

Moral Responsibility for Concepts

Please cite the published version of this paper, which appears in the
European Journal of Philosophy

Rachel Fredericks (ORCID: 0000-0003-4622-0554)

Department of Philosophy & Religious Studies, Ball State University, Muncie, IN, USA

Email: rlfredericks@bsu.edu

Abstract: I argue that we are sometimes morally responsible for having and using (or not using) our concepts, despite the fact that we generally do not choose to have them or have full or direct voluntary control over how we use them. I do so by extending an argument of Angela Smith's; the same features that she says make us morally responsible for some of our attitudes also make us morally responsible for some of our concepts. Specifically, like attitudes, concepts can be: (a) conceptually and rationally connected to our evaluative judgments, (b) in principle subject to rational revision (reasons-responsive), and (c) the basis for actual and potential moral assessments of people that we have good reasons to endorse. Thus, we are *open* to moral appraisal on the basis of having and using (or not using) our concepts when they reflect our evaluative judgments, though even then it is not always appropriate to praise or blame us on that basis.

Keywords: attitudes, concepts, moral responsibility, Angela Smith

1. Introduction

Concepts are neither true nor false, but they can be evaluated: do we have reason to track the distinction drawn by the concept? Should we have this or that concept in our repertoire at all? If so, how ... should [we] construe it? What alternative concepts might we deploy instead? (Haslanger, 2014: 28)

I aim to argue that sometimes we are morally responsible for having and using (or not using) our concepts in the ways that we do. To do so, I extend one of Angela Smith's (2005) arguments for her rational relations view of moral responsibility, according to which we are morally responsible for both our actions and our attitudes insofar as and because they reflect our evaluative judgments (see also Smith 2004, 2007, 2008, 2012 and 2015). Like actions, she says, attitudes can be and often are appropriately connected to mental *activity* to ground moral assessment. This is not, yet, to say anything about what moral assessments are appropriate in any situation. On Smith's view, being morally responsible for something is being *open* to moral appraisal on the basis of it, so being morally responsible for something does not settle the question of whether one merits a positive, negative, or neutral appraisal because of it, much less whether anyone should express praise or blame on that basis.

Similarly, I argue that we are sometimes open to moral appraisal on the basis of our having and using (or not using) concepts in the ways that we do, insofar as and because our doing so reflects our evaluative judgments. The 'insofar as' is crucial. To the extent that having and using (or not using) a concept does *not* reflect one's evaluative judgments, one is *not* morally responsible for it. We are morally responsible for concepts when and only when (and because and to the extent that) they reflect our evaluative judgments. This leaves it open whether (and if so, when and to what extent) we are blameworthy or praiseworthy (or should be blamed or praised) for our concept possession and usage.

To support my conclusion, I argue that our concepts are sometimes conceptually and rationally connected to our evaluative judgments (§ 4.1), that concepts are in principle subject to rational revision (§ 4.2), and that concept possession and usage are the basis for some actual and potential moral assessments of people that we have good reasons to endorse (§ 4.3). First, however, I briefly explain and defend Smith's theory (§ 2) and clarify my assumptions about what concepts are (§ 3). After the main arguments, I discuss *which* concepts are more likely to reflect our evaluative judgments (§ 5) and why we can be morally responsible, not just for using concepts, but also for having them (§ 6). I conclude by responding to two objections (§§ 7–8).

Note that in what follows, I sometimes use 'moral responsibility for concepts' as shorthand for the more cumbersome 'moral responsibility for having and using (or not using) some of our concepts in the ways that we do'. This parallels how we might say 'responsibility for attitudes and actions' when we mean 'responsibility for *having* certain attitudes and *doing* certain actions'.

2. Moral Responsibility

My claim that we are morally responsible for concepts may strike many people as outrageous, since many people believe that (a) we cannot be morally responsible for things we do not directly, voluntarily control and (b) we (as individuals) cannot directly, voluntarily control which concepts we have or how we use them.¹

Regarding (b), I agree that which concepts one has, their exact content, and when and how one uses them are not, in general, under one's full, direct, or voluntary control. This is because (i) concepts are acquired in social contexts that no individual can fully control (see Fricker 2007), (ii) concepts are often acquired, revised, and used unintentionally,² and (iii) the exact content of one's concepts is often opaque even to oneself (which is why conceptual analysis requires effort).

To expand on the support for (b), consider an example. As a child, I was taught that there exists a god of a Christian sort. I did not seek out that concept, I could not

help but form it in my circumstances, and its content was not wholly up to me. While, for reasons not entirely under my direct, voluntary control, I no longer believe that deity exists, the concept is still in my repertoire. More importantly for present purposes, I sometimes find myself using it in ways that conflict with my considered judgments. For example, occasionally, upon realizing that I have wronged someone, the thought 'God is going to punish me' pops into my head. When aware of such thoughts, I think they involve mistaken concept usage, because my considered judgments about appropriate application of the concept GOD have changed. Yet I still sometimes use the concept in that way, since old habits die hard. So, (iv) concept usage is not always under our full, direct, voluntary control, as shown by cases when it conflicts with our considered judgments. Furthermore, insofar as we cannot believe (or cease believing) at will,³ our concept usage is no more under our full, direct, voluntary control when our judgments are correct and consistent than when they are not.⁴ Thus the case for (b) is strong.

As for (a), many people believe that only one's voluntary choice or exercise of control makes one morally responsible for something. That belief underlies many people's thinking about common practices, like saying that someone is not morally responsible for something because she did not mean to do it, did not know what he was doing, had no choice about it, or could not control it.⁵ Smith calls these

volitionalist views of moral responsibility, since they say that exercising the will is a precondition for being morally responsible.⁶

However, (a) is controversial. Some people argue that we *can* be morally responsible for things we do not directly, voluntarily control; of their views, I favor Smith's rational relations view.⁷ As she says, 'what makes an attitude "ours" in the sense relevant to questions of responsibility and moral assessment is not that we have voluntarily chosen it or that we have voluntary control over it, but that it reflects our own evaluative judgments' (2005: 238). One of her supporting arguments involves an investigation of cases designed to show not only that (a) it is common practice to assess people on the basis of attitudes that are not under their direct, conscious, or voluntary control, but also, more crucially, (b) we have good reasons to endorse those assessment practices, (c) a rationalist, rather than volitionalist, account of moral responsibility provides a more compelling rationale for those practices, and thus (d) what makes one morally responsible for an attitude is its reflecting one's evaluative judgments, not its being under one's direct or voluntary control.⁸

For instance, Smith says, we sometimes consider people (including ourselves) responsible, and maybe even blameworthy, for things like forgetting a friend's birthday, even though forgetting is not under people's direct or voluntary control (2005: 236). We do so insofar as and because we think (often for good reasons) that

forgetting reflects something problematic about the forgetful person's evaluative judgments (not just failure to take relevant measures to remember earlier); in such cases, forgetting often reflects judgments about the friend's value and importance. (However, insofar as forgetting does *not* reflect one's evaluative judgments, as, for instance, it might not if one has a serious memory disorder, one would not be morally responsible for it.) Smith concludes that we can appropriately be held responsible not just for our actions, but also our attitudes, since they can reflect our evaluative judgments (our *mental activity*) even when they are not under our direct or voluntary control. Moreover, our being morally responsible for our attitudes means that it is appropriate to ask us to defend or justify them, just as with our actions. A friend might reasonably ask me *why* I forgot her birthday, wanting more than mere explanation; she might want me to justify myself, which I might or might not be able to do.

