

Strawson, Ordinary Language, and the Priority of Holding Responsible over Being Responsible

Mark Balaguer

ABSTRACT: It is often held that P. F. Strawson endorsed a radical and groundbreaking priority thesis according to which *holding someone morally responsible* is prior to (or more fundamental than) *being morally responsible*. I do three things in this paper. First, I argue for a novel interpretation of Strawson according to which he did not endorse a priority thesis that is radical or groundbreaking or original; instead, Strawson's "priority thesis" is just a consequence of his view that the meanings of our words are determined by our usage and intentions and practices concerning those words. Second, I argue against the radical priority thesis that is often (erroneously) attributed to Strawson. Third, I argue that while Strawson's view does not involve a radical priority thesis, it does imply that debates about the nature of moral responsibility (and many other debates about normative ethics, metaethics, and conceptual analysis) are *trivial* in a certain sense.

KEY WORDS: Strawson, moral responsibility, ordinary language, conceptual analysis

Mark Balaguer is Professor of Philosophy at Cal State LA. He is the author of five books: Metaphysics, Sophistry, and Illusion (Oxford University Press, 2021), Free Will (MIT Press, 2014), Free Will as an Open Scientific Problem (MIT Press, 2010), Platonism and Anti-Platonism in Mathematics (Oxford University Press, 1998), and Mathematical Anti-Realism and Modal Nothingism (Cambridge University Press, 2022). He is also the author of numerous papers on a wide range of philosophical topics, including the philosophy of mathematics, metaontology, free will, the philosophy of time, the philosophy of language, and metaethics. <https://orcid.org/0000-0003-2799-7608>; mbalagu@exchange.calstatela.edu.

1. INTRODUCTION

Let a *priority thesis* be any thesis that says that *holding morally responsible*—that is, the thing we do when we hold people morally responsible for things—is prior to (or more fundamental than, or grounds, or some such thing) *being morally responsible*. Some people think that Strawson (1962) endorsed a priority thesis that (a) is *radical*, or *groundbreaking*, or some such thing (for example, Watson [1987: 220] says that it is a “radical claim”¹) and (b) entails that the property *being morally responsible* is very special—different from ordinary properties like *being a bachelor*, *being an electron*, and so on.

In this paper, I do the following four things. In section 2, I develop an interpretation of Strawson—which I call the *ordinary-language interpretation*—on which Strawson endorsed a pretty *pedestrian* priority thesis, a priority thesis that (a) is not at all radical (or groundbreaking, or original, or anything of the sort) and (b) applies to properties like *being a bachelor* and *being an electron* as well as to *being morally responsible*. In section 3, I argue that there is no radical or groundbreaking priority thesis—no priority thesis that is stronger or more robust than the pedestrian priority thesis that I articulate in section 2—that is even remotely plausible. In section 4, I argue that my interpretation of Strawson (that is, the interpretation developed in section 2) is correct. And finally, in section 5, I argue that while Strawson’s view does not involve a radical priority thesis, it does have some interesting and surprising consequences that have gone largely unnoticed. The consequences I have in mind do not have much to do with the nature of moral responsibility; they have more to do with the status of *philosophical debates* about moral responsibility. In particular, I argue in section 5 that if the view that I attribute to Strawson is correct (or, indeed, if the radical-priority interpretation is correct), then debates about the nature of moral responsibility (and, indeed, lots of other debates about normative ethics, metaethics, and conceptual analysis) are *trivial* in a certain sense. *That*, I think, is the surprising and important consequence of Strawson’s view.

2. THE ORDINARY-LANGUAGE INTERPRETATION OF STRAWSON

The ordinary-language interpretation of Strawson (or for short, *the OL interpretation*) says that Strawson endorses (something like) what I will call *the ordinary-language view* (or *the OL view*).² I have argued elsewhere (in Balaguer 2021) that the OL view, or something like it, is *true*. I will not try to argue for that claim here; in what follows, I will just (a) articulate the view, (b) argue that Strawson endorsed the view, and (c) argue that the view does not entail any radical priority thesis but does entail a sort of deflationism about certain kinds of philosophical debates. In any event, the OL view is the conjunction of five theses, the first of which is the following:

Thesis 1: The concept *morally responsible* is the concept picked out by the expression ‘morally responsible.’ And there is nothing special here about *morally responsible* and ‘morally responsible’; we could make exactly analogous claims about *red* and ‘red,’ *electron* and ‘electron,’ *knowledge* and ‘knowledge,’ *water* and ‘water,’ and so on.

This, I think, is pretty trivial. But, again, it is not going to matter here whether the five theses that make up the OL view are actually *true*. The second component of the OL view is the following:

Thesis 2: Which responsibility-like concept is picked out by the expression ‘morally responsible’ is determined by facts about *us*—in particular, by facts about our usage, intentions, and practices concerning the expression ‘morally responsible.’ And once again, there is nothing special here about ‘morally responsible.’ In general, we can say that if ‘C’ is a concept-expressing term in our language, then the concept picked out by ‘C’ is determined by facts about our usage, intentions, and practices concerning ‘C.’ If you like, you can think of the point here in terms of *meaning determination*; in particular, if we take the concept expressed by a (concept-expressing) term to be its meaning, then thesis 2 is just a special case of the general thesis that the meanings of our words are determined by facts about our usage, intentions, and practices concerning those words. So, for example, if you ask whether ‘bachelor’ means *unmarried man* or *eligible unmarried man* or something else, the answer to this question is determined by facts about our usage, intentions, and practices concerning ‘bachelor.’ If ‘bachelor’ expresses the concept *unmarried man* rather than *eligible unmarried man* or some other concept, then that is because of *us*—because *we use* ‘bachelor’ to express that concept. And analogous claims can be made about ‘knowledge,’ ‘water,’ ‘morally responsible,’ and so on.

Putting theses 1 and 2 together, we get the following:

Thesis 3: Which responsibility-like concept counts as the concept *morally responsible* is determined by facts about *us*—about our usage, intentions, and practices concerning the expression ‘morally responsible.’ Suppose, for example, that our usage, intentions, and practices pick out a libertarian concept, that is, a concept that requires a libertarian sort of free will; then that *makes it the case* that this libertarian concept *is* the concept *morally responsible*. (And by the way, I am not thinking here of a scenario in which we *consciously intend* to pick out a libertarian concept of moral responsibility; I am thinking instead of a scenario in which our usage, intentions, and practices pick out a concept that turns out, *in a way that is not obvious to us*, to require libertarian free will. What thesis 3 says is that if our usage, intentions, and practices pick out a libertarian concept of this kind, then that makes it the case that the libertarian concept in question *is* the concept *morally responsible*.) Likewise, if our usage, intentions, and practices pick out a compatibilist concept of moral responsibility—that is, a concept that requires only a compatibilist kind of free will—then that *makes it the case* that this compatibilist concept *is* the concept *morally responsible*. Finally, as before, there is nothing special here about *morally responsible* and ‘morally responsible’; exactly analogous claims can be made about other cases. For example, if our usage and intentions concerning ‘bachelor’ pick out the concept *unmarried man*, then that makes it the case that that concept *is* the concept *bachelor*.

