

## Epistemic Entitlement at the Limits of Doubt

Gary Ebbs  
Draft 9/24/12

Some statements we accept, including the statement that *it is not the case that snow is both white and not white*, are such that we cannot now coherently suppose them to be false, where the “cannot” comes to this: if we try to specify a way in which they may be false, we find we are unable to do so. *In what sense is it epistemically reasonable for us to accept such statements?*

I shall motivate and clarify this question, raised from a first-person, deliberative point of view, sketch three constraints on a satisfactory answer to it, and show that two tempting and representative attempts to construct substantive theoretical answers to the question fail to satisfy the constraints. To satisfy the constraints, I shall argue, we need a different approach. We need to *elucidate* the special role in our inquiries of beliefs that we cannot coherently suppose to be false. I shall end by sketching an elucidation of this epistemological role that satisfies my three constraints and thereby answers my central question.

### 1. A methodological approach to epistemology

My question is rooted in a *methodological* approach to epistemology that I have come to see and appreciate only after years of trying to understand Rudolf Carnap’s, W. V. Quine’s, and Hilary Putnam’s revolutionary alternatives to traditional epistemology. This methodological approach assumes that truth is a goal of inquiry, and that to pursue it we can do no better than to rely on already established beliefs and inferences, and apply our best methods for reevaluating particular beliefs and inferences and arriving at new ones. We may find fault with our grounds for accepting any of our beliefs, inferences, or methods, and revise them accordingly, but only if such findings or revisions are rooted in our latest, best sense of which ones to accept. When the epistemologist’s task is disciplined by this methodological approach, as Quine says, it “differs from others’ ... in no such drastic way as those suppose who imagine for the [epistemologist] a vantage point outside the conceptual scheme that he takes in charge. There is no such cosmic exile.” (Quine, *Word and Object*, p. 275)

Assuming this methodological approach to epistemology, it is natural to affirm that

(I) We have no grip on what it is for a belief to be *epistemically reasonable* apart from our already established beliefs and inferences and our best current methods for reevaluating particular beliefs and inferences and arriving at new ones,

and

(II) There are no *epistemological standards* higher or firmer than the ones we express in our actual endorsements of particular beliefs, inferences, and methods for arriving at new beliefs.

I intend these claims to be partly stipulative of the meanings of “epistemically reasonable” and “epistemological standards”, not as faithful to all aspects of the meanings that these terms are taken to have in traditional or contemporary epistemology. On the methodological approach that I recommend, we own and use our already established beliefs and methods to provide a general schema for applying these abstract and troublesome terms to particular beliefs and inferences that we accept in the course of our inquiries.

The question I shall formulate and try to answer arises when one asks, from one’s first-person, deliberative point of view, while accepting claims (I) and (II), in what sense it is epistemically reasonable for one to accept a statement that one cannot now coherently suppose to be false.

## 2. Transparency of inquiry

A preliminary step toward answering this question is to ask a more general one: In what sense is it epistemically reasonable for one to accept *any* given statement? For a wide range of statements, there is a ready and appealing answer. Suppose, for instance, that prior to looking carefully at a given page, we have no reason to believe there is an occurrence of the word ‘inquiry’ on it and no reason to believe there isn’t. In such circumstances, to determine whether or not to believe that there is an occurrence of ‘inquiry’ on the page, we rely on a vast background of beliefs about such things as what pages are, what occurrences are, what counts as an occurrence of ‘inquiry’ on a page, and so on, and we focus on the question whether or not there is an occurrence of ‘inquiry’ on the page. To answer this question, of course, we examine the page. If we find an occurrence of ‘inquiry’ on it, we will for that reason come to believe that there is one. A belief that we come to form in this way is clearly epistemically reasonable in the sense constrained by (I).

This example suggests a general strategy for clarifying from one’s own first-person, deliberative point of view the sense in which it is epistemically reasonable for one to accept a given statement, or hold the corresponding belief. The leading idea behind the strategy is that for each of us, the question whether to believe that S comes down to the question whether or not S. One might say (following Roy Edgely, Gareth Evans, and Richard Moran, among others) that when we deliberate about whether to believe some statement S that we understand, the question of whether to believe S is *transparent* for us, in the sense that we pass directly *through* it to the question whether or not S.

This point of view on belief is often emphasized in discussions of how a person can come to know what she believes without relying on any *evidence* about what she believes. In contrast, my focus here is on the suggestion that by appealing to the transparency of belief, we can each clarify, from our own first-person, deliberative point of view, what it is for us to believe in a way that is epistemically reasonable in a sense constrained by (I).

It is useful to compare this sort of transparency with our understanding of ‘true’, as defined disquotationally for sentences we can use. Thus defined, the predicate ‘true’ is as clear to us as our uses of the sentences to which we apply it. For instance, ‘there is an occurrence of ‘inquiry’ on page 2’ is true if and only if there is an occurrence of ‘inquiry’ on page 2. Here we use the sentence on the right hand side, thereby directing our attention to page 2, not to the sentence mentioned on the left hand side, or the belief we can express by uttering that sentence.

