

Perceptual Reasons and Defeat

Perceptual evidence about the external world is paradigmatically defeasible. If something looks red to you, it is reasonable to believe that it is red, but if you are wearing rose-tinted glasses, it may not be reasonable at all to believe this, unless you have some independent source of evidence. In this paper, I will compare four models for how to understand this phenomenon. These models differ in their answers to two questions: *what* evidence we get about the external world through perception, and what our having that evidence *consists* in. I like one of these models better than the others, but in this paper my primary concern will be to compare their virtues and vices.

I.1 Background: Reasons and Evidence

I am going to assume in this paper that perceptual experience is a way of acquiring evidence about the world. I will assume without argument that it is more reasonable to believe things that are supported by evidence than to believe without any evidence, or at least that if that is not true, then that is one of those surprising discoveries that philosophy is often claimed to allow us to make. I am also going to move back and forth freely between claims about reasons and claims about evidence. That is because I believe and argue elsewhere¹ that the correct explanation for why it is more reasonable to believe things that are supported by evidence is that evidence is reason to believe, and the only kinds of reasons to believe are evidence.

The claim that evidence is reason to believe is supported by the fact that the same things that are appropriately cited as evidence are also appropriately cited as reasons to believe, and conversely. ‘Reason’ and ‘evidence’ talk also exhibit parallel distinctions between *objective* talk about reasons or evidence that are out there, to be discovered, and *subjective* talk about the reasons or evidence that someone *has*, or from her perspective. It is also possible to do or to believe things for reasons, and similarly, it is possible to believe things on the basis of evidence. These facts, which could be enumerated at much greater length, all constitute strong evidence that reasons and evidence are closely related.

¹ Schroeder [manuscript].

From an ontological perspective, many kinds of thing can be cited appropriately as reasons. For example, the height of the Empire State Building is a reason not to jump off, Robin Jeshion is a reason to study philosophy of language at USC, and the bruise on Alyssa's leg is reason to believe that she has gotten hurt. Heights, people, and bruises all belong, I take it, to different categories. But it is also true that one reason not to jump off of the Empire State Building – and not, intuitively, a different reason – is that it is so high. Similarly, one reason to study philosophy of language at USC is that Robin Jeshion teaches there. And one reason to believe that Alyssa has gotten hurt is that she has a bruise on her leg. So although it makes perfect sense to cite many kinds of thing as reasons, in each case there are adequate paraphrases of what has been reported which attribute the reason in question using a 'that' clause. This may support the conclusion (I think it does) that strictly speaking, reasons should be identified with what is expressed by 'that' clauses – which I take to be propositions. But whether or not it supports this strong conclusion about the ontology of reasons, at least it supports the conclusion that it is a condition of adequacy for any view about what the reason is for someone to do something that it withstand paraphrase in terms of a 'that' clause.

Finally, the same points go for evidence as for reasons. The height of the Empire State Building is evidence that it took a long time to build, Robin Jeshion is evidence that philosophers of language can shed light on rich and subtle social issues, and Alyssa's bruise is evidence that she has gotten hurt. Again, the evidence that it took a long time to build the Empire State Building is that it is so high. Similarly, the evidence that philosophers of language can shed light on rich and subtle social issues is that Robin Jeshion has done so in her work on slurs and derogation. This again illustrates the close parallels between talk about evidence and talk about reasons, and I'll rely on this point in what follows.

I.2 Perceptual Knowledge and Defeat

In this paper I am concerned with the defeasibility of perceptual evidence. But we should be more careful in describing the phenomena in which we will be interested, of which there are two. Both phenomena involve cases in which it appears that someone would know but for some fact. But the kinds of fact but for which she would know are different in each case.

Objective Defeat	Roberta has a visual experience as of something red in front of her. But she is wearing rose-colored glasses.
Subjective Defeat	Roberta has a visual experience as of something red in front of her. But she rationally believes that she is wearing rose-colored glasses.