Shifting gears a bit, the fourth component of the OL view is the following:

Thesis 4: As native speakers of English, we are very good at *using* expressions like ‘knowledge,’ ‘water,’ ‘morally responsible,’ and so on; in other words, we are good at knowing when these expressions apply and do not apply. But we are not very good at knowing the definitions of these expressions. We can, however, *figure out* (or at least make progress on figuring out) what these expressions mean (or what concepts they express) by using our judgments—or our *intuitions*—about when they apply and do not apply. For example, in Gettier cases, we judge that ‘knowledge’ does not apply;

and in Putnam-style Twin-Earth cases, we judge that ‘water’ does not apply. Thus, since we are good at knowing when our expressions apply and do not apply, we can use these intuitions as data points to confirm and falsify theories of what these expressions mean. For example, our Gettier intuitions falsify the justified-true-belief theory of the meaning of ‘knowledge’; and our Twin-Earth intuitions falsify the theory that ‘water’ is a non-rigid term that means *clear tasteless liquid*, and they confirm the theory that ‘water’ is a rigid designator.

The fifth and final component of the OL view is as follows:

Thesis 5: In connection with the term ‘morally responsible,’ in addition to using our detached, impersonal judgments (or intuitions) about when ‘morally responsible’ applies and does not apply, we can also use our *reactive attitudes* to figure out what ‘morally responsible’ means. Thus, for example, if you feel resentment (or moral indignation) toward a person S concerning some act A, then that is a data point about the concept *morally responsible*—or about which concept is picked out by the expression ‘morally responsible.’ It is a data point in more or less the same way that it would be a data point if you had an intuition that S is morally responsible for A. And, again, this is analogous to the way that our Gettier intuitions are data points about the concept *knowledge* and our Twin-Earth intuitions are data points about the concept *water*.

Also, now that we have brought in the idea of the reactive attitudes, it is important to note that we should understand thesis 3 as already being partly about the reactive attitudes. I characterized thesis 3 above as the thesis that which responsibility-like concept counts as the concept *morally responsible* is determined by facts about our usage, intentions, and *practices* concerning the expression ‘morally responsible.’ But I am using the word ‘practices’ to refer to a set of phenomena that includes facts about our reactive attitudes, and I now want to add the following explicit addendum to thesis 3:

Addendum to Thesis 3: Which responsibility-like concept counts as the concept *morally responsible* is at least partly determined by facts about our reactive attitudes—that is, by facts about how we respond emotionally to the actions of other people. So, for example, it might be that facts about our reactive attitudes *make it the case* that some compatibilist concept (or some libertarian concept) is the concept *morally responsible*.

In sum, then, the OL view is just the conjunction of theses 1–5. And the OL interpretation of Strawson is just the view that Strawson endorsed the OL view, or something like the OL view.

If the OL interpretation is right, then Strawson endorsed the following priority thesis (which is essentially just a streamlined version of thesis 3):

Pedestrian-priority: Which concept counts as the concept *morally responsible* is determined by our practices of holding people responsible. And there is nothing special here about *morally responsible*; exactly analogous remarks can be made about concepts like *bachelor*, *electron*, *water*, and so on; for example, which concept counts as the concept *water* is determined by our usage, intentions, and practices concerning the term ‘water.’

It should be clear that this is a priority thesis. After all, on this view, which concept counts as the concept *morally responsible* is determined—not just epistemically determined, but *metaphysically* determined—by our practices of holding people morally responsible. As I

pointed out above, according to this view, if our practices concerning ‘morally responsible’ pick out the concept *C*, then that *makes it the case* that *C* is the concept *morally responsible*.

It is important to note, however, that the OL view (and Pedestrian-priority) are not radical or groundbreaking views. Now, of course, they are controversial—not everyone accepts them—but before we move on, I want to bring out four ways in which the OL view is not radical or groundbreaking.

First, according to the OL view, while facts about our practices determine which concept counts as the concept *morally responsible*, our practices do *not* determine the nature of that concept. Suppose, for example, that our practices pick out a compatibilist concept of responsibility, call it CR. Nonetheless, there is a nearby concept—a libertarian concept of responsibility, call it LR—that we *might have* employed. And if we *had* employed LR, then it would have been the case that we might have employed CR. So, CR and LR are both, so to speak, “there for the taking”—that is, they are both concepts that could be employed by a community of creatures. So, our practices do not determine the nature of CR or LR. All that our practices determine is *which* of these responsibility-like concepts is picked out by the English expression ‘morally responsible’; that is, they determine which of these concepts counts as the concept *morally responsible*.

Second, as I have made clear, according to the OL view, there is nothing special here about the concept *morally responsible*. Analogous remarks can be made about all of our concepts. For example, if we use the word ‘bachelor’ to mean *unmarried man*, then that *makes it the case* that the concept *unmarried man* counts as the concept *bachelor*—because it makes it the case that the concept *unmarried man* is the concept that is picked out by the word ‘bachelor.’

Third, despite the fact that the OL view entails a priority thesis (namely, Pedestrian-priority), it does not imply that the concept *morally responsible*—or *knowledge*, or *bachelor*, or whatever—is nonobjective in any interesting way. It is true that, according to the OL view, facts about our usage and intentions make it the case that the concept *unmarried man* (or whatever) counts as the concept *bachelor*; but whether a given object *O* is a bachelor is determined by objective facts about the nature of *O*—in particular, by whether *O* is an unmarried man (or an eligible unmarried man, or whatever). And according to the OL view, analogous remarks can be made about *morally responsible*, *water*, and so on. Facts about us determine which concepts are expressed by our predicates; but whether some object *falls under* a given concept is determined by objective facts about the nature of that object. So in the case of *morally responsible*, whether a specific person *S* is morally responsible for a specific action or outcome *A* depends on objective facts about *S* and *A*—in particular, on whether *S* and *A* satisfy some sufficient condition for moral responsibility, where the sufficient condition in question is determined by our practices of holding people morally responsible (so, for example, this might depend on facts about whether *S* cared enough about the well-being of the people affected by *A*).

Fourth, in terms of the sociology of the discipline of philosophy, the OL view just is not a radical or groundbreaking view. Thesis 5 is a genuinely novel and interesting twist on the view, but it is just that—a *twist*. The rest of the view, while controversial, is widely believed. *Lots* of philosophers have endorsed views along the lines of theses 1–4, and this was already true back in 1962, when Strawson published “Freedom and Resentment.” Indeed, if anything, views along these lines were more widespread back then, during what might be thought of as the heyday of ordinary-language philosophy.

(Before moving on, I would like to say a few words about Strawson's compatibilism. There is obviously more to Strawson's view of moral responsibility than is contained in the OL view—that is, in theses 1–5. For theses 1–5 do not yet entail compatibilism. To get to compatibilism, we have to add something like the following:

Thesis 6: In our practices, we do not treat the truth of determinism as a good reason to withdraw our feelings of resentment or indignation for cases of wrongdoing—or to withdraw our ascriptions of moral responsibility. And this is an important data point about the concept *morally responsible*. In particular, this gives us good reason to think that that concept—that is, the concept picked out by our usage, intentions, and practices concerning 'morally responsible'—is compatible with determinism.