Smith's view is not without challengers. Some have objected that it is problematically narrow, thin, or incomplete.⁹ David Shoemaker (2011) thinks there are three distinct senses of moral responsibility; he appeals to cases of irrationality, non-rational emotional commitment, and psychopathy to argue that Smith captures only a subset of all moral responsibility, leaving out the kinds found in those cases.

While more substantive responses to such an objection are available (see Smith 2012 and Hubbs 2013), for present purposes, note that if responsibility in Smith's

sense is narrower, thinner, or less comprehensive than what others mean by 'moral responsibility', it is *less* controversial to claim that we are morally responsible, in that sense, for concepts. I do not think that Smith's account is problematic in those ways, but regardless, it is worth showing why we are morally responsible for concepts, even if only in a restricted sense, since doing so highlights a neglected feature of the moral landscape and the argument has implications that are worth exploring.

Nevertheless, others might object to Smith's view because it seems to entail that we are morally responsible even for attitudes that are implanted in us by mad scientists, hypnotists, deities, or bizarre accidents, contra common intuitions. For if someone or something were to implant an attitude in me, I would seem not to be morally responsible for it; because of its causal origin, many would say it is not *mine* in the right sense.

However, even if we are happy to use such far-fetched cases to test theories of moral responsibility (and I am somewhat skeptical about that methodology), we can respond to this worry (see Smith 2005: 261-262 and 2008: 389). For if the implanted attitude is unconnected to my evaluative judgments (say, immediately after implantation of a single attitude), then I am not responsible for it. If it does reflect my evaluative judgments (say, after I have made evaluative judgments that integrate it with my other attitudes), I would be morally responsible for it. In cases of wholesale

attitude implantation, questions about continuity of personal identity become more pressing; I doubt whether *I* would be there, which precludes *my* being responsible.

While there is certainly more to say to defend Smith, moving forward I assume that she is at least largely correct about the criteria for moral responsibility. So, do the same features that make us morally responsible for some of our attitudes also make us morally responsible for some of our concepts? I think so. As Smith says, we are morally responsible for attitudes *insofar as* and because they are (a) conceptually and rationally connected to evaluative judgments, (b) in principle subject to rational revision, and (c) the basis for some actual and potential moral assessments of people that we sometimes have good reasons to endorse. I argue that the same is true of concepts.

3. Concepts

But what are concepts? Starting functionally, concepts are cognitive resources, and as Sally Haslanger says, frameworks of concepts are tools (2012: 23). Our conceptual repertoires allow us to mentally unite tokens and differentiate types, thereby enabling us to categorize and understand things.¹⁰

Whether concepts are mental representations (à la Fodor (1975)), abstract objects (à la Frege (see Zalta 2001)), or even abilities (à la Kenny (2010)), what matters most for my purposes is that concepts are things that we, in common parlance,

possess and use. My concern is with the mental activities that involve or result in persons standing in relations to concepts, specifically the relations denoted by 'possess a concept' and 'use a concept'. In fact, I see concept possession and usage as inextricably linked, which informs my thinking that we are morally responsible for both (a) possessing concepts and also (b) whether, when, and *how* we use (or do not use) them. I say more about why we are responsible for both possessing and using (or not using) concepts in § 6, but I follow Siebel (2004) in thinking that to have a concept is to be able to think about what it is a concept of (which is not to say that one's conception is adequate).

I do not think that ontological debates about exactly what concepts are need to be settled for us to say that the relations of possession and usage can link people to concepts in ways that render the former responsible for the latter. Just as when a person has a propositional attitude, they stand in a relation to a proposition and can thereby be responsible for doing so, when a person possesses or uses a concept, they stand in a relation to it and can thereby be responsible for that. Philosophers who argue that we are morally responsible for our propositional attitudes generally remain agnostic about the exact nature of propositions, and I would prefer to remain similarly agnostic about what a concept is.

However, I cannot be completely agnostic about that here, for two reasons: one having to do with the impossibility of using completely neutral language and one having to do with my focus on *individual* moral responsibility. Regarding the first, whatever verbs I choose to describe our relations to concepts will inevitably fit better with some views about concepts' ontological status than others, so complete agnosticism is not an option. For example, 'grasping concepts' might seem to imply acceptance of the abstract object view of concepts, whereas 'forming concepts' and 'revising concepts' might seem to imply acceptance of the mental representation view, and 'mastering concepts' might seem to imply acceptance of the ability view.

So, I advance my argument under the assumption that concepts are mental representations rather than abstract objects or abilities. I do so because, in addition to the substantial (though not necessarily overwhelming) independent reasons to favor the psychological view that concepts are representations (see Margolis and Laurence 2007 and Sutton 2004), doing so fits best with both (a) Smith's focus on *individual* moral responsibility and (b) common expressions that portray concepts as things that individuals can *possess and use*. I leave it for another occasion to argue for the same conclusion on the assumption that concepts are abstract objects; as I envision it, that involves focusing on collective moral responsibility. For now, I restrict myself to considering how concepts, understood as mental representations, are sometimes

similar enough to attitudes to reflect our evaluative judgments and thus to render us morally responsible for them.

4. Moral Responsibility for Concepts

4.1. *Conceptual and rational connections to evaluative judgments*

The first relevant similarity between concepts and attitudes is that they are both conceptually and rationally connected to evaluative judgments. As Smith writes:

Attitudes such as contempt, jealousy, and regret seem to be partially constituted by certain kinds of evaluative judgments or appraisals. ... There seems to be a conceptual connection between having these attitudes and making, or being disposed to make, certain kinds of judgments. (2005: 250)

For example, having the attitude of regret seems to be conceptually connected to being disposed to judge that something bad happened.

Turning from attitudes to concepts, the mere possession of a concept is similarly *conceptually* connected to making or being disposed to make certain judgments. For instance, to have the concept of red, one must be able to think about redness and disposed to judge that red is a color, is different from green, and so on; without such dispositions, one lacks RED. Similarly, having UNICORN just is (in part) being disposed to judge that a horse with one horn is (or would be) a unicorn.

Concepts come in sets, conceptual repertoires, which have complex internal structures of interconnections. For instance, TRIANGLE is necessarily related to THREE, ANGLE, and SHAPE, and ACTION is necessarily related to AGENT.¹¹ Other connections are contingent; for instance, in my conceptual repertoire, NECTARINE (which represents my favorite fruit) and DELICIOUS are connected, but might not have been.

One might notice that the judgments that I said are conceptually connected to the concepts RED and UNICORN are not *evaluative* judgments, and Smith's argument requires linking attitudes to specifically evaluative judgments. One might think that the attitudes that she mentions necessarily have normative content, whereas the concepts I mentioned do not, and thus that I have a problem.