In Objective Defeat, Roberta cannot know that there is something red in front of her unless she has some independent source of evidence that that is so. Even if she does not realize that she is wearing rose-colored glasses, since they make everything look red, they are incompatible with Roberta having the discriminatory capacity to know that there really is something red in front of her, even if there is. In Subjective Defeat, it seems on the face of it that Roberta cannot reasonably believe that there is something red in front of her, since she rationally believes herself to be in Objective Defeat, and in that case she cannot know. But it seems that knowledge requires reasonable belief, so if it is not reasonable for her to believe it, then it seems that she cannot know it, either – again, unless she has some independent source of evidence, a qualification which I will henceforth ignore.

I want to be a little bit careful about how I describe the data about these two cases, since some of the views that I go on to discuss will disagree about how to describe them. My own view is that both cases can and should be taken at face value – that there are two dimensions along which knowledge can be defeated, both by the facts, independently of what the subject believes about them, and by the subject's other beliefs. I will call these two dimensions *objective* defeat and *subjective* defeat, respectively. I also believe that these two dimensions are independent – cases of objective defeat without subjective defeat are recipes for *Gettier* cases, and cases of subjective defeat without objective defeat undermine the justification or, as I have been putting it in this paper, the *reasonability* condition on knowledge. But as we will see later, some theorists are going to struggle to accommodate subjective defeat. They may still accept that there is an appearance of defeat of knowledge in cases of subjective defeat, but explain this appearance in some other way, because something else nearby, but somewhat different, is true.

2.1 The Classical Theory – What

According to what I will call the *classical theory* of the defeasibility of perceptual evidence, it can be properly subsumed to the non-monotonicity of inference. Non-monotonic inferences are cases where R is a reasonable conclusion to draw from P, but not a reasonable conclusion to draw from P&Q. For example, if you know that Tweety is a bird, it is reasonable to draw the conclusion that Tweety probably flies, but if you also know that Tweety is a penguin, then this is not a reasonable conclusion to draw.

Non-monotonicity in inference is possible whenever inferences are reasonable without being entailed. So according to the classical theory, the evidence about the external world that you get when you have a visual experience as of something red in front of you does not entail anything about the external world. It is, for example, that it *appears to you* that there is something red in front of you, or that you have a *visual experience as*

of something red in front of you, that you are *appeared to redly*, that *there is a red sense-datum*, or the like. There are many ways of developing the classical theory by filling in a more detailed account of the nature of perceptual evidence, according to the classical theory, but for purposes of this paper, they share a package of virtues and vices.

The classical theory is, in the first instance, a theory about subjective defeat. This is because non-monotonicity is a property of reasonable inference, and only subjective defeat concerns reasonable inference, at least directly. The key idea of the classical theory is that it is reasonable to infer that there is something red in front of you from the evidence that it appears that there is something red in front of you, but not from the evidence that you are wearing rose-colored glasses and it appears that there is something red in front of you.

But the classical theory can also be extended to offer a treatment of objective defeat. This was essentially the project of defeasibility analyses of knowledge, such as those offered by Klein [1970], Lehrer and Paxson [1969], Annis [1973], Ackerman [1974], Olin [1976], and Levy [1977]. What all of these theorists noted, was that there seems to be a general pattern whereby knowledge has to be not only reasonable belief, but belief whose reasonability “stands up” to the facts, in some way.² The project explored in their papers was how to say more precisely and in a way free from counterexamples exactly how knowledge must be belief whose reasonability “stands up” to the facts, but all of their accounts agree that this is what goes wrong in cases of objective defeat like Roberta’s that are not also cases of subjective defeat. In such cases, Roberta may reasonably believe that there is something red in front of her, and this belief may be reasonable because she does not realize that she wearing rose-colored glasses, but it is *only* reasonable because she is not aware of this. If you add this proposition to her stock of beliefs, it would no longer be reasonable for her to believe that there is something red in front of her, and that is *why*, on this view, its truth defeats her knowledge.

