

A New Argument against Uniqueness.

Thomas Raleigh

[Forthcoming in *The Philosophical Quarterly* – this is an uncorrected draft, so please do not quote from or circulate this version!]

ABSTRACT: I present a novel argument against the thesis of Uniqueness and in favour of Permissivism. Counterexamples to Uniqueness are provided, based on ‘Safespot’ propositions – i.e. a proposition that is guaranteed to be true provided the subject adopts a certain attitude towards it. The argument relies on a plausible principle: (roughly stated) If S knows that her believing p would be a true belief, then it is rationally permitted for S to believe p. One motivation for denying this principle – viz. opposition to ‘epistemic consequentialism’ – is briefly discussed.

1. Introduction

The recent literature on peer disagreement is a nice example of epistemologists tackling an issue that clearly crops-up ‘in real-life’, the relevance of which extends beyond the esoteric interests of professional philosophers: viz. if you find yourself disagreeing with someone over some question, where you know that you both have exactly the same evidence and the same level of expertise and intelligence to answer the question, should you rationally¹ feel less confident in your own answer in light of this disagreement? An important further issue arising from this debate on peer disagreement is the following: for any set of evidence, is there always exactly one attitude that is the rational attitude to take to a proposition in light of that evidence? Or can there sometimes be evidence in light of which two (or more) attitudes towards some proposition are equally rational? Though the current interest in this latter issue was sparked by the debate on peer disagreement, it is clearly of crucial importance to the enterprise of rational belief formation in general, whether we focus on a single individual or on an inter-subjective situation. Following Ramsey (1931), it is often said that our beliefs form a kind of internal map by which we navigate the world². Put in terms of this map metaphor, the question is whether for any set of evidence there is always a uniquely rational way to draw your map in response, or whether there can

¹ I assume throughout that we are concerned with epistemic rather than practical rationality (or any other kind of rationality). I also assume that we have a clear enough grip on this distinction that we can proceed without trying to define it precisely. Though see section 5 for some difficulties with one common way of trying to draw the distinction.

² Ramsey actually said that the belief that ‘everybody in Cambridge voted’ is ‘a map of the neighbouring space by which we steer’ (Ramsey, 1931, 237).

sometimes be two (or more) equally rational maps that you could draw in response to your evidence?

Take UNIQUENESS³ to be the following thesis:

- For any subject S, proposition p and set of evidence E, *exactly one* of the 3 doxastic attitudes to p – Belief, Disbelief⁴ or Suspension – is rationally permitted for S on the basis of E.

Take A.B.U. (At Best Unique) to be the following thesis:

- For any subject S, proposition p and set of evidence E, *at most one* of the 3 doxastic attitudes to p is rationally permitted for S on the basis of E.

The negation of A.B.U. then is PERMISSIVISM⁵:

- It is possible that there could be some subject S, proposition p and set of evidence E such that *more than one* of the 3 doxastic attitudes to p is rationally permitted for S on the basis of E.

As they stand, these might be thought somewhat imprecise formulations in (at least) a couple of respects. Firstly, in addition to belief, disbelief and suspension, there is also, you might think, a fourth possible option of forming no opinion whatsoever – or ‘withdrawing’⁶. Secondly, one might wonder whether it is best to formulate these theses in an inter-subjective or intra-subjective way⁷ – i.e. if UNIQUENESS is true,

³ Recent advocates of Uniqueness include: White (2005), Christensen (2007), Feldman (2007), Sosa (2010)

⁴ I assume, as I take to be standard, that disbelieving a proposition is equivalent to believing its negation. I.e. $DBp = B\neg p$. However, see Sturgeon (forthcoming) for a denial of this equivalence.

⁵ Recent advocates of Permissivism include: Douven (2009), Kelly (2014), Schoenfield (2012).

⁶ The term ‘withdrawing’ is used by Turri (2012). In fact, it is not clear to me that refusing to adopt any of the 3 doxastic attitudes, so not even suspending judgement, is a state of mind that is subject to the demands of theoretical/epistemic rationality (as opposed to practical rationality). It may be a more or less prudent option to ‘withdraw’, but just refusing to think any further about a certain proposition is not obviously evaluable at all as an intellectual or theoretical move that could be rationally correct or incorrect in light of the available evidence.

⁷ For discussion of the distinction between inter-subjective and intra-subjective versions of these theses, see Kelly (2014).

might it nevertheless be that different subjects with the same evidence could each be permitted to have a *different* uniquely permitted attitude to p ? However, these complications will make no difference in what follows. The argument I will present, if sound, shows that a single subject can be permitted on the basis of her evidence both to believe or to disbelieve a proposition. No matter how advocates of UNIQUENESS or A.B.U. want to incorporate the possibility of ‘withholding’ into their preferred thesis, this is still going to be a problem. And it looks extremely plausible that if a single subject is permitted both to believe or to disbelieve a proposition, then *a fortiori* 2 different subjects with the same evidence could be permitted to each adopt different doxastic attitudes to it. For ease of presentation then, we can safely leave the three theses above as they stand.

At the risk of stating the obvious, for PERMISSIVISM to be true, all that is required is one counterexample to A.B.U. I will argue that, given a plausible assumption, counterexamples can be formed, based on the kind of Moorean propositions that Roy Sorensen (1988) labelled ‘Blindspots’⁸ or on their opposite kind, which I will label ‘Safespots’.

- An attitude A-blindspot for some subject S is a proposition that can be true and that S can take some propositional attitude, A, towards; but not both – i.e. it is bound to be false if S adopts attitude A to it.
- An attitude A-safespot for some subject S is a proposition that can be false and that S can take some propositional attitude, A, towards; but not both – i.e. it is bound to be true if S adopts attitude A to it⁹.

Some examples:

⁸ Bykvist & Hattiangadi (2007) argue that blindspot propositions provide a counter-example to the following ‘truth-norm’ for beliefs: For any S, p : if S considers p , then S ought to (believe that p) iff p is true.