However, I have a two-part response. First, fair enough: I chose RED and UNICORN because they are fairly simple and clear examples. However, those examples do not rule out the existence of other concepts, like JUSTICE, which have normative content and which are linked conceptually (and, as I explain shortly, rationally) to evaluative judgments. Furthermore, and more importantly, in part because concepts are formed in social contexts, their content is rarely, if ever, fully transparent.¹² Since the unrecognized baggage that our concepts carry is often normative, sometimes even familiar concepts that seem to lack normative content

are conceptually or rationally connected to evaluative judgments. For instance, many people think WHEELCHAIR BOUND is purely descriptive and innocuous, but since it falsely implies that all people who use wheelchairs lack freedom and agency (unlike, say, PERSON WHO USES A WHEELCHAIR or WHEELCHAIR USER), it is disrespectful and misleading.¹³

With that in mind, let us shift from *conceptual* connections to *rational* connections. As Smith says:

If we value something and judge it to be worth promoting, protecting, or honoring in some way, this should (rationally) have an influence on our unreflective patterns of thought and feeling. We commonly infer from these unreflective patterns, or from their absence, what a person really cares about and judges to be important. (2005: 247)

Similarly, our evaluative judgments about things' worth should (rationally) influence which concepts we do or do not use in thinking about those things, and the specific ways in which we use those concepts when we do. For patterns of thought are partially constituted by patterns of concept usage; if thought patterns should, rationally, be influenced by evaluative judgments, then concept usage patterns also should be. For example, if I judge that pacifism is virtuous, then I should, rationally, tend to use concepts like BRAVE and REASONABLE rather than COWARD or MISGUIDED when

thinking about pacifists. I could still be rational in thinking of a particular pacifist as being misguided or cowardly, or in thinking, 'Many others see pacifists as misguided cowards', but the point is that our evaluative judgments and what we see as reasons should (in a rational sense) impact which concepts we use, and how we use them.

Of course, we are not always fully rational, as Smith emphasizes:

The view that I am putting forward takes as its starting point the idea that some of our mental states are linked to particular judgments in such a way that, if one sincerely holds a particular evaluative judgment, then the mental state in question should (or should not) occur. The "should" in question here is the should of rationality and, therefore, marks a normative ideal which our actual attitudes may not always meet. (2005: 253)

I am suggesting that sometimes to have or use a concept is to be in a mental state that is, in key ways, similar to having an attitude for which we are morally responsible. For instance, INJUSTICE necessarily has normative content, and possessing it is conceptually connected to being disposed to judge that injustice is bad. Furthermore, there are rational connections between the concepts I have, the evaluative judgments I make, and the ways in which I do or do not use concepts. When I use some concept in my repertoire, it is because I have, consciously or not, made a judgment about when it is appropriate to use it. For example, there is a rational connection between my

possessing the concept INJUSTICE, my using it to categorize an act, and my evaluative judgment that I should not do that act. There is also a rational connection between my possessing the concept INJUSTICE, my evaluative judgment that another act is morally acceptable, and my *refusal* to apply INJUSTICE to that act. Moreover, if I judge that one of Antonin Scalia's legal opinions is a blight on American jurisprudence, then that judgment is rationally connected to my using INJUSTICE to categorize that opinion, and you can appropriately ask me to defend my using INJUSTICE thus.¹⁴ So, as I clarify in the next section, there are rational standards that govern concept usage; insofar as my concept usage reflects and responds to the reasons they provide, and so meets those standards, my concept usage is rational. Thus, like the attitudes for which we are morally responsible, some of our concepts are conceptually and rationally connected to our evaluative judgments.

4.2. *Rational revision*

But the similarities do not end there; concepts are also in principle subject to rational revision, as attitudes are. As Smith writes, 'In order for a creature to be responsible for an attitude, on the rational relations view, it must be the kind of state that is open, in principle, to revision or modification through that creature's own processes of rational reflection' (2005: 256).¹⁵ To see that attitudes are (somewhat) reasons-responsive,

consider regret again. If I feel regret, but then discover that no bad thing has actually occurred, my regret will likely dissolve. Alternatively, if I find that I was wrong about *which* bad thing occurred or how bad it was, the nature of my regret will likely be altered.

Illustrations of the fact that some of one's concepts are similarly (somewhat) reasons-responsive are numerous and varied.¹⁶ For example, over time, without necessarily being aware of doing so or choosing to do so, many Americans have, using rational activity, revised their concept of membership in the Democratic and/or Republican parties. For the strengths of, challenges faced by, and members of political parties change constantly, and people's concepts of said parties change in response: quickly or slowly, in ways that are obvious or recognizable only after careful reflection, if at all. So, some of our concepts, like those of specific political parties, change because reality changes, and because concepts are reasons-responsive.

In contrast, sometimes our conceptual repertoires change, not because external reality changes, but because we gain access to information that, rationally speaking, forces us to conceptualize the same old reality in a new way. For instance, my formative circumstances led me, like many others, to form a biological concept of race. Races, as represented by that concept, are products of natural forces like, say, different kinds of rocks or birds. However, that biological concept has been superseded, in my mind,

with a social concept, which represents races as categories devised by people to rank groups of people on the basis of observed and imagined bodily features. My exposure to relevant scientific and historical information and philosophical arguments like Haslanger's led to this adjustment in my conceptual repertoire, which happened gradually over time. So, my concept of race changed in response to being exposed to new reasons, as have the ways and contexts in which I use it.¹⁷ To use Haslanger's terminology (2012: 375, 378, and 387-390), both my manifest and operative concepts of race have changed.¹⁸ Nevertheless, the biological and social reality that I think about when using RACE have not changed to a degree commensurate with the change in my concept. Other examples of this kind of reasons-responsiveness involve simply adding or subtracting details from our concepts in light of new information, as when it dawned on me that triangles do not just have three sides (as I learned in preschool), but also have three angles (hence the name!). Triangles themselves did not change, but my concept TRIANGLE sure did (though not as much as RACE).¹⁹

Acquiring concepts is also a process that sometimes involves rational activity. Consider Miranda Fricker's (2007: chapter 7) discussion of how the concept of sexual harassment and its name were added to our collective resources by women trying to understand and deal with their own experiences of it.²⁰ They used their rational activity to discover that their experiences fit a pattern, named it, and thereby gave us a

powerful resource for combatting injustice. Thus, our rational activity enables us to acquire new concepts when our conceptual repertoires are inadequate for our purposes.

These examples show that we can revise the content of our concepts, change the contexts and ways in which we use concepts, and acquire new concepts through various more or less direct, more or less intentional, more or less voluntary means, and that such changes involve rational activity. Many philosophers who do conceptual analysis are committed to some version of the claim that concepts are in principle open to rational revision, insofar as they undertake conceptual analysis with the goal of clarifying concepts by shedding the intuitions about them and uses of them that we have reasons to reject. So being presented with new evidence, changing circumstances, or even changing priorities can (and sometimes should) initiate a reasons-responsive process of revising our conceptual repertoires. Since concepts are in principle subject to rational revision, proponents of the rational relations view should see the having and using of a concept, just like the having of an attitude, as something for which one can be responsible. This makes it appropriate to ask people to justify their having and using (or not using) their concepts.

To relate this to control and an earlier example, I cannot simply will myself to eliminate the concept of a Christian god from my repertoire, which, even if I could,

would be inadvisable given the society in which I live. Nor can I consistently use that concept in ways that cohere with my considered judgments by sheer act of will. Nevertheless, its connections to my other concepts, and the ways in which I do and do not use it, certainly reflect my evaluative judgments and my rational activity more generally. (For instance, the ways that I use GOD show that my judgments about appropriate uses of that concept are not wholly consistent.) So, I am morally responsible for that concept and can legitimately be asked to justify my having it and my usage of it, though saying this leaves it open whether I merit a morally positive, negative, or neutral assessment on that basis.²¹

4.3. Everyday practices of moral assessment

Smith's arguments are meant to show that people often make moral assessments of people on the basis of their attitudes. By recognizing this and, more crucially, arguing that we have good reasons to endorse that common practice, she defends the claim that we are morally responsible for our attitudes.