The strength of the classical theory is the elegance with respect to which it subsumes the defeasibility of basic perceptual evidence under the defeasibility of evidence elsewhere, by incorporating it into a general account of non-monotonic inference. Indeed, the literature on nonmonotonic inference in general is replete with examples that take for granted that the classical theory is correct about the defeasibility of perceptual evidence. Its home case is the case of subjective defeat, and although defeasibility theorists in the 1970’s famously discovered that it was hard to pin down exactly what the “standing up to the facts” condition on reasonable belief is that is required for knowledge, so it is intelligible to worry that they will not actually be

² Compare Schroeder [2015].

able to adequately capture all and only the correct cases of objective defeat, classical theorists at least have a well-understood strategy for trying to pin these cases down.

2.2 Consequences of the Classical Theory

However, the weakness of the classical theory has been regarded by many to be the key assumption about the nature of perceptual evidence that is required in order for it to subsume perceptual defeasibility to the monotonicity of inference more generally. And this is the assumption that perceptual evidence does not entail anything about the external world. This assumption, in essence, set all of the central problems of epistemology in the twentieth century, and shaped the space of possible options.

It is the assumption that perceptual evidence does not entail anything about the external world from which it follows that there are skeptical scenarios that are completely consistent with the totality of all perceptual evidence. If perceptual evidence could entail things about the external world, then no such skeptical scenarios would be possible, after all. There would still be scenarios in which we are all brains in vats, or deceived by an evil demon, of course, but these scenarios would be inconsistent with the totality of our perceptual evidence, and so the skeptic's claim that these are compatible with all of our evidence would be false. So in a very real sense, the classical theory's assumption is responsible for a very important aspect of the allure or threat of skepticism about the external world.

The same gap that is occupied by skeptical scenarios is one that non-skeptics, of course, would like to close. And most of the rest of twentieth-century epistemology can be conceived of as responses to how to close this gap. If a bridge premise is needed to close the gap, for example, then it seems that that bridge premise must be either empirical or a priori. If it is empirical, then it must in turn be justified by perceptual evidence, but since by parity of reasoning, that would again require a bridge premise, this path leads toward coherentism. And this is, in fact, the most pressing motive toward adopting coherentism. In contrast, if the bridge premise is a priori, then it turns out that a priori justification is a pre-requisite to empirical knowledge, which amounts to a very strong kind of rationalism.

The way to resist both coherentism and rationalism without retreating into skepticism, therefore, and without giving up on the classical theory, is to hold that no bridge premise is required, in order to make the inference from perceptual evidence to beliefs about the external world, even though these conclusions are not entailed. This is what dogmatists claim. But in order to avoid skepticism, dogmatists must claim that inferences from perceptual evidence to conclusions about the external world are *more* reasonable than inferences to conclusions about the intentions of the evil demon, or about the locations of ones and zeros on

the hard drive of the matrix. Since nothing in the content of the perceptual evidence breaks this symmetry, something else must be appealed to in order to explain the break in symmetry, and this is where epistemological externalism originally came in, in the late 1960's and early 1970's – it looked like every non-skeptical, non-coherentist, non-rationalist view had to give some causal, subjunctive, reliabilist, or otherwise fundamentally externalist answer to what justifies moving from basic perceptual evidence to external world beliefs over alternative conclusions. This was, as Armstrong [1973] argued for these reasons, inevitable. And so, externalists argued, everyone needed to accept the adequacy of such externalist explanations, which obviated the need to explain knowledge or the reasonability of belief in terms of evidence at all.

I don't mean to claim here that any of these views – skepticism, coherentism, rationalism, or pure externalism – are false, let alone that they are so obviously problematic that we need to reject the classical theory of the defeasibility of perceptual evidence in order to avoid them. But the fact that this familiar dialectic and set of choices are a consequence of the classical theory is the most important thing to understand about the classical theory, in order to come to grips with whether it is a package that is worth accepting. So let's turn to what some of its alternatives might look like.

3.1 Doubly World-Implicating Views

Since the classical view both requires and is enabled by the assumption that perceptual evidence does not entail anything about the external world, its alternatives adopt the assumption that perceptual evidence *does* entail things about the external world. They assume, as I will put it, that perceptual evidence is *world-implicating*. But world-implicating views disagree amongst themselves, as we will see, about exactly what sort of content perceptual evidence is, and also about what sort of condition a subject must satisfy, in order to possess this sort of evidence.