These familiar points bear on our understanding of beliefs in contexts where we express them by using sentences. In such cases, we think of ourselves as focused on the world, not on our beliefs (except when our sentences describe our beliefs). Direct illustrations of the transparency of belief are of course presented in language. We note, for instance, that the question whether to believe that there is an occurrence of ‘inquiry’ on page 2 is transparent for us to the question whether or not there is an occurrence of ‘inquiry’ on page 2. When we define ‘true’ disquotationally, or formulate the questions to which our beliefs are transparent, we use our words in a way that appears to “reach out and touch reality”. (I allude here to 2.1511 and 2.1515 of Wittgenstein’s *Tractatus*.) But we have no understanding of such imagery apart from particular uses of sentences that we find clear. Let us dispense with the imagery, then, and think of transparency as in part *displayed* in the very assertions and other uses of language that express our best current theorizing.

What I am calling the transparency of belief takes these observations about disquotational truth and language use for granted, and adds the observation that from our first-person, deliberative point of view, for any given statement S, we think of the question whether to believe S as settled by our answer to the question whether or not S, and we take this answer to be not a mere projection from our antecedent belief that S, but, in a sense to be clarified, an answer that we arrive at by attending to the question whether or not S.

For different statements S, what it is to “arrive at one’s belief that S by attending to the question whether or not S” may differ. Hence the idea that the question whether to believe S is transparent for us is *schematic*. To apply it to a given belief that S, we need to specify the sense in which we arrive at our belief that S by attending to the question whether or not S.

### 3. Bipolarity and unipolarity

The ‘inquiry’ example suggests ways of doing so for some statements S that we take to be bipolar, in the following sense

**Bipolarity:** Person P takes statement S to be bipolar (at t) if and only if P can make sense of S’s being true (at t) and P can make sense of S’s being not true (at t).

If we take a statement S to be bipolar in this sense, and we can specify the sense in which we arrive at our belief that S by attending to the question whether or not S, then we can specify a particular sense in which our answer to the question whether or not S is

transparent for us to the question whether or not S. The statement that there is an occurrence of ‘inquiry’ on page 2 is bipolar, and I can tell by looking whether or not there is an occurrence of ‘inquiry’ on page 2. In other similar perceptual cases, I might specify the sense in which a belief that S is transparent for me by noting that I can tell by listening, touching, smelling, or tasting, whether or not S. In such cases, we can specify the sense in which I arrive at my beliefs by attending to the question whether or not S—I arrive at my beliefs in these cases by looking, listening, touching, smelling, or tasting, whether or not S.

There are of course many statements that we take to be bipolar that we cannot evaluate by perceiving alone. In practice, for any such statement S, we rely on non-perceptual ways of answering the question whether or not S, such as inferring that S by induction, from initial conditions and a well-confirmed, high-level explanatory law, or, more generally, by reasoning to the best explanation. I shall assume that our resources are varied and comprehensive enough, and the requirements for reasonable responsiveness to considerations are weak enough, that, typically, if we take S to be bipolar, we can specify considerations that we take to be relevant to answering the question whether or not S, and thereby filling in the specific sense in which we arrive at our belief that S by attending to the question whether or not S.

Understood in this way, the strategy suggested by the ‘inquiry’ example covers all cases in which a person has (or comes to have) either a reason for believing S, or a reason for believing not S, where these reasons are available prior to and independently of the belief. To categorize these cases, it helps to use a notion of *reason* that satisfies the following constraint:

- (C1) A person has a *reason* for believing that S only if she can think or say *why* she believes that S without presupposing that S.

I come to have a reason in this sense for believing that there is an occurrence of ‘inquiry’ on a given page if, (a) before I look at the page, it is epistemically possible for me that there is such an occurrence and it is epistemically possible for me that there isn’t, and (b) I then come to believe, by looking, that *this* (where I indicate some particular word) is an occurrence of ‘inquiry’ on the page, and I existentially generalize on the content of this perceptual belief. When conditions similar to (a) and (b) hold for a given person P and statement S, we can specify a sense P’s question whether to believe that S is transparent for her to the question whether or not S by citing whatever *reasons* P has (or comes to have) for believing that S, in the sense of “reason” constrained by (C1). (What exactly reasons are is a difficult question that I will not discuss here.)

There are statements that we do not take to be bipolar, however, including statements we cannot coherently suppose to be false, where, as noted above, the “cannot” comes to this: when we try to specify a way in which a statement of this sort may be false, we find we are unable to do so. We take such statements to be unipolar, in the following sense

**Unipolarity:** P takes S to be unipolar (at t) if and only if P can make sense of S's being true (at t) or P can make sense of S's being not true (at t), but not both.