While I think that Smith is correct that people regularly assess themselves and others on the basis of their attitudes, this might seem to pose a challenge for me, since it seems vastly less common to morally assess people on the basis of their concepts. It is quite rare to even try to determine which concepts people possess, though quite

common to try to determine which attitudes people have. Why? There are serious obstacles to determining which concepts people (especially third parties) possess, but the same is true of attitudes. To explain why we rarely assess people on the basis of their concepts as opposed to their attitudes, the more salient points are that in everyday life, (a) people use folk psychology to explain and predict behavior, which involves appealing to beliefs, desires, and emotions, but usually not concepts, (b) most people simply do not think about concepts as such, though most do think about attitudes, and (c) when people do think about concepts, it seems common to assume (probably mistakenly) that most people have more or less the same core set of them, an assumption that seems less likely regarding people's attitudes.

Nevertheless, some of us do occasionally assess people for the concepts that they *use* (both for using them at all and, more importantly, for *when and how* they use them).²² Some of us make moral assessments of people for using concepts with derogatory content to categorize people on the basis of race, sex, nationality, ability, etc.²³ For example, many of us judge people to be callous, insensitive, cruel, culpably ignorant, or otherwise morally bad when they categorize people of Asian descent using ORIENTALS, trans women using SHE-MALES, impoverished people using TAKERS, and undocumented immigrants using ILLEGALS. But liberals are not alone in morally assessing people on the basis of their concept usage, as these examples might

imply. For instance, some conservatives judge people to be profane, offensive, culpably ignorant, or otherwise morally bad when they use MARRIED to categorize same sex couples or use FETUS but not PERSON to categorize unborn humans. Nor do people only make *negative* moral assessments of people's concept usage. For instance, one might consider someone praiseworthy for using PERSON WHO NEEDS HELP instead of TRAFFIC IMPEDIMENT to categorize a motorist stranded in a left-turn lane. To praise or blame in these cases requires an (implicit) attribution of moral responsibility.

Of course, people do not necessarily *know* when others possess particular concepts and use them in particular ways. However, sometimes when people think they know, they ask others to justify their doing so. More importantly, sometimes this is exactly what we should do. The point is not that many people regularly *do* make moral assessments of people on the basis of their concept possession and usage (though some do sometimes). The point is that this practice is sometimes *justified*. In various cases, people have good reasons to make such assessments, just as we often have good reasons to assess people on the basis of their attitudes and actions, whether or not we actually do.

Let's explore the justification for morally assessing people on the basis of their concept possession and usage by considering a more detailed example. I grew up in

Fairfield, Iowa, which is a small town with a university designed to promote the teachings of Maharishi Mahesh Yogi, founder of the Transcendental Meditation movement. Some so-called 'townies', who were unaffiliated with the university, were hostile toward those who were, and referred to members of the university community as 'roos' (or 'rus'), a derogatory term derived from 'guru', which means a person who is, at best, flaky and odd, and at worst, a dangerous, brainwashed moron. Knowing this, when I heard someone say something about 'roos', even when it seemed relatively innocuous and true (like 'Roos eat lots of vegetables'), I found myself making a negative assessment of whoever was using ROO to structure their thinking.

For I (and they) knew that ROO had more in common with FOOL than with RED. It originated in stereotypes that were morally problematic and, generally, epistemically unjustified as well. It was born of fear and dislike, if not hatred, which produced a (likely unconscious) felt need to divide people into 'us' and 'them' and denigrate 'them'. For many people, having disparaging attitudes was part and parcel of using ROO to categorize people, and vice versa (though I discuss a possible exception shortly). For certain attitudes can only be held if one has certain concepts, and the resources in one's conceptual repertoire impact the range of attitudes that one can (and is likely to) have.

Someone might say that the only things of moral significance in this example are derogatory attitudes and the actions that express them, and that I do not need to talk about concepts to capture what people are morally responsible (and potentially blameworthy) for in such cases. I can see why they might suggest that; it would certainly seem to make for a simpler and less revisionary view.

However, I do not think that it is *just* the attitudes and actions that form the basis of moral assessment in at least some such cases (nor do I think it should be). For I reject the underlying assumption that having an attitude is independent of having and using the concepts that figure in it. These are *interdependent* phenomena; having an attitude is a matter of using some of the concepts one possesses together in a certain way, and having a concept is being able or disposed to use it in thinking (in some circumstances) (see Davis 2005a and Siebel 2004).

Even if possessing and using concepts is not partially *constitutive* of having attitudes, possessing and using certain concepts is a *necessary precondition* for having certain attitudes. It is impossible to have attitudes if one lacks or does not use the concepts that figure in them.²⁴ So we would fail to account for some of the complexities of the moral landscape if we said that we are only morally responsible for our actions and attitudes, but not for the concepts that we possess and use in forming those attitudes. It would be like saying that one is morally responsible for shooting

someone without being responsible for having the gun that one used to do so.²⁵ For concepts are the necessary cognitive resources that make thinking about things possible, and without them, we would not think as we do.

But it is not merely using a concept that renders one morally responsible for it. Having a concept but refusing to use it can also reflect one's evaluative judgments (see McCullagh 2011). Specifically, it can reflect the fact that one, consciously or unconsciously, thinks there is good reason not to use it. So not using a concept that one has does not entail a lack of responsibility for it. It can still be attributable to the person; it can still be *theirs* in a way that renders it appropriate to ask them to justify their not using it. But the appropriateness of asking for such a defense entails nothing about whether they will be able to provide a good one, and thus entails nothing about whether they are blameworthy for that concept.

Furthermore, even if having and using concepts were fully separable from having the attitudes in which they figure, sometimes the concepts, not the attitudes, are the root of our problems. For sometimes, having and using a concept (in a certain way) enables problematic attitudes to play a more substantial role in one's patterns of thought, more strongly influence one's behavior, and be propagated to others.²⁶ For example, imagine a child who has heard expressions of derogatory attitudes about people who meditate, but lacks the concept ROO. When the child acquires ROO,

derogatory stereotypes about people who meditate become more salient; it becomes easier for that child to think derogatory thoughts about them, because ROO represents all TM practitioners as sharing some bad essence. It functions as cognitive shorthand for all that is 'wrong' with 'those people'. Moreover, when people start using ROO to mentally categorize people, that significantly increases the likelihood of their describing and treating people as 'roos' (that is, as inferior). That overt behavior then impacts the conceptual repertoires, attitudes, and behavior of others who witness it, including children who do not intend to disparage anyone, but who are too young to think critically about the significance of categorizing people as 'roos'.

More specifically, if one structures one's thoughts using ROO rather than PERSON WHO MEDITATES (or just PERSON), one is more likely to think, say, and do things that produce or reinforce the belief (in oneself and others) that the people in question are quintessentially *followers*, not people who exercise agency by choosing to meditate.²⁷ Simply put, using a concept that implicitly denies (or affirms) people's agency is bound to impact one's judgment, behavior, and the messages one communicates to others regarding them. In such cases, the possession and usage of the concept should be a, if not *the*, primary basis for assessment, at least when it comes to adults whose concept possession and usage reflect their evaluative judgments. So, to be clear, my negative assessments of some people who use the

term 'roo' are partly based on their possession and use of ROO, not *just* their attitudes, speech, or other behavior.²⁸

My focus on the ROO example is not meant to imply that only derogatory concepts applied to moral agents are relevant to moral responsibility and assessment. For example, environmentalists sometimes challenge people who, for Biblical reasons, use DOMINION to think about humans' relation to non-human natural entities. Similarly, one might criticize the use of CURRY to lump together all spiced, sauced dishes from the Indian subcontinent, given that the practice originated with British colonialists, not within the cultures that created said dishes, where people know them by more specific names and types. Even if we do not go so far as to criticize in such cases, we can appropriately ask people to justify their using those concepts to mediate their engagement with the world and consider how using them in those ways can be harmful (by, for instance, contributing to oppression, cultural appropriation or erasure, silencing, marginalization, etc.). The justification we would want in some cases would not be purely epistemic, but moral as well (or instead). Moreover, if (for good reasons) we *refrain* from using CURRY or DOMINION in those ways, we can still be morally responsible for having and *not* using those concepts, because that refraining can reflect our evaluative judgments; they just happen to be judgments that make us more likely to be praiseworthy than blameworthy.