Recall that classical theorists believe that perceptual evidence does not entail anything about the external world. But classical theorists also believe that being in possession of some particular piece of perceptual evidence does not entail anything about the external world. If your perceptual evidence that there is something red in front of you is that it appears to you that there is something red in front of you, then you count as having this perceptual evidence in virtue of it's being *true* that it appears to you that there is something red in front of you. Or perhaps you count as having this evidence in virtue of *believing* that it appears to you that there is something red in front of you. Either way, not only does the *content* of your evidence not entail anything about the external world, the *condition* of your having that evidence does not entail anything about the external world, either.

The first class of world-implicating views that I want to consider disagrees with the classical theory on both of these fronts. These views claim not only that the *content* of your perceptual evidence is world-implicating, but that the relationship that you must stand in to that evidence in order for it to be evidence that can make beliefs reasonable *for you* is world-implicating as well. Let us call these views *doubly* world-implicating.

One example of a doubly world-implicating view of perceptual evidence comes from Williamson [2000]. Williamson holds that your evidence is what you know ($E=K$), that knowledge is the most general factive stative attitude, and that seeing that there is something red in front of you is a factive stative attitude. So it follows that Williamson holds that when you see that there is something red in front of you, that makes available to you as part of your evidence the fact that there is something red in front of you. The *content* of your evidence is a proposition about the world outside of your head, so this view is world-implicating along the content dimension. But the *condition* of this being your evidence is also world-implicating – it is that you *see that* this is the case. And seeing that something the case is a factive relationship to the world – a relationship that only holds if there really is something red in front of you. Let us call this view the *factive content* view.

Another example of a doubly world-implicating view comes from an interpretation of some of John McDowell's views offered by Comesaña and McGrath [2015]. According to this view, which I'll call the *factive attitude* view, in normal cases your evidence that there is something red in front of you is *that you see that there is something red*. And you come to have this reason because it is true – you really do see that there is something red in front of you. The factive attitude view shares with the factive content view a theory about the conditions under which you have perceptual evidence about the world, but disagrees about what that evidence is. But like the factive content view, it holds that both the content of your evidence and the condition of your having that evidence entail something about the external world. Both entail something about the external world because seeing that is a factive relation, and so you cannot see that there is something red in front of you unless there really is something red in front of you. So it is doubly world-implicating.

Doubly world-implicating views avoid the consequences of the classical theory explored in section 2.2. But they do so by undermining the classical theory's elegant subsumption of the defeasibility of basic perceptual evidence to the nonmonotonicity of inference more generally. The inference from 'there is something red in front of me' to 'there is something red in front of me' is not nonmonotonic, and neither is the inference from 'I see that there is something red in front of me' to 'there is something red in front of me'. This non-monotonicity is precisely what avoids the gap between perceptual evidence and conclusions about the external world that the skeptic exploits and that the coherentist, rationalist, and externalist all try to close,

so there is no avoiding the conclusion that avoiding the consequences of the classical theory also means giving up on its elegant treatment of the defeasibility of perceptual evidence.

But doubly world-implicating views – at least, both of the ones that I’ve described here – come with an alternative treatment of perceptual defeasibility. And that is because someone who is wearing rose-colored glasses cannot see that something is red. Seeing that something is red does not only entail that it really is red – it is also incompatible with having insufficient perceptual discriminability to discriminate red from other colors. Since rose-colored glasses interfere with perceptual discriminability, wearing them is incompatible with seeing that something is red. So if you are wearing rose-colored glasses, you do not know that something is red, just by having a visual experience as of it being red.

3.2 Consequences of Doubly World-Implicating Views

Recall that the classical theory is most at home in its treatment of subjective defeat. That is because what it explains, in the first instance, is why certain inferences are unreasonable, in the presence of further background beliefs. In contrast, we can now see that doubly world-implicating views are most at home in their treatment of objective defeat. That is because these views impose an objective condition on the subject’s relationship to the world, in order for her to have a reason at all – that she *see that* something is the case. And seeing that something is red is undermined by wearing rose-colored glasses.