Thesis 6 is obviously separable from theses 1–5; you could endorse theses 1–5 and then go on to reject thesis 6 and claim that *our* concept of moral responsibility—the one that is at work in *our* moral practices and *our* reactive attitudes—is *incompatible* with determinism. But, of course, that is not Strawson's view; he endorsed something like thesis 6. But I am not concerned in this paper with compatibilism or Strawson's stance on compatibilism; so, I am not counting thesis 6 as being part of the OL view.³)

3. WHAT IS THE RADICAL PRIORITY THESIS?

We now have one interpretation of Strawson on the table, namely, the OL interpretation, which takes him to endorse Pedestrian-priority. A second interpretation—which might be called the *radical-priority interpretation*—takes Strawson to endorse the following thesis:

Radical-priority: Holding *morally responsible* is prior to (or more fundamental than, or some such thing) *being morally responsible* in some way that is stronger or more robust than the way in which it is on Pedestrian-priority—that is, in some way that goes beyond the mere claim that which concept counts as the concept *morally responsible* is determined by our practices of holding people morally responsible (and our usage and intentions and practices concerning the expression 'morally responsible').

But it is not clear what the alleged "radical priority thesis" is even supposed to *be*.⁴ Indeed, I will argue in this section that there is no stable/tenable radical priority thesis there at all.

Let me start my argument for this by ruling out two theses that, I think, do not give us a radical-priority thesis of the kind we are looking for here. The two theses are as follows:

Anti-realism: There is really no such thing as someone being morally responsible for something—that is, there is no objective relation of moral responsibility that sometimes holds between persons and actions (or outcomes or whatever). There are just our practices of holding people responsible. There is nothing more—out there in reality—than these practices.⁵

Subjectivist-priority: Sentences of the form 'Person S is morally responsible for action or outcome A' just *mean* that most normal adults would hold S morally responsible for A—or something like that. Or to put the point differently, to say that S is morally responsible for A is to make a claim that is straightforwardly *about* the way that most normal adults would react to the situation regarding S and A.⁶

Anti-realism does not give us a priority thesis that is stronger or more robust than Pedestrian-priority because it does not give us a priority thesis *at all*. *Holding responsible* is not

prior to *being responsible* on this view; rather, if Anti-realism is true, then there is simply no such thing as someone being responsible.

Subjectivist-priority, on the other hand, is a priority thesis (and there is some sense in which it is a stronger, more robust priority thesis than Pedestrian-priority), but Subjectivist-priority is, I think, *not* the priority thesis that we are looking for here. The first point to note in this connection is that subjectivist views of morality are extremely implausible. In particular, they involve implausible empirical claims about what ordinary folk mean when they make moral claims. And Subjectivist-priority is no exception to this. The idea that when ordinary folk utter sentences of the form ‘S is morally responsible for A,’ what they literally say is that *most normal humans would hold S morally responsible for A* just does not hold up to empirical scrutiny. And it is hard to believe that advocates of the radical-priority interpretation (that is, people like Watson, Wallace, and Shoemaker) either (a) believe this subjectivist claim or (b) believe that Strawson believed it. So, while Subjectivist-priority might, in some sense, be a stronger priority thesis than Pedestrian-priority, it seems to me that it is (i) not tenable, and (ii) not what advocates of the radical-priority interpretation have in mind.

Another point worth noting here is that while there is a *sense* in which Subjectivist-priority gives us a stronger kind of priority than Pedestrian-priority does, there is also an obvious sense in which it does not. We can appreciate this by thinking about other predicates that seem to have subjectivist meanings. Consider, for example, ‘magnetic personality.’ This seems to have a subjectivist meaning; sentences of the form ‘Person S has a magnetic personality’ seem to mean that S has a personality that most normal humans would be attracted to—or something like that. But it is not as if there is some weird kind of priority at work in connection with the concept *magnetic personality*; all that is going on here is that the claim that S has a magnetic personality is *about* the ways that normal humans would react to S. And there is an obvious sense in which this is a perfectly objective claim—because (a) it is a claim about how normal people would react to S, and (b) there are perfectly objective facts about this (that is, about how normal people would react to S) that determine whether the claim is true.

Another way to appreciate the point I am making here is to notice that Subjectivist-priority is perfectly compatible with Pedestrian-priority. Pedestrian-priority is obviously compatible with the view that there are subjectivist predicates in our language (that is, predicates that apply to objects just in case normal humans respond (or would respond) to those objects in certain ways). In such cases, advocates of Pedestrian-priority will say that (a) facts about our usage and intentions determine that these predicates have the subjectivist application conditions that they do, and (b) whether a specific subjectivist predicate applies to a specific object O is determined by objective facts about whether ordinary humans really *do* respond (or really *would* respond) to O in the relevant way. So, again, Pedestrian-priority is perfectly compatible with Subjectivist-priority. And I think these considerations motivate the idea that while there is a sense in which Subjectivist-priority gives us a stronger kind of priority than Pedestrian-priority does, there is also a sense in which it does not. And I think they further motivate the claim that Subjectivist-priority is not what advocates of the radical-priority interpretation have in mind.

In what follows, I will assume that neither Anti-realism nor Subjectivist-priority gives us a priority thesis of the kind that we are looking for in connection with Radical-priority. The question I want to ask is whether there is any tenable priority thesis that (a) is stronger

or more robust than Pedestrian-priority and (b) does not just collapse into Anti-realism or Subjectivist-priority. In other words, I want to ask whether there is any tenable way to simultaneously endorse the following three theses:

Realism: People sometimes *are* morally responsible for their actions; that is, there is a real relation of moral responsibility that sometimes holds between persons and actions or outcomes.

Anti-Subjectivism: Sentences of the form ‘Person S is morally responsible for action or outcome A’ do not just *mean* that most normal adults would hold S morally responsible for A—or anything like this. In other words, to say that S is morally responsible for A is *not* to make a claim that is straightforwardly *about* the way that most normal adults would react to the situation regarding S and A. It is rather to make an objective (and non-subjectivist) claim about the situation regarding S and A itself.

Radical-priority: *Holding morally responsible* is prior to (or more fundamental than, or some such thing) *being morally responsible* in some way that is stronger or more robust than the way in which it is on Pedestrian-priority.

I do not think there is any viable way to simultaneously endorse these three views. My argument for this has two premises. Premise (I) says that if we endorse Radical-priority, then we also have to endorse the following:

Subjectivist Determination: Whether a specific person S is morally responsible for a specific action or outcome A is determined not by objective (and non-subjectivist) facts about the situation regarding S and A itself, but rather by whether most normal people would hold S morally responsible for A.

And premise (II) says that there is no viable way to simultaneously endorse Realism, Anti-subjectivism, and Subjectivist Determination.