⁹ Just to be absolutely clear, and again risking statement of the obvious, as I have defined them safespots are propositions that *can be false* and blindspots are propositions that *can be true* – so neither necessary truths nor necessary falsehoods count as blindspots or safespots. Also, safespots and blindspots are propositions towards which the subject *can* take the relevant propositional attitude – so propositions that are just too long and complex for a subject to ‘get in mind’ are ruled out too.

The proposition: ‘It is raining & S does not believe that it is raining’ is a belief-blindspot for S. (I assume here that it is possible to both Believe that p and Believe that one does not Believe that p.) But it is not a hope-blindspot nor a desire-blindspot. And, on the assumption that it is possible for a subject to both disbelieve a proposition and disbelieve that they disbelieve it, nor is this a disbelief-blindspot.

Conversely, ‘ $2+2=4$ & S believes that $2+2 = 4$ ’ is a belief-safespot for S, but not a disbelief safespot for S. (Nor, of course, a hope-safespot or a desire-safespot etc.)

Whereas, the proposition: ‘S is dead’ is both a belief-blindspot and a disbelief-blindspot for S (and a hope-blindspot and a desire-blindspot etc.) And conversely, ‘S is alive’ is both a belief-safespot and a disbelief-safespot – and so on for other attitudes – for S. (I assume here that only the living can have propositional attitudes.)

Notice also that there can be more specific *kinds* of belief-blindspots and belief-safespots. E.g. a proposition of the form: ‘p & S cannot justifiably believe that p’ is not a belief-blindspot for S, but it is a *justified-belief-blindspot* for S. (It can be truly though unjustifiedly believed by S, but it cannot be truly and justifiably believed by S.)

2. The Key Assumption

The existence of Safespots (and Blindspots) becomes a problem for A.B.U. once we grant an apparently very plausible principle, which I will initially state in a slightly over-simple form:

- **PRINCIPLE, version 1.0:** If S knows that her believing p would be a true belief, then it is rationally permitted for S to believe p.

Or in semi-formal terms:

- **PRINCIPLE 1.0:** $[SK (SBp \rightarrow p)] \rightarrow \text{Rationally Permitted: } SBp$

For the time being I will provide no support for this principle apart from simply repeating that it is, I take it, highly plausible and intuitive¹⁰. After all, if you *know* that your forming a belief in a particular proposition is guaranteed to result in a true belief, it looks like you have an exemplary rational basis to hold that belief. Of course, that you *have* a rational basis to believe some proposition, does not guarantee that an actual belief you've formed in that proposition *is* a rational belief, for you might have formed the belief on some other irrational basis.

However, as it stands this initial version 1.0 of PRINCIPLE is not quite right, for although knowledge that a belief that p is bound to be true seems to be an excellent rational basis on which to believe that p, this rational basis could be undermined by (misleading) *higher-order evidence* that this knowledge is not a rational basis or that one lacks such a basis, or by higher-order evidence that you are unable to evaluate what is and what isn't a rational basis for belief. Notice – the sort of higher-order evidence we're now considering would *not* be evidence which tells against the truth of the conditional: If S were to Believe p, then the belief would be true. So it is not evidence that would necessarily undermine S's knowledge that this conditional is true. Rather, this sort of evidence would tend undermine the subject's ability to *use* such knowledge as a rational basis for believing that p.

E.g. you have (false but totally plausible) expert testimony, from some super-intelligent and trustworthy authority on rationality, that your knowledge that your believing that p would bound to be a true belief is *not* a rational basis to believe that p. Or perhaps you have evidence that you have taken a special drug, which selectively impairs your ability to grasp whether the rational basing relation obtains between evidence and beliefs. Etc. Now the topic of higher-order evidence and whether/when it can undermine 1st-order evidence is both fascinating and very difficult¹¹. But it is at least very plausible that in these sorts of cases, the subject would *not* have rational

¹⁰ In section 5 I sketch some unattractive consequences of denying the principle. And in section 6 I show how PRINCIPLE can be thought of as a special case of two famous constraints on rational degrees of belief.

¹¹ See Titlebaum (2014) or Lasonen-Aarnio (2014) for recent arguments denying that higher-order evidence about what one's first-order evidence rationalizes/justifies can undermine one's first-order justification/rationalization. On the other hand, Gibbons (2013) argues that higher-order evidence can so undermine first-order evidence.

permission to believe that p, despite possessing knowledge that such a belief is guaranteed to be true, knowledge that would normally be an ideal rational basis for such a belief, as this basis would here be undermined by the higher-order evidence.

A slightly more complicated but careful formulation of PRINCIPLE then is the following:

- **PRINCIPLE version 1.1:** If S knows that her believing p would be a true belief *and in the absence of any undermining higher-order evidence*, then it is rationally permitted for S to believe p.

Or in semi-formal terms:

- **PRINCIPLE 1.1:** $[SK (SBp \rightarrow p)] \ \& \ \neg(S \text{ has undermining H.O. evidence}) \rightarrow$
Rationally Permitted: SBp

Lets now consider a more specific version of the principle that deals with beliefs *of a specific type* – which, once more, I will initially state in a too-simple formulation:

- **PRINCIPLE version 2.0:** If S knows that a belief *of some specific type* that p (by S) is bound to be a true belief (and in the absence of undermining higher-order evidence), then S is rationally permitted to have a belief that p *of that specific type*.

E.g. If S knows that a belief in p *which has been formed on a Tuesday*, is bound to be a true belief, then S is rationally permitted to have a belief formed on a Tuesday that p.