For example, I take the statement that *it is not the case that snow is white and snow is not white* to be unipolar, since I can make sense of its being true—indeed, I accept it—but I cannot make sense of its being not true (i.e. false): when I try to specify a way in which it may be false, I find I am unable to do so.

If we take a statement S to be unipolar in this sense, we cannot give reasons for believing that S, in the sense of “reason” constrained by (C1), since we cannot think or say why we believe that S without in some way relying on S.

My question, again, is this:

In what sense is it epistemically reasonable to accept this and other such statements, and hold the corresponding beliefs?

Given my methodological assumptions, I shall now argue, this amounts to the question:

In what sense are such beliefs transparent for us?

#### 4. Against methodological analyticity

One will be inclined to resist this reformulation of my question if one believes that some sentences are analytic, or true in virtue of meaning, in a sense that is immediately known to anyone who can use the sentences in discourse. To believe this is to be committed to what I call *methodological analyticity*, according to which we can know of certain specified sentences that we could never come to revise or reject them without changing their meanings.

Suppose S is a statement we take to be unipolar, and  $\sigma$  is the sentence we use to express S. If we take sentence  $\sigma$  to be analytic in the methodological sense I just sketched, then we will be tempted to think that we can infer from our grasp of the meaning of  $\sigma$  that statement S is true, and hence that it is epistemically reasonable for us to accept S. We would therefore be tempted to conclude that S does not really “say anything about how the world is”, and hence that it is inappropriate to seek an articulation of a sense in which the question whether to believe that S is *transparent* for us. It will seem to us that for any statement S, our belief that S is transparent, and hence that in some sense we arrive at our belief that S by attending to the question whether or not S, only if we take S to be bipolar.

If we accept the methodological approach to epistemology that I sketched above, however, we cannot appeal to methodological analyticity to explain why it is epistemically reasonable for us to accept a statement that we take to be unipolar. The central problem is that what we *actually* judge at a given time may conflict with and trump our earlier claims about what we “must” judge in any new situation if we are to count as knowing the meaning, or as having known the meaning, of sentences we use to

express a given statement. I may strongly believe that I could not reject the sentence “It is not the case that snow is both white and not white” without thereby changing its meaning, but I cannot legitimately cite this as a reason for denying that I might later come to be able to specify a way in which the statement that I now express by using that sentence may actually be false. For I am in principle open to discovering that there is some way in which the statement that it is not the case that snow is both white and not white, and the sentence “It is not the case that snow is both white and not white” that I use to express it, may actually be false. And if I do discover such a way, I will thereby also discover that I could in principle come to reject the sentence without changing its meaning, contrary to the assumption that it is methodologically analytic.

This is not to say that I can now specify any way in which the sentence may actually be false, that I have any reason to believe that there is such a way, or even that I can coherently imagine a possible future in which I discover a way in which it may actually be false. It is simply to express an in-principle openness to new discoveries that goes with the methodological approach to epistemology that I sketched above. Given this openness, I cannot explain why it is epistemically reasonable for me to believe that *it is not the case that snow is both white and not white*, for instance, in terms of methodological analyticity.

This argument against methodological analyticity shows that we have no good reason to think that if we take a statement *S* to be unipolar, then our belief that *S* cannot be transparent for us. More positively, I take the argument to show that *even for a statement S that we take to be unipolar, the proper point of view from which to clarify the sense in which a belief that S is epistemically reasonable is the one we occupy when we regard the question whether to believe that S as transparent for us to the question whether or not S.*

I conclude that my question, namely,

In what sense is it epistemically reasonable to accept statements that we take to be unipolar?

amounts to the question

In what sense are such beliefs are transparent for us?

## 5. Entitlements

A first step toward answering this question is to identify a type of epistemological standing that is distinct from having a reason for believing a statement, in the sense of “reason” constrained by (C1), but sufficient for believing it in an epistemically reasonable way, in a sense of “epistemically reasonable” constrained by (I).

Consider, to take a slightly more complex example than the one I have used so far, my belief that no sentence of the form ‘*S* and not *S*’ is true. Both my expression of this belief and the belief itself presuppose a background of beliefs about how to define truth, about

the semantics of negation and conjunction, about what follows from what, and so on. These are not beliefs that I can now suspend or reject without losing my grip on what I say when I use my sentence, “No sentence of the form ‘S and not S’ is true.” Although I have tried to make sense of challenges to my belief that no sentence of the form ‘S and not S’ is true, including challenges from Graham Priest and other dialethiasts, I find I am (so far) unable to describe coherently any situation in which the belief is false. Moreover, all my attempts to argue for the conclusion that no sentence of the form ‘S and not S’ is true rely at some point or other on a premise or inference rule that is either stronger than the conclusion, or equivalent to it. Consider, for instance, the following argument:

Start with two semantic premises: (1) ‘Not S’ is true if and only if S is not true, and (2) ‘S and S’ is true if and only if S is true and S’ is true. Now suppose (toward contradiction) that for some S, ‘S and not S’ is true. Then, by (2), S is true and ‘Not S’ is true. Hence, by (1), S is true and S is not true. *This is a contradiction, hence not true*, and so it is not the case that for some S, ‘S and not S’ is true.