Let me argue for these two premises in turn. I have two arguments for premise (I). My first argument proceeds as follows:

Pedestrian-priority gives us the result that (a) the application conditions of ‘morally responsible’ are determined by our usage, intentions, practices, and attitudes; but (b) whether a specific person S is morally responsible for a specific action or outcome A is determined by objective facts about the situation regarding S and A—in particular, by facts about whether S and A satisfy the application conditions that are determined by our usage, intentions, practices, and attitudes. So, the priority thesis here—that is, the one contained in Pedestrian-priority—concerns clause (a) but not clause (b). Therefore, it seems that in order to obtain a priority thesis that is stronger or more robust than Pedestrian-priority—and, remember, that is exactly what we need to do if we want to endorse Radical-priority—then we will have to replace clause (b) with some sort of priority thesis. But the only way to do this, it seems, would be to say that whether a specific person S is morally responsible for a specific action or outcome A is determined not by objective (and non-subjectivist) facts about the situation regarding S and A, but by whether normal people (or most normal people, or some such thing) would hold S morally responsible for A. But this is just to say that we would need to endorse Subjectivist Determination. And so, premise (I) is true. (In short: if Pedestrian-priority is true, then the only relevant facts (that is, the only facts about *being responsible*) that are not determined by our practices of holding responsible are facts about *individual*

cases—that is, facts about which specific persons are responsible for which specific actions and outcomes—and so to obtain a stronger priority thesis, we would have to say that *those* facts are determined by our practices of holding responsible, which is precisely what Subjectivist Determination says.)

My second argument for premise (I) proceeds as follows:

Pedestrian-priority is a thesis of *meaning* determination. So, if we endorse Radical-priority, we are going to have to endorse a different kind of determination thesis, and the obvious thing to say here is that the relevant sort of determination is something like *grounding*. But grounding is most naturally thought of as a relation that holds between *facts*, not concepts. So the claim that we would be making, if we endorsed a grounding version of Radical-priority, would not be a claim about the concept *morally responsible*; it would presumably be a claim about *specific responsibility facts*, for example, the fact that some specific person S is morally responsible for some specific action or outcome A. More specifically, the claim would presumably be that responsibility facts of this kind—that is, *S-is-responsible-for-A facts*—are grounded not by objective (and non-subjectivist) facts about the situation in question (that is, the situation involving S and A), but rather by subjectivist facts about whether most normal people would hold S morally responsible for A. But this is essentially what Subjectivist Determination says. And so, premise (I) is true.⁷

You might respond to both of these arguments for premise (I) by claiming that advocates of Radical-priority can reject Subjectivist Determination and endorse the following thesis instead:

Normatively Loaded Priority: Whether a specific person S is morally responsible for a specific action or outcome A is determined not by objective (and non-subjectivist) facts about the situation regarding S and A, but rather by whether it would be *fair* (or *appropriate*, or some such thing) to hold S morally responsible for A.

But it is hard to even make sense of this claim. For, *prima facie*, it seems that whether it would be fair to hold S responsible for A is determined by objective (and non-subjectivist) facts about the situation regarding S and A.⁸ But advocates of Normatively Loaded Priority obviously cannot say this, since that would take them right back to Pedestrian-priority. And they also cannot say that the fairness fact here (that is, the fact about whether it would be fair to hold S morally responsible for A) is determined by subjectivist facts about whether most normal people would hold S morally responsible for A—because that would take them back to Subjectivist Determination. But if they cannot say that the fairness fact is determined by objective (and non-subjectivist) facts about the situation regarding S and A, and if they also cannot say that the fairness fact is determined by subjectivist facts about whether most normal people would hold S responsible for A, then what *can* they say? I have no idea—there seem to be no other facts for them to appeal to here. The only other option, it seems, would be to claim that the fairness fact is *fundamental*—that is, that *nothing* determines whether it would be fair to hold S morally responsible for A. But that claim is incompatible with the whole idea behind Radical-priority. And so, I do not think that advocates of Radical-priority can help their cause at all by appealing to Normatively Loaded Priority.

(Shoemaker [2017: 508] endorses the following thesis: “The blameworthy . . . *just is* whatever merits anger (the angerworthy); that is, someone is blameworthy . . . for X if and

only if, *and in virtue of the fact that*, she merits anger for X.”⁹ Shoemaker writes as if this is a radical priority thesis, but I think it collapses into something like Pedestrian-priority. To see why, suppose that some specific person S is blameworthy for some specific action or outcome A. According to Shoemaker, this is grounded in the fact that S *merits anger* for A. But we can now ask what grounds the fact that S merits anger for A. The answer has to be either (i) objective (and non-subjectivist) facts about the S-A situation, thus taking us back to Pedestrian-priority, or (ii) subjectivist facts, for example, the fact that most normal humans would be angry at S for A, thus taking us back to Subjectivist Determination. And it is obvious which option Shoemaker favors here—*option (i)*. For Shoemaker endorses the existence of objective “blameworthiness-makers”—that is, objective conditions that make it the case that certain agents are blameworthy. Now, given this, you might wonder why Shoemaker thinks of himself as endorsing a response-dependent view—or, in my lingo, a priority thesis. The reason, Shoemaker tells us, is that he endorses the following claim: what makes the objective blameworthiness-makers *count* as blameworthiness-makers is that “they are just the sorts of properties to which we humans are built to respond with [anger]” [2017: 510]. But this is precisely the sort of thing that advocates of Pedestrian-priority—and the OL view—would say. Suppose that Strawson is right and the concept of moral responsibility is identical to some compatibilist concept C. According to the OL view, the reason this is true—the reason that C counts as the concept of moral responsibility, and other responsibility-like concepts (for example, libertarian concepts) do *not* count as the concept of moral responsibility—is that C is the concept that is picked out by *our* usage, intentions, practices, and attitudes. But as we have seen, this is just a thesis of meaning determination—and if it is true, it is true of *all* of our concepts. I do not think Shoemaker has given us a radical priority thesis; more precisely, I do not think he has given us a priority thesis that is any stronger or more robust than Pedestrian-priority. (Similar remarks can be made about Wallace because he too thinks that there are objective conditions for the appropriateness of the reactive attitudes—see, for example, Wallace [1994: 77–78]—but I will not take the time to argue for this here.))¹⁰

Let us move on now to premise (II), which, recall, says that there is no viable way to simultaneously endorse Realism, Anti-subjectivism, and Subjectivist Determination. My argument for premise (II) is based on the claim that if we try to simultaneously endorse these three views, then we will have to say things that are wildly implausible at best and possibly just incoherent. Here is an example of the sort of thing we will have to say:

Unstable View: (i) Realism is true—that is, there is a real relation of moral responsibility that sometimes holds between persons and actions or outcomes. Moreover, in particular, Steve is morally responsible for Jane’s injury—that is, Steve bears the morally-responsible-for relation to Jane’s injury. Furthermore, (ii) Anti-subjectivism is true, and so the claim that Steve is morally responsible for Jane’s injury is not a claim about the ways that normal adults would react to the situation regarding Jane’s injury; it is rather a claim about the situation itself—that is, it is a claim about Steve, Jane’s injury, the relations that hold between them, and so on. But (iii) whether Steve is in fact morally responsible for Jane’s injury is not *determined* by objective (and non-subjectivist) facts about the situation regarding Jane’s injury; it is rather determined by whether most normal people would hold Steve morally responsible for Jane’s injury.