Taking ‘SB_tp’ to mean that S has a belief that p is of type t, we could put this in semi-formal terms:

- **PRINCIPLE 2.0:** $[SK (SB_t p \rightarrow p)] \ \& \ \neg(S \text{ has undermining H.O. evidence}) \rightarrow$
→ Rationally Permitted: SB_tp

As it stands this is not quite right, for S may not know whether she is currently in a position to form beliefs of the specific type in question. E.g. when the type of belief in question is: *formed on a Tuesday*, S may not know which day of the week it is. Even if it *is* in fact Tuesday, if S has no idea what day of the week it is, then it would be rationally amiss of her to go ahead and form a belief that p solely on the basis of her knowledge that a belief formed on a Tuesday would be a true belief. What is required, in addition, is that the subject knows that she is in a position to form the specific type of belief in question – i.e. she must know in addition that it is indeed a Tuesday. This additional requirement was not needed in the more general formulation of PRINCIPLE version 1.1, as when the belief that p in question need not be of any specific type, we assume that there are no specific circumstances, which the subject might be ignorant of, that need to obtain in order for her to be able to form a belief that p (of no specific type). I.e. we assume that a subject always knows that she is in a position to form a belief that p where this belief need not be of any further specific type.

So a better formulation of this special case of PRINCIPLE is:

- **PRINCIPLE version 2.1:** If S knows that a belief of some specific type that p (by S) is bound to be a true belief, *and* S knows that *she is currently able to form a belief that p of that specific type* (and in the absence of undermining higher-order evidence), then S is rationally permitted to have a belief that p of that specific type in these given circumstances.

Or in semi-formal terms:

- **PRINCIPLE 2.1:** $[SK (SB_{tp} \rightarrow p) \ \& \ SK (S \text{ is currently able to } B_{tp}) \ \& \ \neg(S \text{ has undermining H.O. evidence})] \rightarrow \text{Rationally Permitted: } SB_{tp}$

The superiority of version 2.1 to version 2.0 is evident when we consider that the type, t, of belief in question could be: *true belief*. For any proposition whatever, we

can know a priori the tautology that: if I have a true belief that p then p is true. But so according to PRINCIPLE version 2.0, a subject would thereby be rationally permitted to believe any/every proposition that happens to be true, even those for which the subject has no evidence! But clearly we need to rule out that it is rationally permitted for a subject to form a belief that happens to be true by sheer lucky guess. This is ruled out by version 2.1 of PRINCIPLE as a subject who forms a true belief that p by sheer luck does not know that she is in a position to form a belief of the type in question – i.e. a true belief.

Unfortunately, we still need to consider one more potential complication. For there could be cases where a subject knows that a belief that p of type t would be true, and knows that she is currently able to form a belief that p of type t, *but* the subject is only able to form a belief that p of type t in such a way that it would *not* be based on her knowledge that such a belief is guaranteed to be true. E.g. perhaps I know how to form a belief that p of type t but I can only do so by first losing my knowledge that such a belief is guaranteed to be true. Or perhaps the only way I am able to form a belief that p of type t requires me to do so on some other irrational basis, or on no basis whatever. In other words, there needs to be the right sort of *connection* between the subject's knowledge that a belief that p of type t would be true and her knowledge that she is currently able to form a belief that p of type t. The subject must be able to form a belief that p of type t in such a way that it would still be *based on* her knowledge that such a belief is bound to be true.

So our final full version of this special case version of PRINCIPLE, with all the bells and whistles, is as follows:

- **PRINCIPLE version 2.2:** If S knows that a belief of some specific type that p (by S) is bound to be a true belief, and S knows that she is currently able to form a belief that p of that specific type *that would be based on her knowledge that such a belief would be true* (and in the absence of undermining higher-order evidence), then S is rationally permitted to have a belief that p of that specific type in these given circumstances.

Or in semi-formal terms:

- **PRINCIPLE 2.2:** $[SK (SB_t p \rightarrow p) \ \& \ SK (S \text{ is currently able to } B_t p \text{ that is based on } K (SB_t p \rightarrow p)) \ \& \ \neg(S \text{ has undermining H.O. evidence})] \rightarrow$
Rationally Permitted: $SB_t p$

Most of these complications to the PRINCIPLE can largely be ignored when we come to present the argument against UNIQUENESS in the next section. But lest it should be thought that version 2.2 was needless complexity or pedantry for its own sake, consider the case where the type, t , of belief in question is: *irrational* belief. And suppose the proposition, p , in question is:

p : S has at least one irrational belief

Now another interesting and tricky question is whether/when a subject can ever deliberately form a belief that p which they know would be irrational. But lets just grant for the sake of argument that a subject could know that she is in a position to form a belief that p of the irrational type. I.e. she knows that she can currently form an irrational belief that she has an irrational belief. She also knows that such an irrational belief that p is guaranteed to be true. So according to version 2.1 of PRINCIPLE, she would be rationally permitted to form an irrational belief that p ! Version 2.2 blocks this unwanted verdict because it requires not just that the subject knows that she can form a belief of the type in question – i.e. an irrational belief that p – but also that this belief is still based on her knowledge that such a belief would bound to be true. And this simply does not appear to be a possibility. Whilst we may grant that a subject could know how to form a belief that she knows would be irrational, we should not grant that there could be a possible way of forming a belief that is both known to be irrational AND is based on knowledge that the belief in question is guaranteed to be true. For, as I stated at the start of this section, knowledge that the belief in question is guaranteed to be true is a paradigm of an excellent rational basis for belief – there is just no such thing as a belief that is both formed on this basis which is also formed on an irrational basis. Notice that the requirement that there is no undermining higher-order evidence is crucial here. For if the subject *did* have such undermining evidence, then it might be possible to form a belief that p , which is based on knowledge that

such a belief is bound to be true, but which is nevertheless irrationally so-based – e.g. if the subject has misleading H.O. evidence that such knowledge is *not* a good rational basis for believing that p. But once this kind of higher-order undermining is ruled out, it seems there is just no way for a belief that p which is based on knowledge that such a belief must be true to also count as a belief that p which was formed on an irrational basis.