This argument begs the question by assuming that there can be no true contradictions. (Priest 1998, p. 418) More generally, my failure to find a non-circular argument for the conclusion that no sentence of the form ‘S and not S’ is true convinces me that I have no *reason* for believing that this, in the sense of “reason” constrained by (C1). The statement that no sentence of the form ‘S and not S’ is true is part of my best current theory, however, and I accordingly take my belief that no sentence of the form ‘S and not S’ is true to be epistemically reasonable, in accord with (I). By claim (II), moreover, there is no higher or firmer perspective from which to judge whether the belief is epistemically reasonable.

We need a word to describe the epistemological standing of such beliefs. I propose that we use ‘entitlement’, subject to the following constraint:

(C2) A person has an *entitlement* (or is entitled) to believe that S if and only if she has no reason (in the sense of ‘reason’ constrained by (C1)) for believing that S—she cannot say *why* she believes that S without relying on S—but it is *epistemically reasonable* for her to believe that S.

I assume that in applying (C2) in accord with assumption (I) we will be guided by our own best judgments about whether a person’s acceptance of the statement is epistemically reasonable. In particular, if I judge that my belief that no sentence of the form ‘S and not S’ is true is epistemically reasonable, and I accept (C2), I commit myself to judging that I am entitled to believe that no sentence of the form ‘S and not S’ is true. (Of course, I need not accept (C2) to have that entitlement.)

It is not in general true, however, that if a person cannot specify a way in which a statement S may be false, then she is epistemically entitled to believe that S. For she might not understand S at all, and hence might not be in a position even to *consider* the statement, let alone to accept, or reject it. How much understanding is required is a

subtle, context-sensitive matter that central to the plausibility of claim that a person has an entitlement to accept a sentence in a sense that is compatible with (C2). Unfortunately, I do not have the space to discuss this matter further here, and so my explanation of entitlement must to that extent remain schematic. The point I want to insist on here is just that some level of understanding is required for entitlement. Due diligence in the effort to specify a way in which a given statement S may actually be false is also required for entitlement to believe that S. I therefore propose the following constraint:

- (C3) If a person understands a statement S well enough to raise the question of whether or not S is to be believed, she tries to specify a way in which S may actually be false, exercises due diligence in this effort, but finds she is unable to specify a way in which S may actually be false, then it is *epistemically reasonable* for her to believe that S.<sup>1</sup>

Again, I assume that in applying (C3) in accord with assumption (I) we will be guided by our own best judgments about whether a person's acceptance of the statement is epistemically reasonable. In a particular case, one might infer, from one's judgment that it is not epistemically reasonable for a given person to accept S, either that she does not understand S well enough to raise the question of whether or not she believes it, or that with due diligence she could specify a way in which S may actually be false. Applying (C3), in combination with (C1), (C2) and assumption (I), is a holistic task, one of weighing particular judgments against others, and revising some of them, if necessary, in order to satisfy the constraints.

I shall henceforth assume that, in a way that is consistent with (C2) and (C3), we are entitled to believe some statements that we take to be unipolar, including the statement that no sentence of the form 'S and not S' is true. To extend the above strategy to such cases we need to specify a sense in which beliefs to which a person is so entitled are transparent for her, and, in light of this specified sense of transparency, to find a related sense in which it is epistemically reasonable for her to accept them.

## 6. The Central Problem

The Central Problem facing any such effort is that the antecedent of (C3) — “a person understands a statement S well enough to raise the question of whether or not S is to be believed, she tries to specify a way in which S may actually be false, exercises due diligence in this effort, but finds she is unable to specify a way in which S may actually be false”— describes conditions that seem irrelevant to the question whether a person's acceptance of S, when seen from her engaged, first-person point of view, is epistemically reasonable. When a person asks herself “Can I specify any way in which S is false”, and, after addressing that question from her own first-person point of view, comes to conclude that she cannot, her conclusion is about a topic—her failure to be able to specify any way

---

<sup>1</sup>What if a person cannot specify any way in which S may be true and she cannot specify any way in which S may be false? In such a case, I think, it is natural to say that she does not understand S. Whether it is natural to say this or not, however, I shall assume it as a constraint on our understanding of 'understands' as it occurs in (C3).

in which S is false—that seems irrelevant to the question whether or not S. As I argued above, however, the proper point of view from which to clarify the sense in which a belief that S is epistemically reasonable is the one we occupy when we regard the question whether to believe that S as transparent for us to the question whether or not S. Hence the conditions specified in the antecedent of (C3) seem irrelevant to the question whether a person’s acceptance of S, when seen from her engaged, first-person point of view, is epistemically reasonable.