To summarize, I have argued that concepts can share three key features with attitudes. First, they can be conceptually and rationally connected to evaluative judgments. Second, they are in principle subject to rational revision. Third, people sometimes make moral assessments of people on the basis of the concepts they have and use (or do not use), and more importantly, upon reflection, we have good reasons to endorse that practice. Since those three features are those in virtue of which we can be morally responsible for something, the same features that can make us morally responsible for our actions and attitudes can also make us morally responsible for our concepts. That responsibility encompasses being responsible for having or not having concepts,²⁹ and for using or not using them in specific ways. However, to claim that someone is morally responsible for their concepts is not yet to make any claim about whether they are blameworthy, praiseworthy, or neither on that basis, nor about whether or how anyone should blame or praise them.

My view both is and is not a significant departure from existing rationalist views about moral responsibility. On the one hand, to any rationalist who already accepts that we can be responsible for mental and physical activities, my conclusion is not so outrageous after all. Using a concept is one of the mental activities characteristic of a rational agent, and having a concept, if not itself a mental activity, at least presupposes mental activity.³⁰ So my view is true to the widely shared belief that moral

responsibility depends on *activity*, which rationalists and volitionalists agree on (despite disagreeing about the *kind* of activity that matters or which mental states we are active with respect to).³¹

On the other hand, my view adds a whole new category of mental activity for which we can be morally responsible. You might think that rationalists gain nothing from this, since (a) they already endorse moral responsibility for attitudes and (b) having and using concepts is nothing over and above having attitudes. However, by saying that we can also be morally responsible for the mental activities involved in having and using concepts, one (a) accepts a previously unrecognized implication of an existing rationalist view, and (b) opens the possibility of more fine-grained moral assessment. Recognizing our moral responsibility for concepts *and* attitudes allows us to target our assessment (and praise or blame) to the most crucial aspects of our agential activity in varying cases. Deciding whether or not someone is morally responsible in such cases is the first step toward deciding what, if any, praise or blame they merit.

5. Which Concepts Are Most Likely to Reflect Our Evaluative Judgments?

I have argued that we are *sometimes* morally responsible (though not necessarily praiseworthy or blameworthy) for having and using (or not using) *some* of our concepts: specifically, that we are morally responsible for our concepts when they reflect our

evaluative judgments. But *which* concepts are most likely to do so? Those seem to merit special moral scrutiny. Due to space limitations, I can only start developing an answer and mention some difficulties that must be faced in improving on it.

There are at least four options for how to approach the question. We might say that *which* concepts are most likely to reflect our evaluative judgments depends on (a) how the concepts were acquired, (b) how they are used, (c) their content or semantic features, or (d) their structure or syntactic features.

My first impulse was to look for a feature common to the examples in § 4, which initially led me to favor approach (c) and say that moral and social concepts (concepts of social groups, practices, etc.) are most likely to reflect our evaluative judgments. For those seemed most likely to meet the three conditions discussed in sections § 4.1-3. Mathematical, logical, and natural kind concepts seemed least likely to fulfill those conditions and thus least likely to reflect our evaluative judgments.

However, that picture is significantly complicated by at least three difficulties. First, we sometimes mistake concepts with moral and/or social content for natural kind concepts. The previously mentioned RACE example is one such case, but philosophers have uncovered many examples of concepts of supposedly natural, essential, and unchanging things are actually concepts of things that are contingent, changing, and socially constructed (in one sense or another). I expect us to find more. So we need

to be humble about our ability to correctly categorize concepts, which makes me cautious about inferring that being morally responsible for concepts is rare.

Second, concepts are used in sets, making it difficult to isolate which evaluative judgments are reflected by the usage of each concept individually, as opposed to the evaluative judgments reflected by the usage of that set of concepts together. For example, AND and THUS do not seem to be moral or social concepts. However, imagine someone who endorses the morally egregious thought that 'x was dressed like a slut *and thus* x deserved to be sexually assaulted'. We have two options – we can say that this thinker is morally responsible for possessing and using all the concepts involved in having that thought, or that they are only morally responsible for some of them. I would prefer to say that they are likely to be responsible for them all, because if the situation is even fairly normal, then in using each of those concepts in that specific way, the subject would be guided (though probably unconsciously) by their evaluative judgments, so the concepts involved thereby become attributable to them.

Reflecting on that case, one might be tempted to say that we are only morally responsible for concept possession and usage insofar as and because they reflect judgments that make morally significant evaluations. But that leads to another complicating factor.

For, third, determining which evaluations are morally significant is a complex and contentious project. To the extent that standards of rationality, prudence, and morality are enmeshed (or identical), this suggests that many of our concepts reflect evaluative judgments. For I can imagine rational and prudential reasons being brought to bear in evaluating nearly every (if not every) instance of concept possession and usage, and if norms of rationality, prudence, and morality are deeply enmeshed, that supports my intuition that many concepts are likely to reflect our evaluative judgments.

That said, my arguments do leave open the possibility of some concept possession and usage that does not reflect our evaluative judgments. For example, pathologically recalcitrant attitudes might involve concept usage that does not reflect a subject's evaluative judgments. Similarly, science fiction examples of attitude implantation might involve subjects who have and use concepts that do not reflect their evaluative judgments. Those cases might make us think that the better approach for determining the likelihood that concepts reflect our evaluative judgments is to focus on how the concepts were acquired (option (a) above).

That approach might also seem to be supported by cases in which someone intentionally seeks to acquire a new concept. Those would seem to provide paradigmatic cases of subjects being morally responsible for possessing concepts, since subjects' intentionally seeking to acquire new concepts seems to reflect

evaluative judgments about the inadequacy of their existing conceptual repertoires for their purposes, if nothing else. But depending on the details, even mental state implantation cases and cases involving intentional concept acquisition might just as well be described as distinctive because of something about the patterned way in which such concepts are used after being acquired. So it is not obvious how best to justify the intuition that implantation cases do not reflect evaluative judgments but intentional acquisition cases do, and thus that we are not responsible for the former, but we are for the latter.

Given those challenges (and others), I hesitate to draw any general conclusions about which concepts are most likely to reflect evaluative judgments, under what conditions they would be most likely to do so, or what proportion of concepts do so. Besides, those are (at least largely) empirical questions. To the extent that they are not, answering them requires accounting for how failures of rationality are possible, which is quite challenging (see Smith 2005: 255). There is certainly room for people to agree that we can be morally responsible for our concepts, insofar as they reflect our evaluative judgments, without (yet) agreeing about exactly which ones (are likely to) meet the relevant criteria.

6. Why We Are Morally Responsible for Having *and* Using Concepts

Readers may have noted that in the main arguments of § 4, I said more about *using* concepts than *having* them. One might challenge me by suggesting that my arguments only show that we can be morally responsible for the former, not the latter.