I do not see how anyone could seriously endorse Unstable View. It may be that this view is incoherent; for it may be that claim (ii) analytically entails the falsity of claim (iii). But

regardless of whether Unstable View is incoherent, it seems wildly implausible. Moreover, I take it that this is *obvious*—that is, that Unstable View wears its implausibility on its sleeve. If the claim that Steve is responsible for Jane's injury is not *about* how normal people would respond to the situation—if it is about the situation *itself*—then why on Earth would the truth value of that claim be determined by how normal people would respond to the situation? It seems that it would be determined by objective (and non-subjectivist) facts about the situation itself. In particular, it seems that the truth value would be determined by whether the situation has the traits that the claim in question (that is, the claim of moral responsibility) *says* that it has. After all, if Realism and Anti-subjectivism are true, then the claim that Steve is responsible for Jane's injury *says* that the Steve-Jane situation is a certain way. And so, it seems that, in this scenario, the truth value of that claim would be determined by whether the Steve-Jane situation *is* that way.

So, I do not think there is any sensible way to simultaneously endorse Realism, Anti-Subjectivism, and Subjectivist Determination. And so, I do not think there is any tenable/stable way to endorse Radical-priority without collapsing into either Anti-realism or Subjectivist-priority—which, again, are not the kinds of priority theses that we are looking for in connection with Radical-priority.

In short, my claim is as follows. If there is a real relation of moral responsibility that sometimes holds between persons and actions or outcomes, and if the claim that a specific person S is morally responsible for a specific action or outcome A is a claim not about the reactions that normal humans would have to the S-A situation but about the S-A situation itself, then the truth value of that claim is determined by facts about the S-A situation itself (in particular, by whether the S-A situation is the way that the claim of moral responsibility says that it is) and not by facts about the reactions that normal humans would have to the S-A situation. And so, it seems to me that premise (II) is true.¹¹

I think that my argument here generalizes—that is, it applies not just to *morally responsible* but to *all* of our concepts. In other words, I do not think there are any ordinary-language concepts about which Realism, Anti-subjectivism, and Radical-priority are all true. And I would like to end this section by making two points about the concept of a *strike* in baseball—which some people seem to think is analogous to *morally responsible* in a radical-priority sort of way.¹²

First point: the word 'strike' (as it is used in the context of a baseball pitch that the batter did not swing at) seems to be *ambiguous*. On sense 1, it means *passed through the strike zone*; and on sense 2, it means *was called a strike by the umpire*. We seamlessly go back and forth between the two meanings. If I am keeping score at a baseball game, and I miss a pitch and ask you, "Hey, was that last pitch a ball or a strike?" then I clearly have sense 2 in mind. I do not care whether the pitch was in the strike zone; I just want to know whether the umpire called it a strike. But if I see that the umpire called a pitch a strike, and I turn to you and say, "*What the . . . ?* That was a terrible call. Do you think that pitch was a strike?," then I clearly have sense 1 in mind.

Second point: sense 2 is just a straightforward *subjectivist* sense, and the sort of priority that we get in connection with sense 2 is a straightforward subjectivist kind of priority, and so it is not a priority of the kind that we are looking for in connection with Radical-priority. So, there is nothing useful to advocates of Radical-priority here; on the contrary, if anyone should claim that *morally responsible* is analogous to *strike*, it is *subjectivists*, not Radical-priority theorists.

It is also worth noting here that in connection with ‘morally responsible,’ there does not seem to be any analogue of sense 2 of ‘strike.’ It is just an empirical fact about our language that sentences of the form ‘S is morally responsible for A’ do not say that some person(s) hold(s) (or would hold) S morally responsible for A (or that the *speaker* holds S morally responsible for A). And this, of course, is just to say that subjectivism about moral responsibility is false.

4. TEXTUAL EVIDENCE

I now want to argue for the OL interpretation of Strawson—that is, for the claim that Strawson endorsed something like the OL view and Pedestrian-priority. And I want to start by arguing that there is no good reason to reject this interpretation.

4.1. No Textual Evidence against the OL Interpretation

There are a few passages in “Freedom and Resentment” that some people have taken to motivate the view that Strawson endorsed some sort of anti-realism about moral responsibility or some sort of radical priority thesis. But these passages actually fit perfectly with the OL interpretation. Probably the most widely cited of these passages is the following:

Passage 1: Only by attending to this range of attitudes [that is, the reactive attitudes] can we recover from the facts as we know them a sense of what we mean, i.e. of *all* we mean, when, speaking the language of morals, we speak of desert, responsibility, guilt, condemnation, and justice. (Strawson 1962: 78)

This passage might seem at first glance to motivate the view that Strawson endorsed an anti-realist view of moral responsibility or a radical priority thesis. But it actually fits *better* with the OL interpretation than with anti-realist and radical-priority interpretations. According to the OL interpretation, Strawson held that our reactive attitudes can be used as data points to figure out what moral responsibility is, or what ‘morally responsible’ means; thus, passage 1 is precisely the sort of thing that Strawson might say if the OL interpretation were right.

As I read “Freedom and Resentment,” there are only two passages in that paper that even seem like they *might* be incompatible with the OL interpretation:

Passage 2: The making of the demand [that is, the moral demand] is the proneness to such attitudes [that is, disapprobation and indignation]. (Strawson 1962: 77)

Passage 3: The existence of the general framework of attitudes itself is something we are given with the fact of human society. As a whole, it neither calls for, nor permits, an external ‘rational’ justification. (Strawson 1962: 78)

In passage 2, Strawson seems to be saying that to (morally) demand that someone behave in a certain way just *is* to be disposed to feel certain ways if the person does not behave in the given way. And this might seem to push us toward the idea that Strawson endorsed an anti-realist view of moral responsibility or a radical priority thesis. But I think this is a mistake. Notice that passage 2 is a claim about moral *demands*. It is not a claim about moral assertions, for example, assertions to the effect that some person is morally responsible for some action or outcome. A demand has the form: “Do A.” Even full-blown moral realists would want to say that demands of this kind are not the sorts of sentences that express

propositions or have truth values. So whatever Strawson says about demands of this kind, it is not incompatible with the OL interpretation; for the OL interpretation does not take any stand at all on what Strawson thought about moral demands; it only takes a stand on what Strawson thought about assertions of moral responsibility.

What about passage 3? Strawson says here that our practices of holding people morally responsible do not require any external justification—and, in fact, *cannot* be externally justified. And this might seem to be suggestive of a priority thesis that is stronger than Pedestrian-priority. But, in fact, this is not true; as I will now explain, passage 3 fits quite well with the OL view.

According to the OL view, the concept *morally responsible* is the concept expressed by our term ‘morally responsible.’ But there are other responsibility-like concepts that we might have employed, and so you might think that our practice of employing the concept that we *do* employ—that is, the concept *morally responsible*—instead of some *other* responsibility-like concept stands in need of justification. To make this a bit clearer, let us pretend (to simplify things) that there is exactly one compatibilist responsibility-like concept and exactly one libertarian responsibility-like concept, and let us call these two concepts *compatibilist-responsible* and *libertarian-responsible*, respectively. According to a compatibilist version of the OL view (that is, according to the conjunction of the OL view and thesis 6), *compatibilist-responsible* is the concept that we express with the term ‘morally responsible,’ and so it *is* the concept *morally responsible*. But one might raise the following worry for advocates of this view:

The external-justification worry: Let us grant that we express the concept *compatibilist-responsible* with our term ‘morally responsible.’ Even so, the question remains whether we *should* express that concept with ‘morally responsible.’ If deterministic creatures do not *deserve* to be blamed or praised for their actions, then we *should* be employing the concept *libertarian-responsible*—that is, we should be using ‘morally responsible’ to express the concept *libertarian-responsible*. And so, our practice of employing the concept *compatibilist-responsible* stands in need of justification.