How then can there be *any* way to satisfy constraints (C1)-(C3)?

This is not a question of conceptual analysis, but of explication. To answer it we need to find a way to apply our constraints (C1)-(C3) that addresses, and avoids, the central problem of confusing facts about our *psychological* limitations with aspects of our inquiries in virtue of which our acceptance of unipolar statements is epistemically reasonable.

## 7. Impure entitlements and an analogy with perception

Following Crispin Wright, one might think of our entitlement to believe that no sentence of the form ‘S and not S’ is true as like an entitlement to perceptual belief—an entitlement that presupposes that our senses are reliable unless we have some special reason in a given context for doubting that they are. For Wright, an entitlement is “a kind of warranted acceptability which originates quite otherwise than in the existence of *evidence* for the truth of the proposition accepted.” (Wright 2004, p. 163, n5) If we equate “evidence for the truth of the proposition accepted” with “reasons for the truth of the proposition accepted”, in the sense of ‘reason’ constrained by (C1), then Wright’s generic characterization of entitlement is compatible with constraints (C1)-(C3). Wright’s account of entitlement goes beyond anything settled by these constraints, however. The heart of what is distinctive about his account of entitlement is his account of its role in our cognitive projects. According to Wright,

We should view each and every cognitive project as irreducibly involving certain elements of adventure. In the end, I have to take on trust — and in that sense, *take a risk on* — the reliability of my senses and cognitive powers in general just as I take a risk on the continuing reliability of the steering, and the stability of the road surface every time I ride my bicycle. (Wright 2004, p. 162)

These examples suggest that entitlements, in Wright’s sense, amount to *epistemic permissions* to take risks on the reliability, or the continuing reliability, of our senses and other “cognitive powers”. I call such entitlements *impure*, since we cannot have them in cases where the notions of risk or reliability do not apply.

The best ordinary examples of reliable belief-forming mechanisms of the sort that may yield impure entitlements are those we associate with perception. Let us therefore now investigate whether we can specify a sense in which my belief that no sentence of the form ‘S and not S’ is true, and other beliefs like it, is transparent for me by looking for a

similarity between impure *perceptual* entitlements and the entitlements constrained by (C2) and (C3).

One way of trying to specify such a similarity is suggested by the different roles that our senses play in cases in which we see that there is an F and cases in which we see there is no F, where ‘F’ stands in for some predicate.

Recall the simple perceptual case in which I examine a given page, see that *this* (indicated word) is an occurrence of ‘inquiry’ and infer that there is at least one occurrence of ‘inquiry’ on the page. I believe that my senses are reliable, but do not need any justification for this belief in order to see that *this* (indicated word) is an occurrence of ‘inquiry’. In short, I have an impure entitlement for relying on my perceptual belief that *this* (indicated word) is it an occurrence of ‘inquiry’.

Now consider a page on which there is no occurrence of ‘inquiry’. Suppose that when I examine the page, I do not take myself to see any such occurrence, and I therefore conclude that there is there is no such occurrence. In these circumstances, one might say, I *see that* there is no occurrence of ‘inquiry’ on the page. But there is a crucial difference between this use of “see that” and the first one. In the first, I both see *x*, for some *x* that is an occurrence of ‘inquiry’, and *see that* this (indicating *x*) is such an occurrence. In the second, there is no such *x* for me to see, and we say that I *see that* there is no such *x*, where we understand this to mean, in part, that if there *were* such an *x*, I *would have* seen it, and I *would thereby have seen that* it is an occurrence of ‘inquiry’.

More generally, in the first sort of case, I both see *x*, for some *x* that is an F, and *see that* this (indicating *x*) is an F. In the second sort of case, there is no *x* that is F for me to see, and we say that I *see that* there is no F, where we understand this to mean, in part, that if there *were* an *x* that is F, I *would have* seen it, and I *would thereby have seen that* it is an F.<sup>2</sup>

With the above counterfactual distinction in mind, then, consider the following analogy:

---

<sup>2</sup>One might think that one could eliminate the counterfactual formulation “if there *were* an *x* that is F, I would have seen it, and I *would* thereby have seen that it is an F” if we take “I see that there is no F” to mean “For every *x* that is a candidate for being an F, I saw *x* and that *x* was a G (for some G) such that *x*’s being G precludes its being F.” This suggestion appears promising, but I don’t think it can succeed. For suppose I accept this analysis, and I conclude, for a given predicate F, that for every *x* that is a candidate for being an F, I saw *x* and that *x* was a G (for some G) such that *x*’s being G precludes its being F. I should then be prepared to rule out two possibilities: (first) there is a candidate for being an F that I did not see, and (second) that there is such a candidate and it is an F. To rule out these possibilities, I would have to convince myself (first) that *there is no candidate for being an F that I did not see*, and (second), *there is no candidate for being an F that I did not see to be a G (for some G) such that *x*’s being G precludes its being F*. To accept either of these claims it seems that I must suppose that *if there were a candidate for being an F on the page, I would have seen it, and I would have thereby seen that it is a candidate for being an F on the page*. But this italicized claim is of the form “If there were an H, I would have seen it, and I would have thereby seen that it is an H”, where ‘H’ stands in for ‘*a candidate for being an F*’. The proposed analysis therefore does not ultimately avoid the sort counterfactual formulation I proposed above.