As I said in § 3, I think that having and using a concept are not fully separable phenomena. One's having a concept is necessarily connected to (if not constituted by) a pattern of using (and not using) that concept (including the ability or disposition to do so under counterfactual conditions). This gives us reason to think that if we can be morally responsible for one, we can be for both. If I became convinced that I was wrong about the metaphysical inseparability of having a concept and using it in a patterned way, then I would probably concede that we are morally responsible only for using (or not using) concepts, not for merely having them. For I think that being morally responsible for a concept requires a rational connection between the concept and one's evaluative judgments, and if *per impossibile* one has a concept yet neither uses nor does not use it, nor has an ability or disposition to do so under any circumstances, then that concept would not be rationally connected to one's evaluative judgments, and one would not be morally responsible (that is, in principle open to moral assessment) for it.

I also gave reasons in § 4.3 for thinking that one can be morally responsible for having but *not* using a concept. The idea is that instances of non-usage are part of

one's pattern of usage and can reflect one's evaluative judgments. However, there are additional reasons to say that we can be morally responsible for both having *and* using (or not using) concepts. One has to do with ongoing metaphysical debates about concept individuation and how to determine when a person (a) has a concept that is somehow different from what it should be (a problematic concept) as opposed to when (b) they have the concept they should, but apply it in a problematic way. Similarly, there is debate about how to distinguish between (a') two people having different concepts and (b') their having the same concept, but using it differently.³²

If we can be morally responsible for having *and* using concepts, then we can say that when (a) is the correct diagnosis for a case, the subject is potentially worthy of criticism for *having* the concept. Whether they are potentially worthy of *moral* criticism will depend on the sense in which the concept is problematic. However, when (b) is the correct diagnosis, they are potentially worthy of criticism for *using* the concept. Of course, they might be worthy of criticism for both, if both types of errors obtain. However, we need to establish subjects' morally responsibility *before* we entertain the separate question of whether they are blameworthy for those errors.

One might protest that it is unfair to say that one is morally responsible for merely having a concept if one could not help forming it and cannot stop having it. For example, one could not have been expected to know, upon picking up this paper,

that one would encounter ROO. But upon reading the paper, one probably cannot stop oneself from acquiring and continuing to possess that concept. One might find something unfair about saying that one is morally responsible for having ROO under such conditions.

My response has two parts. First, this seems to confuse being morally responsible with being blameworthy. I do not think that questions of fairness are relevant to questions of whether someone is morally responsible for something. Saying that x is morally responsible for y is, in my view, saying that y is attributable to x in the way necessary for it to make sense to ask x a justificatory question about their relation to y. Saying that y is attributable to x is not fair or unfair, whether y is a concept, attitude, action, or body part. However, questions of fairness are often relevant to blameworthiness and whether people should express blame. Fairness is salient to whether the attribution of y to x is grounds for blame; it can be unfair, arrogant, and/or counterproductive to blame people for what they cannot control. Second, and maybe more importantly, the unfairness concern seems to presuppose a volitionalist account of moral responsibility, which I reject. Since I think we are morally responsible for whatever reflects our evaluative judgments, not for whatever is a direct, voluntary product of our wills, this unfairness objection begs the question against my view.

7. Objection: Eliminating Conceptual Resources

One might object if one thinks that my view entails that we ought to remove from our repertoires all concepts that have derogatory content or that are necessarily connected to derogatory attitudes, even if that requires more than just an act of will. That could be seen as problematic for two different reasons.

First, so many of our concepts are somehow connected to derogatory attitudes, and so many of those connections are somewhat opaque to us, that to eliminate them would take lots of work. Not only would we need to do the hard work of identifying and revising many concepts, but we would probably also need to develop many new concepts to compensate for those foregone resources. Committing to such wholesale conceptual overhaul is unappealing, given the workload involved.

But it is an open question whether requiring substantial revisions to our conceptual repertoires is a problem. I think this makes my view attractive, for I tend to think that we have lots of room to improve in every domain of life. But to justify any answer to that question requires arguing for an independent principle that states how much conceptual revision it is appropriate to require of people, and I know of no such argument.

Second, and more importantly, one might worry about my view because activists need to possess derogatory (and otherwise problematic) concepts, refer to them,

and sometimes use them in careful, strategic ways to combat injustice. For as Fricker shows, sometimes the lack of a concept, like SEXUAL HARASSMENT, constitutes and perpetuates injustice. To reduce or eliminate an unjust practice like sexual harassment, we need that concept *and* the concepts used by harassers to describe and categorize people as targets for harassment. So we probably need concepts with derogatory content, like SLUT and BITCH, for that purpose.

Regarding that second point, I agree that we need *lots* of concepts to reduce and, I hope, eventually eliminate unjust practices. Having ROO certainly helped me begin to understand group-based bias and animosity, and begin developing strategies for resisting them. But I do not claim that we should eliminate all concepts that have derogatory content or are connected to derogatory attitudes. Neither the fact that we can be morally responsible for our concepts, nor the fact that we can be morally blameworthy for having and using (or not using) morally problematic concepts entails that we should eliminate rather than revise any given problematic concept.

In fact, my view entails that we have reason to keep *revised versions* of derogatory concepts in our repertoires as long as related injustices and other problems exist (or are likely to) and so need to be resisted (or prevented). There is no inconsistency in saying this, for, like Smith, I want to keep the claim that a person is morally responsible for an action, an attitude, or (in my case) a concept separate

from claims about whether that person is blameworthy or praiseworthy for it (or neither). As Smith says:

[T]he rational relations view is an account of the conditions of responsibility in the sense of moral attributability, that is, the conditions under which something can be attributed to a person in the way that is required in order for it to be a basis for moral appraisal of that person. Merely claiming that a person is responsible for something, therefore, does not by itself settle the question of what appraisal, if any, should be made of the person on the basis of it. (2005: 266)

This is crucial. For determining that someone is morally responsible for a concept is just the first step toward determining whether they thereby merit praise, blame, or neither. When morally assessing someone for one of their concepts, we must respond to numerous factors specific to the person and context to determine whether having and using (or not using) that concept reflects on them positively, negatively, or neither (see Smith 2007: § 4).

I am inclined to think that nearly everyone who *possesses* a certain concept is morally responsible for that, but that some are morally blameworthy and others praiseworthy or neutral for how they *use* (or do not use) that same concept. So adults who use ROO to distance themselves from people who meditate, thereby

conceptualizing such people as inferior, are, I submit, both morally responsible and probably somewhat blameworthy for doing so. Similarly, upon acquiring ROO, I became morally responsible for having it. However, since I always (a) understood that the concept was created and used to unfairly denigrate all members of a group, (b) avoided using it to structure my thinking, (c) refused to use the term 'roo' to refer to people, and (d) tried to articulate disapproval of those who did, I do not think that my possessing ROO made me blameworthy. However, I probably did not merit praise; my possessing ROO was probably morally neutral, though whether that is so depends on various factors, some of which I am poorly positioned to assess. Anyone who possessed it but did not make any relevant evaluative judgments (if that is possible) would not have been responsible for possessing or using it, and thus certainly not blameworthy.