To briefly recap: the core, intuitive idea is simply that when you know that some belief, perhaps of some specific kind, is *guaranteed* to be true, that knowledge is surely an excellent rational basis on which to form such a belief. And then so long as we've ruled out various tricky possibilities – such as higher-order evidence undermining that rational basis, ignorance as to whether you are able to form the kind of belief in question, or some strange situation which prevents you from using the knowledge in question as the basis for your belief – we can safely say that you are rationally permitted to go ahead and form that belief.

3. The Argument

A belief-safespot is a proposition that can be false but which is guaranteed to be true if the subject believes it. A belief-of-type-t-safespot is a proposition that can be false but which is guaranteed to be true if the subject forms a belief of type t in that proposition. The problem for A.B.U. and UNIQUENESS now arises as there can be a proposition that you can know to be a belief-of-type-t-safespot without actually (yet) believing it to be true – indeed, whilst knowing that it is (currently) false. Thus there can be a situation where you are permitted to believe that the safespot is false, where that belief is of some *other* type, but you also know that a belief of type t would be true, and so (by PRINCIPLE) you are also permitted to have belief of type t in the safespot. Thus you are, at the same time and on the same total evidence, permitted either to believe (type t) or to disbelieve (not type t) the one same proposition. Hence UNIQUENESS and ABU are both false and PERMISSIVISM is true.

For example, consider the following proposition:

q: S has at least one non-innate belief.

Notice that this is *not* a belief-safespot for S, as if all of S's beliefs are innate, including her belief in q, then q would be both false and believed by S. But q is a *non-innate-belief-safespot* for S. For q can be false and q can be non-innately believed by S, but it cannot be falsely and non-innately believed by S.

Now, it seems that S could innately know that q is a non-innate-belief-safespot for her – i.e. $SK_i(SB_{ni}q \rightarrow q)$, where K_i = innate knowledge and B_{ni} = non-innate belief. And S can know innately that any new belief she forms will be non-innate – so she knows that she is in a position to form a belief that q of that specific type. Moreover, such a non-innate belief that q could perfectly well be *based on* her innate knowledge that such a non-innate belief that q must be true. (We can just stipulate that there is no undermining higher-order evidence.) And so, by PRINCIPLE 2.2, S is rationally permitted to $B_{ni}q$ – i.e. to have a non-innate belief that q.

But it seems possible that S could also have excellent evidence¹² that q is (currently) false. S might have excellent evidence that all her current beliefs are innate (including her belief that q is a non-innate-belief-safespot for her). Indeed, S might innately *know* that q is false. E.g. perhaps S is an android, furnished with many innate beliefs, who has the ability to form new non-innate beliefs but who has only just been turned on for the first time innately *knowing* both that q is (currently) false and that q is a non-innate-belief-safespot for her. And so it seems that S is also rationally permitted to *innately* disbelieve that q (i.e. to $B_i\neg q$, i.e. to DB_iq). So S is rationally permitted to adopt 2 different doxastic attitudes to p – an innate disbelief or a non-innate belief. Hence, to repeat, A.B.U. and UNIQUENESS are both false, PERMISSIVISM is true.

¹² I do not wish to assume the truth of 'evidentialism' here. If you think that the epistemic rationality of belief can be partially determined by non-evidential factors, then we can just stipulate that these extra factors also obtain in this case.

Nothing in the foregoing hinges on the specifics of the innate vs. non-innate distinction. For it seems clear that further examples of this kind could be manufactured by appeal to other pairs of contrasting belief types. E.g. Consider: beliefs formed after coming to know that you are over 20 years old vs. beliefs formed before coming to know that you are over 20 years old. And consider the proposition: S has formed a belief after coming to know that she is over 20 years old. S could know *since before she turned 20 years old* that: S has formed no beliefs *after coming to know that she is over 20 years old*. But S could also know, again since before the age of 20, that the negation of this proposition – S has formed at least one belief after coming to know that she is 20 years old – *would be true* if first believed by S *after* learning that she is over 20 years old. And so, assuming that S knows that she has now turned 20 and so knows that she is now in a position to form beliefs of the type: believed-only-since-learning-I-am-over-twenty, it seems that S is rationally permitted (at a single time) either to have a disbelief-formed-before-the-age-of-twenty that [I have formed a belief after coming to know I am over 20 years old], or to have a belief-formed-after-learning-I-am-over-twenty that [I have formed a belief after coming to know I am over 20 years old].

In general and semi-formal terms then, the issue is that safespots¹³ allow for the possibility that both of the following conditions obtain:

$SK_{t1} \neg p$

So, Rationally Permitted: $SDB_{t1} p$

$SK_{t1} (SB_{t2} p \rightarrow p) \ \& \ SK_{t1} (S \text{ is in a position to } B_{t2} p \text{ that is based on } K_{t1} (SB_{t2} p \rightarrow p))$
 $\& \ \neg(S \text{ has undermining H.O. evidence})$

So, Rationally Permitted: $SB_{t2} p$ (by PRINCIPLE 2.2)

¹³ I presented the argument in terms of safespot propositions. But it could equally have been presented in terms of a blindspot. The negation of q is a non-innate-belief blindspot for S:

$\neg q$: S has no non-innate beliefs.

S could have excellent (innate) evidence that $\neg q$ is true, and indeed innately know that $\neg q$. So it seems S is rationally permitted to innately believe $\neg q$. But S could also innately know that a non-innate disbelief in $\neg q$ – i.e. a non-innate belief that $\neg \neg q$ – would bound to be true. So, by PRINCIPLE, S is rationally permitted to non-innately believe that $\neg \neg q$ – i.e. disbelieve that $\neg q$.

