about the strength of 1st-order evidence whether p and/or evidence about one's current cognitive state) indicates that the 1st order evidence and one's current cognitive relation to it does not allow you to tell whether p and so points towards Sp – at least from the subject's point of view.

- 1st order evidence is overall weak and inconclusive whether p and so, in itself, points towards Sp, *but* H.O. evidence indicates that the 1st order evidence and one's cognitive relation to it strongly indicates that one can indeed tell that p, and so points towards Bp – at least from the subject's point of view.

In these kinds of cases, where 1st order and higher-order evidence conflict, it is not clear whether the complex agnostic state – consisting of both doxastic neutrality about p and the meta-cognitive belief that one cannot tell whether p – will be rational/justified or not.

Is this bad for (BEL-M)? I think not.

For I think that once we grant the possibility of such cases, it creates tricky problems for *everyone*. All parties must accept that in cases such as (I) the subject has *both* some evidence that, in itself, would tend to rationalise/justify Bp *and* some evidence that, in itself, would tend to rationalise/justify B(my evidence does not support Bp). And in cases such as (II), the subject has evidence that in itself tends *not* to rationalise/justify Bp and evidence that in itself tends to rationalise/justify B(my evidence supports Bp). So all parties will have to accept that subjects can be put in the rather Moore-Paradoxical situation of having total evidence which simultaneously bears in these rationally conflicting directions – both towards believing that p and towards believing that one cannot tell whether p (or that one ought not believe that p). How to weigh up or adjudicate between the conflicting rational bearings of 1st order and higher-order evidence is then a further, very deep and difficult question[25] – which I will not attempt to answer here!

But there is no apparent reason here to think that (BEL-M) *in particular* faces some kind of especial difficulty with these cases. It is not obviously any *more* counter-intuitive to allow, as (BEL-M) does, that in these cases of 1st-O and HO conflict, a subject can have *some degree* of justification/rationalistion to Bp and *some degree* of justification/rationalisation to Sp. Indeed it is a point in favour of the present theory that it's two components – the meta-cognitive belief and the connected state of Doxastic Neutrality – are at least able to account for the existence of both of these conflicting rational 'forces'. Thus the view can account for why, say, in a case where one lacks reflective/introspective access to one's 1st-order evidence, this would be at least *some reason* in favour of suspending judgement. For the lack of reflective/introspective access will tend to rationalize/justify a belief that 'I cannot currently tell whether or not p based on my evidence'.

8. Conclusion

I have argued that a belief-based account of suspending judgement, (BEL-M), offers the best explanation of how suspension is an attitude that can be rendered epistemically rational/justified (at least to some extent) by one's evidence and how/why it counts as a genuine opinion on the question whether p. And I have argued that treating suspension as a sui generis attitude that is entirely distinct from belief (full or partial) leaves it quite mysterious how/why evidence can render the attitude more or less epistemically justified, more or less epistemically rational. Nor is it at all clear how such an alethically uncommitted

[25] The literature on Higher-Order evidence has exploded over the past few years. Some important papers include: Christensen (2010), Lasonen-Aarnio (2014), Titlebaum (2015), Horowitz (2014).

state could be what makes the difference between an agnostic's genuine neutral opinion whether p and merely having no opinion whatsoever.

Now, I do not pretend to have demonstrated that it is *impossible* for a sui generis theorist to provide some further, alternative explanation for how/why one's evidence can render suspending more or less justified/rational. Back at the start of section 2, I briefly considered the possibility of rejecting constraint (I) and insisting instead that suspension is an attitude that is not subject to the demands of *epistemic* rationality at all, only practical rationality applies. This might provide a bullet-biting option for the (SG) theorist.

Another possibility: I suppose that an (SG) theorist might try to insist that the link between evidence and the justification/rationality of suspending judgement is just a *brute* normative fact, entirely distinct from and unconnected with the way that evidence bears on the rationality/justification of belief (or other alethically committed attitudes). After all, if suspension is a sui generis state, entirely distinct from belief, perhaps it is subject to its own sui generis epistemic norms?