Furthermore, our needing a concept for present purposes implies nothing about whether we always will. I needed ROO as a child to understand and cope with the local social landscape, but things change. I certainly did not need it in the same way once I moved away, and I have heard that the social division in my hometown has diminished. It may be that no one uses that concept anymore (though I have heard that some people who meditate have reclaimed 'roo' and now self-identify as roos). Maybe no one needs that concept to understand and struggle against the idiosyncratic social

division that I grew up with, because maybe the problem has been resolved, changed, or will be soon. If so, I would be delighted.

8. Objection: Disanalogy between Attitudes and Concepts

Someone might object to my claim that concepts are reasons-responsive by noting one way that attitudes and concepts differ. One might say that attitudes are reasons-responsive if and only if they can and should disappear once we alter the judgments on which they depend. So, for instance, my hope for victory can and should be eliminated once I decide that defeat is inevitable, because defeat's inevitability is inconsistent with victory's possibility, and hope requires believing in the latter.

However, one might say, concepts are not like that: despite someone's coming to believe that ROO is not appropriately applied to anyone, that person likely will, and maybe should, retain that concept. Since concepts do not and should not necessarily disappear, unlike the hope described above, one might say that concepts are not linked to judgments as attitudes are, and thus are not reasons-responsive.

However, this objection relies on an overly narrow account of reasons-responsiveness. Utterly eliminating an attitude from one's mind is only one way to respond to reasons; there are many other ways that involve changing, but not eliminating, mental states. For instance, hope is reasons-responsive even if, in light

of a judgment about defeat's inevitability, one does not eliminate it, but merely reduces its intensity or changes its object (say, from winning to doing one's best). While such a person might have conflicting attitudes, and thus not be *perfectly* reasons-responsive, that does not mean that nothing reasons-responsive is going on.

Furthermore, attitudes and concepts can both be reasons-responsive without being reasons-responsive in exactly the same way. For one thing, they are structurally different; concepts are used in having attitudes, so in one sense, attitudes are structurally more complex than concepts. Many attitudes, for instance, are propositional and purport to describe the world, so they are responsive to reasons for thinking that they do so *correctly* (or not). The possibility of attitudes describing things incorrectly is why their elimination is often the response one has most reason to make. But isolated concepts are not propositional and thus do not purport to describe the world (though they can be used as tools in attempts to do so). Therefore, the kinds of reasons that concepts respond to are somewhat different from those that attitudes respond to. Hence, we use different criteria to make judgments about them and tend to describe attitudes as correct or fitting, but concepts as apt or useful. So, I can accept and, to some extent, explain the proposed disanalogy between concepts and attitudes while maintaining that both are (somewhat, somehow) reasons-responsive.

9. Conclusion

Assuming that concepts are mental representations, and building on Angela Smith's rational relations account of the criteria for moral responsibility, I have argued that because concepts are like attitudes in three key respects relating to rational activity, we are sometimes morally responsible for having and using (or not using) our concepts in the ways that we do, despite the fact that such things are not under our full or direct voluntary control. For having and using (or having but not using) concepts can involve taking a stand about how we ought to think about things and what matters.³³ Since it makes sense to ask us to defend such a stand, we should be prepared to do so. Concept possession and usage (or non-usage) invite such justificatory questions in ways that, for instance, kidney possession and usage do not, because concepts can reflect evaluative judgments and kidneys cannot.

If I am correct, there is plenty more to discuss regarding our moral responsibility for concepts; I particularly look forward to thinking more about *collective* moral responsibility for concepts, given the social factors that influence concept acquisition.³⁴

References

Arpaly, N. (2006). *Merit, Meaning, and Human Bondage: an Essay on Free Will*.

Princeton, NJ: Princeton University Press.

- Barnes, E. (2016). *The Minority Body: A Theory of Disability*. Oxford: Oxford University Press.
- Camp, E. (2013). Slurring Perspectives. *Analytic Philosophy*, 54, 330–349.
- Davis, W. (2005a). Concept Individuation, Possession Conditions, and Propositional Attitudes. *Noûs*, 39, 140–166.
- Davis, W. (2005b). Concepts and Epistemic Individuation. *Philosophy and Phenomenological Research*, 70, 290–325.
- Del Pinal, G. and Spaulding, S. (2018). Conceptual Centrality and Implicit Bias. *Mind & Language*, 33, 95–111.
- Dotson, K. (2012). A Cautionary Tale: On Limiting Epistemic Oppression. *Frontiers*, 33, 24–47.
- Duhau Girola, L. (2012). The Myth of Concept Publicity. *Ideas Y Valores*, 61, 101–113.
- Fodor, J. (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Fricke, M. (2007). *Epistemic Injustice: Power & the Ethics of Knowing*. New York: Oxford University Press.
- Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. New York: Oxford University Press.
- (2014). Social Meaning and Philosophical Method. *Proceedings & Addresses of the American Philosophical Association* 88, 16–37.

- Hieronymi, P. (2008). Responsibility for Believing. *Synthese*, 161, 357–373.
- Hubbs, G. (2013). Answerability without Answers. *Journal of Ethics & Social Philosophy*, 7, 1–15.
- Isaacs, T. (2011). *Moral Responsibility in Collective Contexts*. New York: Oxford University Press.
- Kane, R. (1998). *The Significance of Free Will*. New York: Oxford University Press.
- Kenny, A. (2010). Concepts, Brains, and Behaviour. *Grazer Philosophische Studien*, 81, 105–113.
- Levy, N. (2005). The Good, the Bad and the Blameworthy. *Journal of Ethics & Social Philosophy*, 1, 2-16.
- Margolis, E. and Laurence, S. (2007). The Ontology of Concepts – Abstract Objects or Mental Representations? *Noûs*, 41, 561–593.
- McCullagh, M. (2011). How to Use a Concept You Reject. *The Philosophical Quarterly* (1950–), 61, 293–319.
- McKenna, M. (2008). Putting the Lie on the Control Condition. *Philosophical Studies*, 139, 29–37.
- Mills, C. (1997). *The Racial Contract*. Ithaca and London: Cornell University Press.
- Nordby, H. (2004). Incorrect Understanding and Concept Possession. *Philosophical Explorations*, 7, 55–70.

- Scanlon, T.M. (1998). *What We Owe to Each Other*. Cambridge, MA: Belknap Press of Harvard University Press.
- Shoemaker, D. (2011). Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility. *Ethics*, 121, 602–632.
- Sidgwick, H. (1981). *The Methods of Ethics*. Indianapolis: Hackett.
- Siebel, M. (2004). A Puzzle about Concept Possession. *Grazer Philosophische Studien*, 68, 1–22.
- Smith, A. (2004). Conflicting Attitudes, Moral Agency, and Conceptions of the Self. *Philosophical Topics*, 32, 331–352.
- (2005). Responsibility for Attitudes: Activity and Passivity in Mental Life. *Ethics*, 115, 236–271.
- (2007). On Being Responsible and Holding Responsible. *The Journal of Ethics*, 11, 465–484.
- (2008). Control, Responsibility, and Moral Assessment. *Philosophical Studies*, 138, 367–392.
- (2012). Attributability, Answerability, and Accountability: In Defense of a Unified Account. *Ethics*, 122, 575–589.
- (2015). Attitudes, Tracing, and Control. *Journal of Applied Philosophy*, 32, 115–132.

Sutton, J. (2004). Are Concepts Mental Representations or Abstracta? *Philosophy and Phenomenological Research*, 68, 89–108.

Taylor, R. (1970). *Good and Evil*. New York: Macmillan.

Thagard, P. (1990). Concepts and Conceptual Change. *Synthese*, 82, 255–274.