4. Two Clarifications

The conclusion of the argument is that there can be a situation in which a subject is both rationally permitted to believe a proposition and is rationally permitted to believe its negation. This does not, of course, mean that the subject is rationally permitted to be in the state of simultaneously both believing a proposition and disbelieving it, (nor does it mean that the subject is permitted to believe a contradiction). Permissivists who think that both belief or disbelief in some proposition, p , can be rationally permitted on the same evidence, will clearly wish to deny that this entails it can be rationally permitted to both believe and disbelieve p at the same time. And so they will presumably want to deny that the rational permissibility of belief is closed under conjunction. I.e. they will need to deny that:

Rationally Permitted (SB p) & Rationally Permitted (SB q) \rightarrow Rationally Permitted (SB p & SB q)

But of course the following sort of closure principle for permissibility in general is clearly invalid:

It is permitted to: (do X) & it is permitted to: (do Y) \rightarrow it is permitted to: (do X & do Y)

E.g. S is allowed to marry Jack & S is allowed to marry Jill does not entail that S is allowed to be simulataneously married to both Jack and Jill. So there seems to be no obvious theoretical cost for denying this kind of closure principle in the specific case of rationally permitted doxastic attitudes.

My presentation of the argument in the previous section assumed that a proposition can have different truth-values at different times: the proposition q was innately *known* to be (currently) false, but it was also innately known that q would *turn true* if

non-innately believed. And whilst this sort of view – that propositions have truth-values *at a time* and which can *change* over time – is quite respectable, dating back at least to Aristotle, there will also be those theorists who hold that propositions have their truth-values eternally or timelessly. Fortunately however, the assumption that propositions can change truth-value over time is not essential to the argument.

E.g. consider the following proposition:

q*: S forms at least one non-innate belief (at some time or other).

Let's suppose that this proposition is eternally true. S can know the conditional claim that: if S non-innately believes q* then q* is (eternally) true. And S can know that she is currently able to form a non-innate belief in q*. But S might also have lots of misleading but rationally convincing (innate) evidence which suggests that q* is (eternally) false – S will never in fact ever form a non-innate belief, despite being perfectly *able* to do so. (Here I am assuming that it is possible to know that you are able to do X despite having excellent evidence that you will not in fact ever do X. This may not be totally uncontroversial, but nor is it, I hope, an overly radical or unorthodox assumption to rely on.)

Her evidence then rationally permits the false (innate) belief that q* is false – i.e. it permits her to innately disbelieve that q*. But, by PRINCIPLE, S is also rationally permitted to non-innately believe that q*.

So the safespot proposition in question, q*, does not have to change truth-value over time in order to provide a counter-example to A.B.U. In particular, these counter-examples don't depend on the subject *knowing* that the safespot is (currently) false. It is enough that the subject rationally (though perhaps falsely) disbelieves the proposition in question.

5. Two Possible Objections

Advocates of A.B.U. or of UNIQUENESS then apparently need to deny PRINCIPLE. I will now briefly consider one possible motivation for such a denial.

It might be objected that by tying rational permissibility to a subject's knowledge that a belief *would* be true, PRINCIPLE assumes, or at least is motivated and made plausible by, a form of *epistemic consequentialism* that advocates of UNIQUENESS (or ABU) might want to reject.

- (EC) EPISTEMIC CONSEQUENTIALISM: The ultimate epistemic goals/values, in virtue of which the epistemic rationality of holding any particular belief, on the basis of any particular set of evidence, is to be determined are: (i) acquiring true beliefs, (ii) avoiding false beliefs.

An anti-consequentialist about epistemic rationality will insist, against (EC), that the norms imposed by evidence are fundamentally concerned simply that the subject believes what the evidence indicates *is* (now, actually) true. Such norms are not means to some further end concerning one's set of beliefs; they are not concerned with the *results* of respecting one's evidence in this way. So in particular the rationality of forming a belief in accord with one's evidence is not fundamentally to be explained in terms of its promoting the goals of gaining accurate beliefs and avoiding inaccurate ones; rather, the epistemic rationality of forming some belief *just is* determined simply by whether (and the extent to which) the evidence indicates that the proposition in question *is true*.

- (AC) EPISTEMIC ANTI-CONSEQUENTIALISM: The epistemic rationality of forming any particular belief, in response to any particular set of evidence, is determined solely by which propositions the evidence indicates (more or less strongly) to be (currently/actually) true or false.

[And hence rationality is not fundamentally determined by reference to the promotion of the doxastic goals/values (i) or (ii), in the statement of (EC) above.]

This is not the place to mount a full discussion of epistemic consequentialism. But I will note that blindspots seem to yield a particularly unintuitive result for the variety

of anti-consequentialism just stated. For your evidence could very strongly indicate that some blindspot proposition is (now, actually) true. But this same evidence could also clearly indicate that a belief in the proposition is bound to be false. It sounds very strange then to say that your evidence here rationally permits belief in the blindspot, when it manifestly indicates to you that the proposition is bound to be false if you believe it.

In other words, if one rejects PRINCIPLE on general anti-consequentialist grounds, it would seem that one should also be committed to rejecting the following:

- CONVERSE PRINCIPLE 1.0¹⁴: $[SK (SBp \rightarrow \neg p)] \rightarrow \neg \text{Rationally Permitted: SBp}$

If anything, this CONVERSE PRINCIPLE seems to be even more plausible than the original PRINCIPLE. When you *know* that your believing something would bound to be a false belief, then it is not rationally permitted for you to have that belief. But if, in adherence to (AC), one insists that one's evidence rationalises one's beliefs *solely* in virtue of what it indicates *is actually* true/false, then so long as your evidence indicates strongly enough that a blindspot is currently true, you should be permitted to believe it even though you know this belief would be false.