Just as we sometimes take ourselves to *see that* there is no F, in the sense that we fail to see any *x* that is an F, but would have seen such an *x*, and would have seen that *x* is an F, if there were such an *x* there to see, so, for some statements S, we take ourselves to *see that* S, in the special sense that we fail to see any way in which S may actually be false, but *would have* seen such a way, and recognized it as such a way, if there were one there to see.

The analogy is imperfect in at least three ways. First, where S is the statement that *no sentence of the form 'S and not S' is true*, the natural semantic comparison to the visual case would be that we fail to see any *sentence* of the form 'S and not S' that is not true, not that we fail to see any *way* in which the statement that *no sentence of the form 'S and not S' is true* may actually be false. For our purposes, however, this semantic analogy is not relevant, because we do not take the statement to be justified by induction from our observations about the truth values of particular sentences of the form 'S and not S'. Second, in any ordinary perceptual sense of "see", a *way* in which a sentence or statement may actually be false is not the sort of thing that we see, and, though we can be said to *see that* there is such a way, the "see that" here is not perceptual, but closer in meaning to "realize that". Third, if we suppose that there is an F that we fail to see, then we are committed to concluding that the statement that there is no F is false, whereas if we suppose that there is a way in which a given statement S may actually be false, we may still coherently suppose that S is true. This last point does not entirely spoil the analogy, however, since each supposition (that there is an F that we fail to see, and that we fail to see a way in which S may actually be false) has an *epistemic* consequence. Our supposition that there is an *x* that is F that we fail to see commits us to denying that have any *conclusive reason* for believing that there is no F, and our supposition that there is a coherent way in which a given statement S may actually be false commits us to concluding that we have no *entitlement* to believe that S.

On its most straightforward reading, the analogy seems to require that the entitlements constrained by (C2) and (C3) be impure in the way that perceptual entitlements are. In an ordinary perceptual case, we explain how our failure to see an *x* and to see that *x* is F is compatible with the transparency of our belief that there is no such F in part by citing the reliability of our senses. Such explanations are externalist in the sense that to understand them is to adopt, or at least imagine oneself adopting, a third-person point of view on the perceptual entitlement, a point of view from which the reliability of the perceiver's senses could be described.

The problem with this approach is that if we are unable to specify any way in which a given statement may actually be false, we are also unable to adopt, or imagine ourselves adopting, a third-person point of view from which we can describe our acceptance of the statement as the result of an ability to respond to considerations that we take to be relevant to answering the question whether or not S, one way or the other. To imagine such a perspective is to suppose there is a "cosmic exile" of the sort that our methodological assumptions about inquiry rule out.

The problem, in short, is that for any statement S that we take to be unipolar, we are by assumption unable to distinguish between alternatives to which we could take our answer the question whether or not S to be reliably and reasonably responsive. This apparently prevents us from using the analogy to articulate a way in which being entitled to accept a given belief, in the sense constrained by (C1)-(C3), is compatible with the supposition that the belief is transparent for us.

It might appear that this problem arises only if we to suppose, wrongly, that to be entitled to accept a given belief in Wright's sense one must be in position to *specify* the reliable cognitive mechanisms that explain one's entitlement. In fact, however, the problem does not rest on this false supposition. The problem, instead, is that *as theorists of entitlement*, we have no idea of how a belief (whether ours or someone else's) could be viewed as the result of a reliably reasonable response to alternative considerations if we are unable to distinguish between alternative considerations to which we could take our answer the question whether or not S to be responsive. One might think that in one's role as a theorist of entitlement, one can always suppose that there are alternative considerations to which our beliefs are responsive. But recall, again, that we are unable to specify any ways in which our belief that no sentence of the form 'S and not S' is true may be false—any alternatives to the case in which it is actually true. One cannot coherently and simultaneously suppose both that *there are* such alternatives and that *one cannot make sense of their being* any such alternatives. For this and other similar beliefs, we cannot give any theoretical content to the idea that the beliefs are the results of reliably reasonable responses of the sort that explain our perceptual entitlements.