Advocates of the OL view can (and, I think, should) respond to this worry by saying something like the following:

Words like ‘deserve’ and ‘should’ are interdefinable with words like ‘responsible.’ It is true—and, indeed, *analytic*—that deterministic creatures do not libertarian-deserve to be blamed and praised for their actions. And it is also true (and analytic) that we libertarian-should use ‘morally responsible’ to express the concept *libertarian-responsible*. But it is also true (and analytic) that (a) we compatibilist-should use ‘morally responsible’ to express the concept *compatibilist-responsible*, and (b) deterministic creatures sometimes compatibilist-deserve to be blamed and praised for their actions. Now, you might wonder whether deterministic creatures ever *really* deserve to be blamed and praised for their actions; and you might wonder how we *really* should use ‘morally responsible.’ But if the OL view is true, then the meanings of ‘really deserve’ and ‘really should’ are determined by facts about *us*—about our usage and intentions and practices concerning the words ‘deserve’ and ‘should.’ So, we do not need an external justification of our moral practices, and in fact, the kind of justification that is being asked for in the above external-justification worry is impossible to achieve. All we can do, and all that we need to do, is to say, *in English*, that we *should* use our

moral words—that is, *really morally should* use them—in precisely the ways that we *do* use them.

This response to the external-justification worry assumes that the word ‘should’ is being used here in its *moral* sense—that is, to express a *moral* should. Now, we can also talk about an *all-things-considered* should, and you might counter the above response to the external-justification worry by saying something like the following:¹³

The all-things-considered worry: It might be analytic that we *morally* should use our moral words in precisely the ways that we do use them. But it is not analytic (or even obviously true) that we *all-things-considered* should use our moral words in the ways that we do use them. And so even if it is true that we currently use ‘morally responsible’ to express the concept *compatibilist-responsible*, we can ask whether we all-things-considered should continue to do this or whether we all-things-considered should change our practices and start using ‘morally responsible’ to express the concept *libertarian-responsible*.

To this, advocates of the OL-view can give a very Strawsonian response:

The idea that we could just *choose* to start expressing different concepts with our moral words is a fiction. And if we *did* have a choice about whether to do this, then—since, as we have seen, we do not have a *moral* reason to do it—the only rational way to settle the question of whether we all-things-considered should do it would be to think about what the gains and losses to our lives would be. And it is hard to see how the fact that deterministic creatures do not libertarian-deserve blame or praise for their actions would be relevant to our deliberations here—especially since (assuming that the compatibilist version of the OL view is true) deterministic creatures do sometimes *deserve* (that is, *really, morally* deserve) blame and praise for their actions.

Here is Strawson saying something very similar:

[I]f we could imagine what we cannot have, viz, a choice in this matter, then we could choose rationally only in the light of an assessment of the gains and losses to human life, its enrichment or impoverishment; and the truth or falsity of . . . determinism would not bear on the rationality of this choice. (1962: 70)

So, I think that passage 3 fits perfectly with the OL view and the OL interpretation of Strawson. And more generally, I do not think Strawson says anything in “Freedom and Resentment” that is incompatible with the OL interpretation.

4.2. Positive Evidence for the OL Interpretation

Strawson does not say in “Freedom and Resentment” that he endorses the OL view. But if we look at his other works, we find clear evidence that he endorsed something like the OL view. He is most clear about this in his 1969 paper “Meaning and Truth,” but he seems to have held something like this view throughout his career. For instance, in his 1950 paper “On Referring,” he says that “the meaning [of an expression] is the set of rules, habits, conventions for its use in referring” (10); and he says that to give the meaning of an expression—or as he also calls it in this paper, the *analysis* of an expression—is to “explain and illustrate the conventions governing the use of the expression.” In short, Strawson’s view in “On Referring” is that *to give the meaning or analysis of an expression is to give the conventional rules for using that expression*.

Strawson argues for an OL-style view in much more detail in “Meaning and Truth.” He argues there for a Gricean theory of meaning determination. To understand this theory, we first need to distinguish *speaker meaning* (or as Strawson calls it, *communication-intention*) from *expression meaning* (or as Strawson calls it, *linguistic meaning*). The latter is just the meaning of an expression—so this is what we are trying to construct a theory of when we try to specify the meaning of an expression, or to analyze the concept expressed by an expression. Speaker meaning, on the other hand, is psychological; in particular, it is *what a particular speaker means by an expression on a given occasion of use*. And like Grice (1957), Strawson is very clear that he takes this to be an *intention*; indeed, he calls it a *communication-intention*.

The central idea behind Grice’s theory of meaning is that we can give a theory of expression meaning by, *first*, giving a theory of speaker meaning that does not make use of the notion of expression meaning, and *then* giving a theory of expression meaning in terms of speaker meaning. Or as Strawson puts it, the idea is

to present your general theory of meaning in two stages: first, present and elucidate a primitive concept of *communication* (or communication-intention) in terms which do not presuppose the concept of *linguistic meaning*; then show that the latter concept can be, and is to be, explained in terms of the former. (1969: 172)

And later in the same paper, Strawson explains that the way to carry off the second part of this project is to develop a theory on which the expression meaning (that is, the linguistic meaning) of an expression is the set of conventional rules for how to speaker-mean things with that expression—or, as Strawson puts it, rules for how to use the expression to “fulfill . . . [your] communication-intentions” (1969: 173).

Strawson calls these rules “the meaning-determining rules of the language” (1969: 176). So in his view, the meaning of an expression is determined by the rules for using the expression, and these rules of use are determined by facts about *us*—in particular, by facts about our intentions, habits, conventions, and so on. And so, putting these two points together, we get the result that, according to Strawson, the meanings of our expressions are determined by these facts about us. And given this, it seems that Strawson endorsed the first three theses of the OL view—that is, theses 1–3 from section 2—or something very much like those theses; for those three theses essentially just say that the meanings of our expressions are determined by facts about our usage, intentions, practices and so on.

Now, I think that just about anyone who endorsed theses 1–3 would also endorse something like thesis 4. But in case there is any doubt about whether Strawson endorsed thesis 4, he says the following in his 1992 book *Analysis and Metaphysics*, while talking about conceptual analysis:

In one sense—to repeat an example—we understand the concept of knowing, we know what knowing is or what the word ‘know’ means; for we know how to use the word correctly. In one sense we understand the concept of personal identity, we know what sameness of person is, we know what the words ‘same person’ mean; for we know in practice how to apply the concept. . . . But in another sense, perhaps, we don’t understand the concepts, don’t know what personal identity is, can’t say what the word ‘know’ means. We have mastered a practice, but can’t state the theory of our practice. We know the rules because we observe them and yet we don’t know them because we can’t say what they are. In contrast

with the ease and accuracy of our use are the stuttering and blundering which characterize our first attempts to describe and explain our use. (7)

So, clearly, Strawson endorsed something very much like thesis 4 as well as theses 1–3. When we go to analyze a concept—for example, *knowledge* or *morally responsible* or whatever—what we are trying to do is construct a theory of something that is, in a very real sense, *in our heads*. We are trying to come up with a theory of how to use the words that express the given concepts. But as Strawson points out, we *already know* how to use these words; what we cannot do is *explain* how we use them; and that is what we are trying to do when we try to analyze a concept, or characterize the meaning of a word.