I don't have anything much to say in response to this sort of appeal to brute, sui generis norms except to repeat that it seems to leave it entirely unexplained *why evidence* – that which indicates truth or falsity – should bear on the epistemic rationality of *this* specific alethically uncommitted attitude but does *not* so bear on the epistemic rationality of other alethically uncommitted states, such as imagining that p, hoping that p etc. To insist that this is just a brute normative fact is effectively to concede that there is no explanation. Likewise, this kind of move apparently offers no illuminating explanation as to what the difference between suspending judgement and merely having no opinion (doxastic neutrality) consists in.

Of course, explanations must end somewhere and perhaps all theories are bound to have some kind of brute or primitive element in their story. But I hope that those readers who have followed my presentation of the belief-based account this far would accept that the sort of *intelligibility* that (BEL-M) provides as to *why* suspending is an attitude that can be rendered more or less justified by your evidence, and how it differs from mere doxastic neutrality, is much to be preferred to any appeal to special, brute, sui generis norms[26].

[26] As well as being presented at the conference on Reasonable Doubt at Birkbeck College, material from this paper was presented at events in Oslo, Barcelona, St. Andrews, Cambridge, Cologne, and at meetings of the European Epistemology Network in Madrid and in Paris. I am grateful to the audiences on all those occasions for their feedback. By far my biggest debts of gratitude are to Anna-Maria Eder and to Julien Dutant, both of whom generously provided me with extensive written comments which really transformed this paper and which saved me from a number of embarrassing mistakes. Of course they should not be assumed to endorse any of the views advanced here, nor are they responsible for any errors and infelicities that remain. Many thanks also to Bill Brewer, Peter Brössel, Jessica Brown, Jonathan Knowles, Clayton Littlejohn, Anders Nes, Carlota Serrahima, Justin Snedgar, Jan Willem Wieland and Timothy Williamson for helpful advice and criticism. I am very grateful to two anonymous referees for this journal, whose insightful reports substantially improved this paper. Finally, I would like to thank Jane Friedman, whose fascinating papers on suspending judgement originally led me into this topic.