## 8. Conceptual Schemes

Hilary Putnam, whose examples from the history of science convinced many philosophers to abandon methodological analyticity for the reasons I sketched above, has long insisted that "... the difference between statements that can be overthrown by merely conceiving of suitable experiments and statements that can be overthrown only by conceiving of whole new theoretical structures—sometimes structures, like Relativity and Quantum Mechanics, that change our whole way of reasoning about nature—is of logical and methodological significance, and not just of psychological interest." (Putnam 1962, p. 249; see also Putnam 1975, Putnam 1983a-b, and Putnam 1994.) Let us now consider whether we can recruit some of Putnam's suggestive remarks about theoretical structures, or what he also sometimes calls conceptual schemes, in an effort to satisfy constraints (C1)-(C3).

Putnam's central thesis is that for each of us, there are some statements such that (first) to specify a way in which they may actually be false, we would need to develop a substantially new conceptual scheme, and (second) we do not now know even how to begin on this project. Suppose there are such statements, and that S is one of them. Then, according to Putnam, the supposed facts about S that I just listed *explain* why it is epistemically reasonable for us to accept S without having any reason (in the sense of (C1)) to believe S. I call this the *conceptual scheme explanation* of (C3). Given (C2), it may seem to amount to an explanation of why we are entitled to believe S.

Like the account of impure entitlements that I discussed above, Putnam's conceptual scheme explanation is best suited to situations in which we can adopt a third-person point of view on the beliefs we seek to explain. According to Putnam, for example, it was reasonable for a scientist in the eighteenth century to hold immune to disconfirmation the statement that physical space is Euclidean, a statement that we now regard as false. The interest of such cases for Putnam is that

Before the development of general relativity theory, most people, even most scientists, could not imagine any experiences that would lead them to give up, or that would make it rational to give up, Euclidean geometry as a theory of actual space.... (Putnam 1983b, 99.)

By describing the methodological roles of such statements in another person's inquiries, or in our own inquiries at some previous time, Putnam aims to show that some statements are so basic for a person at a given time that it is both reasonable for her to accept them, and unreasonable for her to reject them unless she develops radically new theories and concepts.

I agree with Putnam that this sort of *retrospective* assessment of our, or another person's, acceptance of statements must be part of an account of when it is epistemically reasonable for a person to accept a statement, even one that we now regard as false. Unfortunately, however, an appeal to conceptual schemes does not help us to understand our own *current* entitlements to accept a statement S that we cannot now coherently suppose to be false. When it comes to our own current entitlements, Putnam's approach faces a dilemma. For either (i) the conceptual schemes explanation is just another way of expressing our observation that we cannot now specify a way in which S may be false, in the sense that when we try, we fail to be able to specify such a way, or (ii) it aims to go beyond that observation—it aims to explain the observation as a consequence of features of our conceptual scheme—and hence falls afoul of a version of the point, noted above, that there is “no cosmic exile”.

If (i), the conceptual schemes explanation has no more content than our observation that we are unable to specify a way in which S may be false, and hence does not clarify or explain that observation.

If (ii), however, then the conceptual schemes explanation fares no better than our earlier attempt to regard our failure to be able to specify a way in which a given statement may actually be false as analogous to our failure to see that, for instance, there is no occurrence of the word ‘inquiry’ on a given page. The problem with that earlier attempt was that if we are unable to specify any way in which a given statement S may actually be false, we are also unable to adopt, or imagine ourselves adopting, a third-person point of view from which we can describe our acceptance of S as the result of an ability to respond to alternative considerations that we take to be relevant to the question whether or not S. There is a similar problem with option (ii). To make sense of option (ii) we would have to be able to imagine a third-person point of view on our own system of

beliefs, a point of view from which the features of our conceptual scheme—beliefs or inference patterns—that we would have to revise in order to specify a way in which a given statement S may actually be false can be identified and described *as obstacles* to our specifying such a way. The problem is that to see some of our beliefs or inference patterns as *obstacles* to specifying a way in which S may actually be false, we would have to be able to specify some way in which S may actually be false. But if we could do this, we would not take ourselves to be entitled to accept S, in the sense of “entitlement” constrained by (C1)-(C3). We can accept the conceptual schemes explanation only if we aim to explain our entitlement to accept a statement S by imagining a third-person point of view on our situation that we cannot imagine if we have that entitlement. This aim is contradictory, and hence option (ii) falls apart. And this leaves us with option (i), where we began.

#### 9. An elucidation by subtraction

Recall that my goal is to elucidate the idea that beliefs we cannot now coherently suppose to be false are transparent for us by finding some way to satisfy constraints (C1)-(C3), and that the Central Problem that must be solved to achieve this goal is that the conditions described in the antecedent of (C3) are conditions about a person that seem irrelevant to the question whether or not S is transparent for her, and that therefore *also* seem irrelevant to whether her acceptance of S, when seen from her engaged, first-person point of view, is epistemically reasonable. We have now seen that we cannot solve the Central Problem by searching for a relevant similarity between our acceptance of beliefs that we cannot now coherently suppose to be false, on the one hand, and perceptual beliefs, on the other, or by conjecturing that we would have to overcome deeply entrenched conceptual obstacles in order to specify a way in which such a belief may actually be false.