Verdejo, V. and de Donato Rodríguez, X. (2015). Partial Understanding and Concept Possession: A Dilemma. *Ratio*, 28, 153–162.

Wallace, R.J. (1996). *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.

Wolf, S. (1990). *Freedom within Reason*. Oxford: Oxford University Press.

Zalta, E. (2001). Fregean Senses, Modes of Presentation, and Concepts. *Philosophical Perspectives*, 15, 335–359.

¹ I am concerned with *moral* responsibility throughout; I do not argue that we are *epistemically* responsible for our concepts, though I suspect we are, given that we are epistemically responsible for our beliefs and concepts are used in having beliefs.

² By ‘revising a concept’, I intend to encompass adjustments to its content and to the ways in which we use it (or do not).

³ See Hieronymi (2008).

⁴ Perhaps concept usage is only under our direct control when we choose to entertain a thought.

⁵ I think that in many such cases, people lack *blameworthiness*, not moral responsibility.

⁶ See Isaacs (2011: 79); Kane (1998); Levy (2005: 2–16); Sidgwick (1981: 60–61); Taylor (1970: 241–252); and Wallace (1996: 131–132) for volitionalist views.

⁷ For rationalist views, see Scanlon (1998) and Arpaly (2006). I appreciate how such views make epistemic and moral responsibility more clearly continuous; Hieronymi (2008) highlights this particularly well.

⁸ Two other benefits of Smith's view are its ability to (i) account for the moral self's boundaries and (ii) explain why intuitions that seem to support volitionalism are common. See Smith (2005: 263).

⁹ On the supposed shallowness of real self views like Smith's, see Wolf (1990); McKenna (2008); and Smith (2008).

¹⁰ See Thagard (1990) for more roles that concepts play in our thinking.

¹¹ One might ask whether saying that TRIANGLE is conceptually connected to ANGLE entails that one must possess the latter to possess the former. For it seems that children might identify, intentionally draw, and talk about triangles without having ANGLE. That seems right, but those children would not have *the same conception* of a triangle as I do (see Davis 2005a and 2005b on the difference between concept

possession and mastery, and the difference between concepts and conceptions). We are both morally responsible for our concepts, but her youth makes her more likely to have an excuse if hers is problematic. For present purposes, I need not be able to identify when a specific person possesses a specific concept, since I only want to show that when someone does have and use a concept, it can be attributable to her in a way that makes her open to moral appraisal on that basis. I lack space to explain or adjudicate debates about how to understand possession conditions of concepts, how to differentiate concepts, or what count as mistaken applications of concepts, but see also Nordby (2004), Siebel (2011), and Verdejo and de Donato Rodríguez (2015).

¹² See Haslanger (2012: 375, 378, and 387–390) on the distinction between manifest and operative concepts.

¹³ For relevant discussion of terms that are and are not (implicitly or explicitly) normatively laden, see Barnes (2016: 173–176).

¹⁴ You could ask me to defend my judging that the opinion is unjust, but that reframes the same request, since that judging *simply is* a particular way of using INJUSTICE.

¹⁵ I cannot capture all the subtleties of Smith's view here, nor the full significance of 'in principle' in this quote. For Smith (2004) recognizes that attitudes can be recalcitrant, especially when we hold conflicting attitudes.

¹⁶ The following discussion of rational revision raises questions, which I cannot answer here, about *what* is being revised. Some say it is concepts, but others say conceptions. On the distinction, see Siebel (2011). Some say conceptual content is being revised, but others say it is conceptual structure or functional or inferential role.

¹⁷ Have I revised my concept of RACE or chosen between two different concepts that are available for my use? I leave that question unresolved, since both are reasons-responsive processes. I also set aside a ship of Theseus-like problem about conceptual identity and change.

¹⁸ This rarely prevents me from understanding people using a biological concept of race, since I can remember or discover changes in my conceptual repertoire in hindsight, which enables such understanding. Revising my concept does, however, rationally commit me to judging (when appropriately prompted) that there is something problematic about my 'earlier draft' of the concept. On partial understanding and concept possession, see Nordby (2004) and Verdejo and de Donato Rodríguez (2015).

¹⁹ I do not think there is a bright line between those two kinds of conceptual change. For example, one's concept of marriage might have changed *both* because of changing (legal) reality and *also* because of new ways of understanding unchanged aspects of reality (like love, commitment, and justice).

²⁰ See also Barnes (2016: 44) on the social construction of DISABILITY, which implies a similar process.

²¹ Moreover, if ought implies can, then insofar as we are generally unable to eliminate concepts from our repertoires, continuing to possess even morally problematic concepts can be morally permissible.

²² See this paper's epigraph, where Haslanger asks whether we should retain our concepts, allowing for a moral interpretation of 'should'.

²³ Some might object to my treating slurs as expressions of distinct concepts, preferring to treat them as expressions of non-slurring concepts coupled with derogatory attitudes like disgust or contempt, and locating the problem in the attitudes. I do think of slurring concepts as distinct, with problematic content built in (see Camp 2013), but even if I am wrong to do so, elsewhere I aim to show that the kinds of criticisms that I level against slurring concepts also apply to some commonly used non-slurring concepts that *appear* to be purely descriptive and morally unproblematic.

²⁴ This raises questions, which I briefly discuss in § 7, about the possibility and desirability of rehabilitating or reclaiming (rather than eliminating) concepts like QUEER, BITCH, etc. (and associated terms) that have been linked to morally problematic attitudes and practices.

²⁵ This might make sense in some unusual cases, just as we might not be responsible for concepts in some unusual cases. However, shooters will virtually always be at least partially responsible for having the guns they use, even if others share that responsibility.

²⁶ A similar idea apparently animated Charles Mills when he wrote:

concepts are crucial to cognition: cognitive scientists point out that they help us to categorize, learn, remember, infer, explain, problem-solve, generalize, analogize. Correspondingly, the *lack* of appropriate concepts can hinder learning, interfere with memory, block inferences, obstruct explanation, and perpetuate problems. I am suggesting, then, that as a central concept the notion of a Racial Contract might be more revealing of the real character of the world we are living in, and the corresponding historical deficiencies of its normative theories and practices, than the raceless notions currently dominant in political theory. (1997: 6-7)

Relatedly, for discussion of an overlooked type of biases encoded in the dependency networks of our conceptual representations, see Del Pinal and Spaulding (2018). See also Camp (2013).

²⁷ Those two concepts may or may not be coextensive, but they certainly represent the items in their extension differently.

²⁸ Of course, in practice, it will often be hard to tell whether the (main) moral problem is the concept(s), attitude(s), behavior(s), or some precise combination of those.

²⁹ On the epistemic and moral significance of *lacking* concepts and disagreement about the primary harm of such a lack, see Fricker (2007) and Dotson (2012). If I am correct, we have a moral responsibility not only to revise or eliminate problematic concepts, but also to *acquire* liberatory and otherwise worthwhile concepts.

³⁰ Here 'rational agent' means one capable of rational activity, not one who is always perfectly rational.

³¹ Volitionalists could support my conclusion by arguing that having and using concepts involves choice or another exercise of will.

³² See, for example, Duhau Girola (2012).

³³ See the Brandom quote in footnote 18 of McCullagh (2011).

³⁴ Many thanks to Ben Almassi, Zac Cogley, Jeremy Fischer, Angela Smith, my colleagues at Ball State, anonymous reviewers at multiple journals, and audiences at WOGAP, the Rocky Mountain Ethics Congress, and the APA Eastern for insightful comments.