We have seen how PRINCIPLE is in conflict with (AC) – safespots can provide examples where the evidence can indicate both that a proposition is false and that it would be true if believed. In these cases, PRINCIPLE insists, against (AC), that what is rationally permitted to believe is *not* determined solely by what the evidence indicates to be true/false. (And likewise, blindspots can provide examples where the evidence can indicate both that a proposition is true and that it would be false if believed. In these cases CONVERSE PRINCIPLE insists, against (AC), that what is

¹⁴ For simplicity, I have here stated CONVERSE PRINCIPLE in a formulation that does not include a clause mentioning higher-order evidence. Perhaps some philosophers will want to allow that there can be cases where despite *knowing* that a belief that p is bound to be false, you could nevertheless be rationally permitted to form such a belief because you have misleading higher-order evidence that it *is* rationally permitted (or even obliged). In which case, again, we would need a slightly more complicated formulation such as the following:

CONVERSE PRINCIPLE 1.1: $[SK (SBp \rightarrow \neg p) \ \& \ \neg(S \text{ has countervailing H.O. evidence})] \rightarrow \neg \text{Rationally Permitted: SBp}$

rational to believe is not determined solely by what the evidence indicates to be true/false.)

But notice, accepting PRINCIPLE (or CONVERSE PRINCIPLE) does not obviously require going so far as endorsing (EC). I.e. accepting that there are *some* cases in which the consequences of forming a belief are *relevant* to assessing epistemic rationality, does not obviously require accepting that in *every* case such consequences are the *sole* or *ultimate* determinants of epistemic rationality. And this may be just as well for the plausibility of PRINCIPLE and CONVERSE PRINCIPLE, because consequentialism may well have its own problems – see e.g. Berker (2013) for arguments against (EC)¹⁵. (So my suggestion that (AC) faces a problem dealing with blindspots and safespots should certainly not be taken as an endorsement of (EC).)

Now, another sort of reaction someone might have to the argument of this paper is that appealing to such tricky, self-referential propositions as blindspots or safespots is a kind of sophistry: the debate between advocates of UNIQUENESS and PERMISSIVISM arose from considering what to do in real-life situations of peer disagreement. Insofar as they take themselves to provide good answers to these sorts of pressing, real-life epistemic concerns, advocates of UNIQUENESS or of A.B.U. should not be worried by the existence of *recherché* counter-examples, based on self-referential propositions that have no real-life use outside of philosophy departments.

In reply, it is tempting to give the following somewhat flippant response: since when was being *recherché* or self-referential a bar to being a good counter-example in analytic philosophy? But as a less flippant response, I would make two points. Firstly, it is not obviously true that blindspot propositions never have any real-life relevance. In his classic 1988 treatment of the topic, Sorensen provides the example of the ‘efficient market hypothesis’¹⁶ in economics, which he argues is a form of ‘complex

¹⁵ Though see Ahlstrom-Vij & Dunn (2014) for a defence of epistemic consequentialism against Berker’s argument.

¹⁶ This hypothesis is, very roughly, that the market is so efficient at gathering and distributing information that any information relevant to the price of stock is near-instantly reflected in its price, so that it is impossible to turn a profit on the basis such information.

blindspot’ – it is a hypothesis that is well-supported by evidence and accepted by many economists, but in order for it to be true it must be disbelieved by at least a significant proportion of economic agents (see Sorensen, 1988, 109-114). He also shows how ‘blindspotting’ – i.e. providing evidence for a proposition that is a blindspot for your opponent – can be an important strategy in ‘prisoner’s dilemma’ type games (see Sorensen, 1988, chs 9-10).

Secondly, and more importantly, blindspots and safespots serve to highlight an important distinction between two possible understandings of what it is to aim for accurate beliefs:

- Aiming to form a belief that accurately represents how the world *actually/currently* is (i.e. before the belief is actual)
- Trying to form beliefs that *will* be an accurate representation of the world *once they have been formed* (i.e. once the belief is actual).

Blindspots and safespots show how these two aims can come apart, as they provide the clearest illustration of how the truth-value of a proposition can sometimes depend on the attitude we adopt towards it. The difference between these two aims has not always been sufficiently carefully noted. I will limit myself to just two examples.

Richard Foley makes the (consequentialist) claim that epistemic rationality is fundamentally based on the goal of “now believing those propositions that are true and now not believing those propositions that are false” (Foley 1993, 19). This seems to miss that there are the potentially distinct goals of believing propositions that will be true, and not believing propositions that will be false, *once the beliefs have been formed*. If our goal was simply to now believe propositions that are currently true then, it should be rationally permitted to believe a currently true, non-believed blindspot, even though we know full well that the proposition would be false if we believed it. [Likewise, with a currently false, non-believed safespot – according to Foley’s fundamental goal we should simply be concerned to now not believe this false

proposition, despite knowing that the proposition would be true if we went ahead and believed it.]

And in a more recent paper on epistemic reasons, Sharon Street writes:

‘We intuitively make a distinction between (1) reasons to believe P that **tell in favour of the truth of P**, and (2) reasons to believe P that in no way tell in favour of the truth of P, but instead point to something to be gained by believing P, or cite some other practical consideration in favour of holding that belief. As is customary, let’s call the former sort of reason epistemic reasons for belief, and the latter sort of reason pragmatic reasons for belief.’ (Street, 2009, 216-7, emphasis added)

But at the end of the next paragraph Street writes:

‘...in this paper I will concentrate exclusively on the case of epistemic reasons - that is, considerations that count for or against a belief in virtue of the way they **bear on that belief’s truth or falsity**.’ (ibid, 217, emphasis added)

Blindspots and Safespots show that these 2 definitions of an epistemic reason are not equivalent – evidence can indicate that a blindspot proposition is true without telling in favour of the truth of one’s belief in that proposition. Conversely evidence can indicate that one’s belief in a safespot would be true without also telling in favour of the safespot proposition actually being true. It is also worth noting that there appear to be cases of an epistemic reason telling against a belief that p, where the reason *neither* indicates anything about the truth-value of the proposition, *nor* anything about whether the belief that p would be true or would be false. Suppose that you have no evidence whatsoever for or against p. Now the world’s leading expert on the general topic of p, who has carefully considered the question better than anyone else and whom you have every reason to trust, gives you the following testimony: you ought to suspend judgement whether p¹⁷. This would seem to be an epistemic reason that counts against believing that p (or disbelieving that p). It certainly doesn’t look like a *pragmatic* reason not to believe that p. But the testimony that you ought to suspend

¹⁷ The example is drawn from Turri (2012), which focuses on the much trickier case where the subject’s only evidence is expert testimony is that she ought *not* to suspend judgement whether p.

judgement whether p gives you no information whatsoever as to whether p is true or false, nor does it give you any information as to whether a belief that p would be true or would be false. All it tells you is that we currently don't have any good basis to judge either way – i.e. it tells against the epistemic rationality of believing that p (or disbelieving that p) without in any way telling for or against either the truth/falsity of that belief, or the truth/falsity of that proposition.