But subjective defeat involves background beliefs that might (at least, when subjective defeat does not overlap with objective defeat) be false. Having a background belief that you are wearing rose-colored glasses does not prevent your vision from actually working properly. So it does not prevent you from seeing that there is something red in front of you. So whereas the doubly world-implicating views get an immediate explanation of why cases of objective defeat are cases in which the subject lacks reasons for her belief (and hence don’t know), they have no such obvious explanation for why cases of subjective defeat are cases in which agents lack reasons for their beliefs.

However, defenders of doubly world-implicating views can say something derivative about cases of subjective defeat. They can say, for example, that although in cases of subjective defeat, the subject really does know, perhaps it is reasonable for her to believe that she does not know, or that she probably does not know. Alternatively, they can defend the principle that knowledge is incompatible with reasonably believing yourself not to know. So since cases of objective defeat are incompatible with knowledge, and cases of subjective defeat involve reasonably believing yourself to be in a case of objective defeat, cases of subjective defeat will be derivatively incompatible with knowledge. Like the method by which classical theorists attempt

to transfer their account of subjective defeat over to objective defeat, it is intelligible to doubt whether any of these accounts will quite work out, but it is very clear that there are a range of overlapping strategies available, here.

More worrisome is the *way* in which doubly world-relative accounts explain the defeat of knowledge in the objective case. They explain it by showing that the subject has no reason at all to believe, in those cases. But reasons seem to be more directly connected to reasonability than to knowledge. So doubly world-relative accounts seem to predict not only that someone whose perceptual belief suffers from objective defeat does not *know*, but that in the absence of other evidence, her belief is not *reasonable*, either. They distinguish, that is, between ‘good’ and ‘bad’ cases of perceptual belief not only with respect to whether the subject knows, but also with respect to whether, and why, the subject’s belief is *reasonable*.

Now, it is open to doubly world-relative accounts to give an *alternative* account of the evidence that is available to a subject who suffers from objective defeat, in order to explain why her belief is reasonable after all, even though she lacks the evidence that is possessed by someone in the good case. A theorist who goes this way is an *epistemological disjunctivist* – that is, she tells different stories about the reasonability of perceptual beliefs in the good and bad cases. Another way of being an epistemological disjunctivist is to deny that perceptual beliefs are reasonable in the bad case of objective defeat at all, and explain why they *seem* reasonable by pointing out that the subject does not know that they are not reasonable.

So doubly world-relative views of the defeasibility of perceptual evidence have twin vices – they have to explain why knowledge seems to be defeated in subjective defeat cases, and they have to explain why reasonability seems *not* to be defeated in objective defeat cases, and if they do try to explain why beliefs can be reasonable in objective defeat cases, they must do so without undermining their own explanation for why the subject doesn’t *know*, in such cases. These are clearly different costs and advantages from the classical theory.

4.1 Singly World-Implicating Views

The classical theory and doubly world-implicating views disagree about two things. They disagree about whether the *content* of perceptual evidence can be world-implicating, and they disagree about whether the *condition* that a subject must satisfy, in order to have perceptual evidence, can be world-implicating. But the features of their view that doubly world-implicating views cite as their advantage over the classical theory derive from their view about the *content* of perceptual evidence, and the features that cast their treatment of defeasibility into doubt derive from their view about the *condition* requires to possess that evidence. The idea

behind singly world-implicating views is to split this difference by adopting the view that the contents of perceptual evidence are world-implicating, but denying that the condition required in order to possess this evidence is world-implicating.

Let's take these points one at a time. To see that the advantages of the doubly world-implicating views turn on what they say about the *content* of perceptual evidence, it suffices to recall that the problem with the classical view was that it required the assumption that skeptical hypotheses can be compatible with the totality of your perceptual evidence. It is necessary and sufficient to face this problem that the *content* of perceptual evidence – the proposition such that your perceptual evidence can be accurately reported as being *that p* – entails something about the external world.