References

- Adam Carter, J. & Bondy, P. (eds.) (forthcoming) *Well-Founded Belief: New Essays on the Basing Relation,* London: Routledge.
- Alston, W. (1989) *Epistemic Justification: Essays In The Theory Of Knowledge,* Ithaca, NY: Cornell University Press.
- Archer, A. (2018) 'Wondering about what you know', *Analysis* 78(4): 596–604.
- Audi, R. (1986) "Belief, Reason and Inference", *Philosophical Topics* 14(1): 27-65.
- Bergmann, M. (2005) 'Defeaters and higher-level requirements'. *Philosophical Quarterly* 55: 419–36.
- Booth, A. (2012) 'Two reason why epistemic reasons are not object-given reasons' *Philosophy and Phenomenological Research* 89 (1): 1-14
- Christensen, D. (2009) 'Disagreement as Evidence: The Epistemology of Controversy'. *Philosophy Compass*, 4(5): 756–767.
- Christensen, D. (2010) 'Higher-Order Evidence,' *Philosophy and Phenomenological Research* 81 (1): 185-215.
- Conee, E. (1982) 'Utilitarianism and Rationality' *Analysis* 42: 57
- Conee, E. (1992) 'The Truth Connection', *Philosophy and Phenomenological Research* 52 (3): 657-669.
- Crawford, S. (2004) 'A solution for Russellians to a puzzle about belief'. *Analysis*, 64, 223–229.
- Davidson, D. (1980) *Essays on Actions and Events*. Oxford: Clarendon Press.
- Decker, J. (2012) "Disagreement, Evidence, and Agnosticism" *Synthese* 187 (2):753-783.
- Elga, A. (2007) "Reflection and Disagreement". *Noûs* 41 (3):478–502.
- Evans, G. (1982) *The Varieties of Reference*. (ed.) McDowell, J. Oxford: Oxford University Press
- Evans, I. (2013) 'The Problem of the Basing Relation' *Synthese* 190(14): 2943-2957.
- Feldman, R. (2007). "Reasonable Religious Disagreement" in L. Antony ed. *Philosophers Without God.* Oxford: Oxford University Press. 197-214.
- Fitelson, B. (2012) "Evidence of Evidence is Not (Necessarily) Evidence', *Analysis*, 72 (1): 85-88
- Friedman, J. (2013a) 'Suspended Judgment'. *Philosophical Studies* 162 (2):165-181.
- Friedman, J. (2013b) 'Rational Agnosticism and Degrees of Belief'. *Oxford Studies in Epistemology* 4:57.
- Friedman, J. (2017) 'Why Suspend Judging?', *Nous*, DOI 10.1111/nous.12222
- Gibbons, J (2013) *The Norm of Belief*. Oxford: Oxford University Press.
- Hájek, A. (1998) 'Agnosticism meets Bayesianism'. *Analysis*, 58:199–206.
- Horowitz, S. (2014) 'Epistemic Akrasia', *Noûs* 48 (4): 718-744.
- Joyce, J. (2005) "How Probabilities Reflect Evidence", *Philosophical Perspectives*, 19: 153–178.
- Kelly, T. (2014) "How to Be an Epistemic Permissivist" in Steup, M. and Turri, J. eds. *Contemporary Debates in Epistemology*, 2nd edition. Malden, MA: Blackwell Publishing.
- Korcz, K. (1997) 'Recent Work on the Basing Relation', *American Philosophical Quarterly* 34 (2): 171 - 191.
- Lasonen-Aarnio, M. (2014) 'Higher-Order Evidence and the Limits of Defeat', *Philosophy and Phenomenological Research* 88 (2): 314-345.
- Lehrer, K. (1971) 'How Reasons Give us Knowledge, Or The Case of the Gypsy Lawyer', *Journal of Philosophy* 68(10): 311-313
- Lehrer, K. (1990) *Theory of Knowledge*, Boulder, CO: Westview.
- McHugh, C. (2012) 'The Truth Norm of Belief', *Pacific Philosophical Quarterly* 93: 8–30
- Nelson, M (2010) 'We Have No Positive Epistemic Duties'. *Mind* 119 (473): 83-102.
- Neta, R. (forthcoming) 'The Basing Relation', *The Philosophical Review*.
- Pollock, J. & Cruz, J. 'The Chimerical Appeal of Epistemic Externalism' in Schantz, R. (ed.), *The Externalist Challenge*. De Gruyter. 125-42.
- Rosenkranz, S. 'Agnosticism as a Third Stance'. *Mind* 116 (461): 55-104.
- Sturgeon, S. (forthcoming) 'Epistemic Attitudes'.

- Sorensen, R. (1987) ‘Anti-Expertise, Instability, and Rational Choice’, *Australasian Journal of Philosophy* 65 (3): 301 – 315
- Sorensen, R. (2009) ‘Meta-Agnosticism: Higher Order Epistemic Possibility’, *Mind* 118 (471): 777-784.
- Swain, M. (1979) "Justification and the Basis of Belief", in Pappas (ed.) *Justification and Knowledge*, Dordrecht, Holland: D. Reidel.
- Tang, W. (2016) ‘Reliabilism and the Suspension of Belief’, *Australasian Journal of Philosophy* 94 (2):362-377.
- Titelbaum, M. (2015) ‘Rationality’s Fixed Point’. *Oxford Studies in Epistemology* 5: 253–294.
- Turri, J. (2009) ‘The Ontology of Epistemic Reasons’. *Noûs* 43 (3): 490-512.
- White, R. (2005) “Epistemic Permissiveness.” *Philosophical Perspectives* 19 (1): 445–459.
- White, R. (2010) “Evidential Symmetry and Mushy Credence” in T. Szabo Gendler and J. Hawthorne eds. *Oxford Studies in Epistemology Vol.3*. Oxford: Oxford University Press.