6. Degrees of Belief

Returning to Ramsey's image of belief as an internal map, blindspot and safespot propositions show that if one wants an accurate map with which to navigate the world, it is not always enough to update and revise one's map by taking account of how the world actually/currently is. One must sometimes also take into account how the world (including one's map) *would/will be* after the revisions to the map have been made – for the very act of updating one's map can influence which propositions are true/false. Blindspots and safespots are extreme cases of this, where the act of believing p *entails* that the proposition is false or entails that it is true. But there can also be cases where forming a belief that p merely makes p *more likely* to be true or more likely to be false. In 'The Will to Believe' William James (1897) provides some examples: my believing that she will fall in love with me may (sometimes) make it slightly more likely that she will indeed fall in love with me¹⁸ – though such a belief does not, alas, entail that she will love me! Likewise, our believing that we can defeat the train-robbing highwaymen can make it more likely that we will indeed defeat the train-robbing highwaymen¹⁹. James does not consider the converse sort of case, but perhaps my forming a belief that I am about to go to sleep can sometimes make it slightly *less* likely that I am about to go to sleep.

Consider then a further variation on PRINCIPLE framed in terms of degrees of belief rather than full belief:

¹⁸ 'How many women's hearts are vanquished by the mere sanguine insistence of some man that they *must* love him!' (James, 1897, 25)

¹⁹ See section IX of James, 1897.

- **PRINCIPLE 3.0:** if S knows that [*if* she believes that p to a degree = n, then p *would* have a probability = n], then (modulo undermining evidence) S is rationally permitted to believe that p to a degree = n.

[Notice that PRINCIPLE 1.0 is effectively just the special case of PRINCIPLE 3.0 where $n = 1$.]

Now we could have a situation, analogous to a safespot, in which S's evidence indicates that p *actually* now has some probability, n, but the same set of evidence could also indicate that p *would have* some other probability, $n \pm x$, *if* S increased/decreased her confidence in p to $n \pm x$. And thus again we could have a situation where the evidence rationally permits more than one attitude. E.g. it could be that one has a degree of belief that $p = 0.7$, a level of confidence which is rationally appropriate given one's evidence whether p, but that same evidence might also indicate that *were* you to increase your confidence that p to, say, 0.8, the likelihood that p would be = 0.8.

We could also frame a version of CONVERSE PRINCIPLE for degrees of belief:

- **CONVERSE PRINCIPLE 3.0:** if S knows that [*if* she believes that p to a degree = n, then p *would not* have a probability = n], then (modulo undermining evidence) S is *not* rationally permitted to believe that p to a degree = n.

And we could also have a situation, analogous to a true blindspot, where one's evidence indicates that a proposition, p, has a probability of n, but the same set of evidence also indicates that p *would have* some *other* probability, $n \pm x$, *if* S sets her confidence in p to n. E.g. the evidence indicates that p has a probability of 0.7, but it also indicates that *were* you to form a credence of 0.7 in p, then p *would have* a probability significantly different than 0.7.

I think the plausibility of Principle 3.0 and Converse Principle 3.0 is enhanced when we note that they seem to be importantly similar both to Lewis's (1980) famous Principal Principle and to Van Frassen's (1984) Reflection Principle.

- PRINCIPAL PRINCIPLE (simplified!): $Cr(A | Ch(A) = n) =_{\text{OUGHT}} n$
- REFLECTION (simplified!): $Cr_{t_1}(A | Cr_{t_2}(A) = n) =_{\text{OUGHT}} n$

[Cr stands for the subject's subjective credences, Ch is the objective chance. t_2 is a time later than t_1 , so these subscripts are no longer about types of belief.]

The statement of Principal principle, above, is simplified from Lewis's original (and from his later New Principal principle) as Lewis restricts the principle to propositions and evidence that are about events/states up to a certain time t . And the statement of Reflection is also over-simple as it is now widely accepted that you could rationally flout Reflection if you have good reason to think that your later self at time t_2 will be epistemically worse off – drugged, or forgetful or whatever. So some kind of restriction is needed to cases where you reasonably expect to be in at least as good a state epistemically at t_2 as at t_1 .

Putting Principle 3.0 and Converse Principle 3.0 into similar-style formulae concerning conditional credences we get:

PRINCIPLE 3.1:

$$Cr(A | Ch[A | (Cr(A) = n) = n] = n) =_{\text{PERMITTED}} n$$

CONVERSE PRINCIPLE 3.1:

$$Cr(A | Ch[A | (Cr(A) = n) \neq n] \neq n) \neq_{\text{PERMITTED}} n$$

Now one very important difference between my principles and those of Lewis and Van Frassen is that Principal Principle and Reflection are, I take it, about what an

agent's credences *ought* rationally to be (i.e. they are norms that impose an *obligation*), whereas my PRINCIPLE in all of its versions only concerns what an agent's credences *may* rationally be (i.e. they are norms that provide rational *permission*). Hence the subscripts 'OUGHT' and 'PERMITTED' have been added to the respective formulae above to make this difference explicit.