I propose a different approach. I propose that we elucidate the sense in which beliefs we cannot now coherently suppose to be false are transparent for us by *subtracting* the sorts of confusions engendered by the two approaches we have examined above from our own engaged, practical understanding of the role in our inquiries of such beliefs.

The argument against methodological analyticity makes clear that if transparency is understood from within our inquiries and expressed in uses of sentences that we find clear, it goes hand-in-hand with an openness, in-principle, to revising any of our beliefs, even those that we cannot now coherently suppose to be false. We are strongly tempted to try to explain our current inability to specify any way in which a given statement or belief may actually be false in such a way that from the explanation, it follows that no reasonable person could *ever* come to specify a way in which the statement or belief may actually be false. But if we accept the primacy and authority of our ongoing inquiries, such efforts are as hopeless as trying to lift oneself up by one’s bootstraps. If we accept the methodological approach to epistemology, and claims (I) and (II), as I have interpreted them, we are committed to regarding all our beliefs as in principle revisable, in a methodological sense that can only be grasped and elucidated from the point of view

of a person engaged in inquiry, by combining a disquotational understanding of truth with a rejection of methodological analyticity.

To say a belief we cannot now coherently suppose to be false is revisable in this sense is to express a methodological stance that goes hand-in-hand with our rejection of methodological analyticity. It is not to say that we can specify a way in which the belief may be false, or even that we can specify a way in which we might eventually come to be able to specify a way in which the belief may actually be false, but to express one's openness, in principle, to reevaluating and revising it.

This methodological stance toward beliefs we cannot now coherently suppose to be false is all that is left when we subtract the confusions exposed above about the sense in which such beliefs are transparent for us from our own engaged, practical understanding of their role in our inquiries. I conclude that for a statement *S* that we cannot now coherently suppose to be false, to take one's answer to the question whether or not *S*, and one's corresponding belief that *S*, to be transparent to whether or not *S*, is just to find the answer obvious, in the ordinary sense that on considering whether or not *S* one accepts *S*, while simultaneously committing oneself to being open, in principle, to reevaluating and revising one's belief that *S*. This is a special sort of commitment that we can only properly express by focusing on the question whether or not *S* from our own first-person, deliberative point of view. By elucidating in this way the sense in which the question whether or not *S*, and our corresponding belief that *S*, can be transparent us to the question whether or not *S*, we can skirt the Central Problem I raised above, and thereby also elucidate a notion of entitlement that satisfies constraints (C1)-(C3). For statement *S* that we cannot now coherently suppose to be false, I conclude, our *entitlement* to believe *S*—or, as we sometimes say, our *seeing that S*—comes to this: on considering whether or not *S*, we accept *S*, even take *S* to be obvious, and, although we are unable to specify any way in which *S* may actually be false, we are open, in principle, to discovering such a way, and, in response to such a discovery, reevaluating and revising our belief that *S*.

## References

- Edgley, Roy, 1969. *Reason in Theory and Practice*. London: Hutchinson University Library.
- Evans, Gareth, 1982. *The Varieties of Reference*. Oxford: Clarendon Press.
- Moran, Richard, 2001. *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton, New Jersey: Princeton University Press.
- Priest, G. 1998. "What is so Bad about Contradictions?" in *The Journal of Philosophy* 95/8 (1998), pp. 410-426
- Putnam, H. 1962, "It Ain't Necessarily So," reprinted in Hilary Putnam, *Mathematics, Matter, and Method: Philosophical Papers: Volume 1* (Cambridge: Cambridge University Press, 1975) 237-249. First published in *The Journal of Philosophy*, 59/22 (1962).
- 1983a "Two Dogmas' Revisited," in Hilary Putnam, *Realism and Reason: Philosophical Papers: Volume 3* (Cambridge: Cambridge University Press, 1983) 87-97.
- 1983b "There is at Least One A Priori Truth," in Hilary Putnam, *Realism and Reason: Philosophical Papers: Volume 3* (Cambridge: Cambridge University Press, 1983) 98-114.
- 1983c "Analyticity and Apriority: Beyond Wittgenstein and Quine," in Hilary Putnam, *Realism and Reason: Philosophical Papers: Volume 3* (Cambridge: Cambridge University Press, 1983) 115-138.
- 1994 "Rethinking Mathematical Necessity," in Hilary Putnam, *Words and Life*, edited by James Conant (Cambridge, Mass.: Harvard University Press, 1994) 245-263.
- Wittgenstein, Ludwig 1921: *Tractatus Logico-Philosophicus*, translated by D. F. Pears and B. F. McGuinness. London: Routledge and Kegan Paul, 1961. First published German in *Annalen der Naturphilosophie*, 1921.
- Wright, C. 2004. "Intuition, Entitlement and the Epistemology of Logical Laws," in *Dialectica* 58/1 (2004), pp. 155-175.