In contrast, the problems for the doubly world-implicating views turn on what they say about the condition of *having* perceptual evidence. Since they say that this condition entails something about the external world, it follows that this condition cannot be shared across good and bad cases, and hence that the story about why belief is reasonable must be different across the good and bad cases, since the same evidence is not available to justify it. In contrast, if the condition under which a thinker counts as *having* a certain piece of evidence entails nothing about the external world, then it will cross-cut the good case/bad case distinction, and so we will be able to tell a non-disjunctive story about why thinkers are reasonable to believe in the way that they do.

So that is what motivates singly world-implicating views. But there is one major reason why philosophers have not explored singly world-relative views. And that is that it is common to assume that evidence must be *true*. If evidence must be true, then singly world-implicating views are impossible. For if evidence entails something about the external world, and having that evidence entails that it is true, then the condition of having the evidence must entail something about the external world, as well. So proponents of singly world-implicating views must reject the view that evidence (or for that matter, reasons) must be true.

This can sound bizarre, but it is not such an unnatural view. We may distinguish between objective and subjective reasons, and correspondingly between objective and subjective evidence. Objective reasons count in favor of what it is advisable to believe, but subjective reasons count only in favor of what it is reasonable to believe. Objective reasons must be true, because what it is advisable to believe (or do) depends on what is true, but subjective reasons need only be believed (perhaps rationally), because what it is reasonable to believe (or do) depends only on what you (rationally) believe.

4.2 Comparing Singly World-Implicating Views

We can distinguish between two obvious varieties of singly world-implicating views of the defeasibility of perceptual evidence – one analogue of each of the doubly world-relative views considered in section 3. Recall that according to the *factive content* view, your perceptual evidence when you see that there is something red in front of you is that there is something red in front of you, and the condition under which this is your evidence is that you see that there is something red in front of you. If we relax the condition that the perceptual condition for having this evidence must be a *factive* perceptual relation, we are led to the *non-factive content view*, according to which your perceptual evidence is the same – namely, that there is something red in front of you – but the condition under which this counts as your evidence is that you have a visual experience *as of* something red in front of you.

Similarly, recall that according to the *factive attitude view*, your perceptual evidence when you see that there is something red in front of you is that you see that there is something red in front of you, and the condition under which this is your evidence is that it is *true*. Again, we can relax the world-implicating condition of *truth* that this view requires in order for this to be part of your evidence, and replace it with some other, more subjective condition, such as that it *seems* true to you. Let us call this the *non-factive attitude view*.

Of all of the views that have been discussed here, the non-factive content view offers by far the least promising account of the defeasibility of perceptual evidence. Let us take the case of objective defeat first. In this case, you have a visual experience as of something red in front of you, but unbeknownst to you, you are wearing rose-colored glasses. According to the non-factive content view, you have available as evidence the proposition that there *is* something red in front of you. But this evidence is both available to you, and *true*. So it is hard to see why it can't ground knowledge that there is something red in front of you, especially because you have no contrary evidence. But it shouldn't be possible to know, on the basis of visual evidence, that there is something red in front of you, if you are wearing rose-colored glasses. The non-factive content view has no obvious answer to this problem. It *loses* the elegant answer that the doubly world-relative views gave to this problem, because it relaxes the constraint on having evidence too far.

The non-factive content view also fails to adequately treat cases of subjective defeat. In this case, you have a visual experience as of something red in front of you, but you also realize that you are wearing rose-colored glasses. According to the non-factive content view, as with the doubly world-implicating views, this latter belief does not interfere in any way with your satisfying the condition required to have world-implicating evidence – in this case, your evidence being that there is something red in front of you. The

doubly world-implicating views, as we saw, might attempt to try to explain why it is nevertheless not rational to believe that there is something red in front of you in a way that was derivative from their treatment of objective defeasibility. But as we've just seen, the non-factive content view does not have a successful treatment of objective defeasibility. And so it can't borrow this strategy, either.