Finally, given that Strawson endorsed theses 1–4, it should be clear from what he says in “Freedom and Resentment” that he endorsed thesis 5 as well. Thesis 5 is really just an addendum to thesis 4 to the effect that we can use our reactive attitudes to figure out (or to help us figure out) what we mean by ‘morally responsible.’ And as evidence that Strawson endorsed this view, I can do no better than to requote the passage that is often used to motivate the claim that he endorsed a radical priority thesis:

Only by attending to this range of attitudes [that is, the reactive attitudes] can we recover from the facts as we know them a sense of what we mean, i.e. of *all* we mean, when, speaking the language of morals, we speak of desert, responsibility, guilt, condemnation, and justice. (Strawson 1962: 78)

In sum, then, I think we have good reason to endorse the OL interpretation of Strawson—that is, to think that Strawson endorsed theses 1–5, or something like those theses.

You might object here by claiming that even if the above considerations suggest that Strawson endorsed a general OL-style view, they are perfectly compatible with the claim that he endorsed a subjectivist or anti-realist (in particular, non-cognitivist) view of moral responsibility. For you could endorse a general OL-style view and then go on to claim that *we use* ‘morally responsible’ in a subjectivist or non-cognitivist way. But (a) there is no *evidence* that Strawson endorsed an anti-realist or subjectivist view of moral responsibility—he certainly did not say in “Freedom and Resentment” that he endorsed such a view—and (b) I think we have good reason to think that he did not endorse either of these views. Indeed, I think the fact that Strawson did not *say* that he endorsed either of these views already gives us good reason to think that he did *not* endorse either of them. Moreover, I think there is strong textual evidence for the positive claim that Strawson was a moral realist; but I will not try to argue for this here.

You might also object by claiming that that even if I am right that Strawson endorsed the OL view, I need to argue that this is *relevant* to his stance in “Freedom and Resentment”; more specifically, you might think I need to argue that Strawson intended, in “Freedom and Resentment,” to be applying his general theory of meaning and conceptual analysis to the case of moral responsibility. But I think we can reasonably infer that Strawson was doing this from the fact that he couched his discussion in terms of *what we mean* when we talk about moral responsibility, and more often, in terms of the *concept* of moral responsibility. Indeed, one of Strawson’s main goals in “Freedom and Resentment” was to make a point about the concept of moral responsibility. And given this, it seems reasonable to expect that his views of meaning and conceptual analysis would be relevant to understanding his view.

5. DEFLATIONISM ABOUT NORMATIVE ETHICS AND METAETHICS

If I am right, then Strawson did not endorse a radical priority thesis. But it is worth noting that Strawson's view (if my interpretation of him is correct) does have an important and largely unnoticed consequence. For if the OL view is true, then it follows that certain philosophical questions and debates are *trivial* in a certain sense. Consider, for example, the following questions:

Q1: What is moral responsibility? That is, what is the right analysis of the concept of moral responsibility?

Q2: Is indeterminism a necessary condition for moral responsibility?

Q3: Is cognitivism or non-cognitivism true of moral responsibility?

If the view that I am attributing to Strawson (that is, the OL view) is true, then these three questions (and other questions like them) are entirely settled by empirical facts about the usage, intentions, practices, and attitudes of ordinary folk. Not just *epistemically* settled, but *metaphysically* settled. If ordinary folk use 'morally responsible' to pick out a libertarian concept, then that *makes it the case* that the libertarian concept in question *is* the concept *morally responsible*; and if ordinary folk use 'morally responsible' to pick out a compatibilist concept, then that makes it the case that *that* concept is the concept *morally responsible*; and if ordinary folk use 'morally responsible' in a non-cognitivist way, then that *makes it the case* that non-cognitivism about moral responsibility is true; and so on.

I am not suggesting that this makes *Q1–Q3* *uninteresting*. My claim is just that if the OL view is true—or, indeed, if a radical priority thesis is true—then *Q1–Q3* are empirical questions about folk psychology. So, they are not philosophically *deep* questions. And they are not about mind-independent reality. On the contrary, they are entirely settled by contingent, empirical facts about the heads of ordinary folk—facts about how ordinary folk happen to use their words, and how they happen to react emotionally to various kinds of actions, and so on.

Similar remarks can be made about normative ethical questions about moral properties like *wrongness* and *goodness*. For example, if the OL view is true, then questions about whether utilitarianism is true (and whether Kantianism is true, and so on) are metaphysically settled by empirical facts about the usage, intentions, practices, and attitudes of ordinary folk. And likewise for various kinds of metaethical questions, for example, the question of whether cognitivism or noncognitivism is true of our talk of moral wrongness and moral goodness. And perhaps most notably, if the OL view is true, then similar remarks can be made about not just *moral* conceptual-analysis questions like 'What is moral responsibility?' and 'What is moral wrongness?', but *all* conceptual-analysis questions, regardless of whether they are about moral or nonmoral concepts—for example, questions like 'What is free will?', 'What is knowledge?', 'What is consciousness?', 'What is an essential property?', and so on.

If the OL view is true, then there is nothing philosophically deep about *any* of these questions; and there is nothing about the nature of mind-independent reality at issue in connection with these questions. In fact, if the OL view is true, then all of these questions boil down to empirical questions about ordinary folk—about how they happen to use their words, and how they happen to react emotionally to certain kinds of actions, and so on.

It is worth noting, however, that the OL view does *not* imply that *applied* ethical questions are settled by facts about folk meaning. On the contrary, if the OL view is true, then applied ethical questions are settled by objective facts about the actions and people and so on that these questions are about. For instance, if the OL view is true, then the question of whether meat-eating actions are morally wrong is settled by objective facts about whether these actions have the property of moral wrongness—that is, the property picked out by our usage, intentions, and practices concerning ‘morally wrong’; and the question of whether people who eat factory-farmed meat are morally responsible for the pain that the animals in those farms experience is settled by facts about whether these people stand in the moral-responsibility relation (that is, the relation picked out by our usage, intentions, and practices concerning ‘morally responsible’) to the pain that the animals in question experience; and so on.¹⁴

Moreover, the point that I am making here—that while the OL view implies that questions like *Q1–Q3* (and various other conceptual-analysis questions, and normative ethical questions, and metaethical questions) are settled by facts about folk meaning,¹⁵ it does not imply that applied ethical questions are settled by such facts—is deeply related to the central point of this paper, that is, the point that while the OL view implies a *pedestrian* priority thesis, it does not imply a *radical* priority thesis, and it does not lead to any substantive form of anti-objectivism about moral responsibility. Putting these two points together, we can say this: while the view that I am attributing to Strawson does not lead to any important sort of deflationism about moral responsibility, it does lead to an important sort of deflationism about certain kinds of philosophical questions about moral responsibility.

Acknowledgements

Thanks are due to Michael McKenna, David Shoemaker, and Patrick Todd for commenting on earlier versions of this paper. Thanks are also due to Galen Strawson for helpful discussion and for pointing me toward some of his father’s writings. Finally, I presented a version of this paper at a conference at Florida State University, and I’d like to thank the members of the audience there for providing useful feedback, especially Al Mele, Randy Clarke, and John Martin Fischer.