Having noted this difference, I think that there is a real and important similarity with Principal Principle, which is a requirement that one's subjective probabilities should match up with (what you take to be) the objective chances. PRINCIPLE 3.1 & CONVERSE PRINCIPLE 3.1 can be thought of as covering the special case where we are dealing with the objective chance of some proposition *conditional* on your having some specific degree of belief in that proposition.

Likewise I think there is a genuine similarity with Reflection, which is a principle that effectively requires you to accept your own future unconditional credences – again, assuming that you take this future self to be in as good an epistemic position or better. PRINCIPLE 3.1 spells out one instance of when a future/counterfactual self will be in a good epistemic position. You *may* rationally accept your future or counterfactual self's unconditional credences if you know that this future/possible credence *would be* in line with the objective chances. Conversely: you are not rationally permitted to accept your future/counterfactual credences if you know that these future/possible credences would *not* be in accord with the objective chances.

7. Conclusion

To recap briefly: the general moral here is that a set of evidence may not only indicate what is or is not the case, but also what would or would not be the case if you have certain beliefs – i.e. if your map is drawn in a certain way. Epistemic rationality can depend on both kinds of indication²⁰. And where these two kinds of indications come

²⁰ Given that a map which has been drawn so as to accurately reflect the state of the world *prior* to any updating will generally be *less useful* for subsequently navigating the world than a map which accurately reflects how the world is *after* the update to the map is complete, one might be tempted to dismiss aiming for the latter as a matter of (mere) *practical* rationality, as opposed to epistemic rationality. But this would be wrong. It is easy to imagine scenarios in which forming *either* kind of

apart there may be two equally rational options for how to draw one's 'map of the world'. UNIQUENESS and A.B.U. fail to accommodate this possibility.

Advocates of UNIQUENESS have, of course, given their own arguments in favour of the thesis and against PERMISSIVISM²¹; nothing I have said is supposed to indicate what flaws there are, if any, in those arguments. And so if one were sufficiently strongly convinced on independent grounds that UNIQUENESS must be correct, one might treat the foregoing argument of this paper as a reason to reject PRINCIPLE. But in any case, whatever the ultimate theoretical costs or benefits of denying PRINCIPLE, I hope that the need to deny it is at least an interestingly non-obvious and, *prima facie*, implausible consequence of UNIQUENESS and of A.B.U.

References

- Ahlstrom-Vij, K. & Dunn, J (2014) 'A Defence of Epistemic Consequentialism', *The Philosophical Quarterly*, 64 (257): 541-551
- Berker, S. (2013) 'Epistemic Teleology and the Separateness of Propositions.' *Philosophical Review* 122 (3): 337-393.
- Bykvist, K. and Hattiangadi, A. (2007) 'Does thought imply ought?' *Analysis* 67: 277-85.
- Christensen, D. (2007) "Epistemology of Disagreement: The Good News", *Philosophical Review* 116 (2): 187-217.
- Douven, I. (2009) "Uniqueness Revisited." *American Philosophical Quarterly*, 46 (4) 347-362.
- Feldman, R. (2007). "Reasonable Religious Disagreement" in L. Antony ed. *Philosophers Without God*. Oxford: Oxford University Press. 197-214.
- Foley, R. (1993) *Working without a Net: A Study of Egocentric Epistemology*. Oxford: Oxford University Press.
- Gibbons, J. (2013) *The Norm of Belief*, Oxford: Oxford University Press.
- James, W. (1897) *The Will to Believe and Other Essays in Popular Philosophy*, Cambridge, MA: Harvard University Press.
- Kelly, T. (2014) "How to Be an Epistemic Permissivist" in M. Steup & J. Turri (eds.) *Contemporary Debates in Epistemology*, 2nd edition. Malden, MA: Blackwell Publishing.
- Lasonen-Aarnio, M. (2014) "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88: 314 – 45

'map' would not be in your best practical interests. E.g. your evidence strongly indicates that p and you do in fact innately believe that p, but the evidence also indicates that were you to non-innately disbelieve that p then not-p. However, an evil demon is going to appear tomorrow and kill you and your family if you either (innately) believe or (non-innately) disbelieve that p. Neither the (innate) belief nor the (non-innate) disbelief is practically rational here. And yet, whatever your practical concerns, there is a clear, non-pragmatic sense in which both beliefs are *rationalised simply by the evidence*, for both beliefs are ways of aiming at accurately representing the world that are appropriately based on that evidence.

²¹ See footnote 3 for references.

- Lewis, D. (1980) “A Subjectivist's Guide to Objective Chance,” in R. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, Volume II, Berkeley: University of California Press.
- Ramsey, F. P. (1931) ‘General Propositions and Causality’, in *The Foundations of Mathematics and other Logical Essays*, (ed.) R. B. Braithwaite, London: Kegan Paul, 237-55.
- Schoenfield, M. (forthcoming) “Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief”, *Noûs*.
- Sorensen, R. (1988) *Blindspots*. Oxford: Clarendon Press.
- Sosa, E. (2010) “The Epistemology of Disagreement” in *Armchair Philosophy*, Princeton, NJ: Princeton University Press.
- Street, S. (2009) “Evolution and the Normativity of Epistemic Reasons”, *Canadian Journal of Philosophy*, 39: sup. Vol. 1, 213-48.
- Sturgeon, S. (forthcoming) ‘Epistemic Attitudes’.
- Titlebaum, M. (2014) “Rationality’s Fixed Point”, *Oxford Studies in Epistemology* 5: 253-94.
- Turri, J. (2012). “A Puzzle about withholding”, *Philosophical Quarterly*, Vol. 62: 247, 355-364.
- van Frassen, B. (1984) “Belief and the will”, *Journal of Philosophy*, 81: 235–256.
- White, R. (2005) “Epistemic Permissiveness.” *Philosophical Perspectives* 19 (1): 445–459.