In contrast, the non-factive attitude view offers elegant, parallel, treatments of both objective defeasibility and subjective defeasibility. It is able to do this, in effect, because it borrows *more* from the doubly world-implicating views – it takes on board their idea that factive perceptual relationships to the environment play an important role in basic perceptual evidence. The fact that you are wearing rose-colored glasses defeats your visual evidence that there is something red in front of you, on this view, because it cannot be *true* that you see that there is something red in front of you if you are wearing rose-colored glasses. This suffices to defeat your knowledge, because knowledge cannot be based on false lemmas. But it doesn't defeat the reasonability of your belief, because reasonability of belief is not defeated by false lemmas. So this is exactly the right kind of explanation of the objective defeat to avoid leading to epistemological disjunctivism.

The non-factive attitude view also offers an elegant treatment of subjective defeat. In this case, you believe that you are wearing rose-colored glasses, but part of your evidence about the world is that you see that there is something red in front of you. But as we've been reminded again and again throughout this paper, these claims cannot both be true. So either you should give up your belief that you are wearing rose-colored glasses, or you shouldn't infer anything from the proposition that you see that there is something red in front of you. So either your belief that you are wearing rose-colored glasses is not rational, or, if it is, it is not reasonable to infer anything from this particular piece of perceptual evidence, and hence it is defeated as a sort of reasonable belief about the world.

5 Summing Up

We've now explored four accounts of the defeasibility of perceptual evidence: the classical account, the two doubly world-implicating views, and the non-factive attitude view. (The non-factive content view, though it occupies a place in the logical space occupied by the other views, can hardly be described as an account of the defeasibility of perceptual evidence, since its problem is precisely that it has no such account.) Classical views have the advantage of fitting seamlessly into general accounts of non-monotonic inference, but carry with them a commitment to a restricted space of possible options in general epistemology. Their challenge is to extract an adequate treatment of objective defeat from their elegant treatment of subjective defeat. Both doubly world-implicating views lead to commitments about the differences in what explains reasonable belief

in good and bad cases. Their challenge is to extract an adequate treatment of subjective defeat from their elegant treatment of objective defeat. The non-factive content view shares neither of these commitments, and it offers parallel elegant treatments of both objective and subjective defeat. Are those sufficient reasons to think that it is true? Sufficient, I think, to wonder, at least in the absence of defeating evidence.³

³ Special thanks to Jessica Brown, Mona Simion, and Shyam Nair.

References

- Ackerman, Terrence [1974]. 'Defeasibility Modified.' *Philosophical Studies* 26(5-6): 431-435.
- Annis, David [1973]. 'Knowledge and Defeasibility.' *Philosophical Studies* 24(3): 199-203.
- Armstrong, David [1973]. *Belief, Truth, and Knowledge*. Cambridge: Cambridge University Press.
- Barker, John A. [1976]. 'What You Don't Know Won't Hurt You?' *American Philosophical Quarterly* 13(4): 303-308.
- Comesaña, Juan, and Matthew McGrath [2015]. 'Perceptual Reasons.' *Philosophical Studies* 173(4): 991-1006.
- Johnsen, Bredo [1974]. 'Knowledge.' *Philosophical Studies* 25(4): 273-282.
- Lehrer, Keith, and Thomas Paxson [1969]. 'Knowledge: Undefeated Justified True Belief.' *Journal of Philosophy* 66: 225-237.
- Levy, Stephen [1977]. 'Defeasibility Theories of Knowledge.' *Canadian Journal of Philosophy* 7(1): 115-123.
- Olin, Doris [1976]. 'Knowledge and Defeasible Justification.' *Philosophical Studies* 30(2): 129-136.
- Schroeder, Mark [2015]. 'Knowledge is Belief for Sufficient (Objective and Subjective) Reason.' *Oxford Studies in Epistemology* 5: 226-252.
- _____ [manuscript]. *Reasons First*. Book manuscript in progress.
- Williamson, Timothy [2000]. *Knowledge and Its Limits*. Oxford: Oxford University Press.