Notes

1. Others to endorse interpretations along these lines include Wallace (1994) and Shoemaker (2017). See also Fischer (1994: 211–13), Nelkin (2011: 28–30), McKenna (2012: 31–55), Beglin (2018), and Hieronymi (2020: 76ff.) for discussions of this view.
2. De Mesel and Heyndels (2019) interpret Strawson in a way that is similar in certain ways to the interpretation that I put forward in this section, but their interpretation is also different from mine in various ways.
3. Todd (unpublished manuscript) and De Mesel and Heyndels (2019) make similar points about compatibilism being separable from the priority thesis. For pushback against this separability, see (Beglin 2018).

4. Todd (2016) also expresses puzzlement about what exactly the radical priority thesis says; but the argument he develops is quite different from the argument I develop in this section.
5. Bennett (1980) thinks that Strawson endorsed an anti-realist (in particular, non-cognitivist) view of moral responsibility.
6. I do not know of anyone who thinks that Strawson endorsed Subjectivist-priority.
7. You might think that the relevant sort of determination is *truthmaking*, not grounding. But this would not change anything. For truthmaking is a relation that holds between sentences (or propositions) and facts, and so if we endorsed a truthmaking version of Radical-priority, we would presumably be saying that specific moral-responsibility sentences—for example, ‘Person S is morally responsible for action or outcome A’—are made true not by objective (and non-subjectivist) facts about the specific situation regarding S and A, but rather by facts about whether most normal people would hold S morally responsible for A. But, again, this is essentially what Subjectivist Determination says.
8. In other words, if it is *fair* to hold S responsible for A, that must be because of facts about *what S did* (and what S’s motivations were, and what the consequences of S’s actions were, and so on).
9. Wallace (1994: 91) endorses a nearly identical claim.
10. It is worth noting here that Pedestrian-priority is perfectly compatible with the view that moral responsibility—that is, the concept picked out by our usage, intentions, practices, and attitudes—requires all sorts of complicated interactions between agents and victims. For example, it may be that in order for Steve to be morally responsible for Jane’s injury, Steve has to know that Jane expects him to have some degree of good will toward her. Even if this is true, it does not necessarily give us a radical priority thesis because, again, it is perfectly compatible with Pedestrian-priority that this is built into the concept of moral responsibility that is picked out by our practices.
11. It is worth noting that these considerations hold regardless of whether we think of the relevant sort of determination as *truthmaking* or *grounding*. If the claim that Steve is responsible for Jane’s injury does not *say* that most normal people would hold Steve responsible for Jane’s injury—if it says something about the Steve-Jane situation *itself*—then, assuming that this claim is true, it is made true by, and grounded by, objective (and non-subjectivist) facts about the Steve-Jane situation itself, for example, by the fact that Steve did not care enough about Jane’s well-being, or some such thing.
12. Shoemaker made this comparison in private correspondence with me. The analogy is also discussed by Todd (unpublished manuscript).
13. Pam Hieronymi (2020: 57) makes a similar distinction between a moral worry and a non-moral worry, and she suggests that Strawson can respond to the two worries in ways that are similar to the ways in which I am saying they can respond.
14. The OL view does leave open the possibility that *some* applied ethical disputes are merely verbal. For example, it leaves open the possibility of an applied ethical dispute in which the two disputants use some moral word to pick out two different concepts, and the moral assertion under dispute is true in the language of one disputant and false in the language of the other, so that the two disputants could figure out who was right by figuring out what the relevant moral word meant *in English*. In such cases, we can say that the dispute could be *epistemically settled for the two disputants* by facts about

folk meaning. But (a) the dispute would still be *metaphysically* settled by objective (non-meaning) facts in the manner described in the text; and (b) for whatever it is worth, my own view is that merely verbal moral disputes like this are pretty rare—for example, typical disputes about veganism and abortion and the death penalty do not seem to me to be merely verbal.

15. The OL view does not imply that *all* metaethical questions are settled by facts about folk meaning. For example, if the OL view is true, then metaethical disputes between moral realists and error theorists about whether the property of moral wrongness is actually instantiated are settled not by facts about folk meaning, but by objective (non-meaning) facts about the world.

References

- Balaguer, M. 2021. *Metaphysics, Sophistry, and Illusion*. Oxford: Oxford University Press.
<https://doi.org/10.1093/oso/9780198868361.001.0001>
- Beglin, D. 2018. “Responsibility, Libertarians, and the ‘Facts as We Know Them’: A Concern-Based Construal of Strawson’s Reversal.” *Ethics* 128: 612–25.
<https://doi.org/10.1086/695990>
- Bennett, J. 1980. “Accountability.” In *Philosophical Subjects*, ed. Z. van Straaten, pp. 14–47. Oxford: Clarendon Press.
- De Mesel, B., and S. Heyndels. 2019. “The Facts and Practices of Moral Responsibility.” *Pacific Philosophical Quarterly* 100: 790–811.
<https://doi.org/10.1111/papq.12276>
- Fischer, J. M. 1994. *The Metaphysics of Free Will*. Oxford: Blackwell.
- Grice, H. P. 1957. “Meaning.” *Philosophical Review* 66: 377–88.
<https://doi.org/10.2307/2182440>
- Hieronymi, P. 2020. *Freedom, Resentment, and the Metaphysics of Morals*. Princeton, NJ: Princeton University Press.
<https://doi.org/10.23943/princeton/9780691194035.001.0001>
- McKenna, M. 2012. *Conversation and Responsibility*. Oxford: Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199740031.001.0001>
- Nelkin, D. 2011. *Making Sense of Freedom and Responsibility*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199608560.001.0001>
- Shoemaker, D. 2017. “Response-Dependent Responsibility; or, A Funny Thing Happened on the Way to Blame.” *Philosophical Review* 126: 481–527.
<https://doi.org/10.1215/00318108-4173422>
- Strawson, P. F. 1950. “On Referring.” In *Logico-Linguistic Papers*, ed. P. F. Strawson, 1–27. London: Methuen & Co. Ltd.
- Strawson, P. F. 1962. “Freedom and Resentment.” In *Free Will*, ed. G. Watson, 59–80. Oxford: Oxford University Press.
- Strawson, P. F. 1969. “Meaning and Truth.” In *Logico-Linguistic Papers*, ed. P. F. Strawson, 170–89. London: Methuen & Co. Ltd.
- Strawson, P. F. 1971. *Logico-Linguistic Papers*. London: Methuen & Co. Ltd.
- Strawson, P. F. 1992. *Analysis and Metaphysics*. Oxford: Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780198751182.001.0001>

- Todd, P. 2016. "Strawson, Moral Responsibility, and the 'Order of Explanation': An Intervention." *Ethics* 127: 208–40. <https://doi.org/10.1086/687336>
- Todd, P. n.d. "Strawson, Libertarianism, and 'External Justifications' of Our Responsibility Practices." Unpublished manuscript.
- Wallace, R. J. 1994. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- Watson, G. 1987. "Responsibility and the Limits of Evil: Variations on a Strawsonian Theme." In *Agency and Answerability: Selected Essays*, ed. G. Watson, 219–59. Oxford: Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199272273.003.0009>

