



ANCHORING A REVISIONIST ACCOUNT OF MORAL RESPONSIBILITY

BY KELLY ANNE MCCORMICK

JOURNAL OF ETHICS & SOCIAL PHILOSOPHY

VOL. 7, No. 3 | DECEMBER 2013

URL: WWW.JESP.ORG

COPYRIGHT © KELLY ANNE MCCORMICK 2013

Anchoring a Revisionist Account of Moral Responsibility

Kelly Anne McCormick

REVISIONISM ABOUT MORAL RESPONSIBILITY is the view that we would do well to distinguish between what we think about moral responsibility and what we ought to think about it, that the former is in some important sense implausible and conflicts with the latter, and so we should revise our concept accordingly. A successful revisionist theory of responsibility must thereby meet at least three challenges. First, it must meet the *diagnostic challenge* of identifying our folk concept and providing good reason to think that some important features of this concept are implausible.¹ Second, it must meet the *motivational challenge* of explaining why, in light of this implausibility, our folk concept ought to be revised rather than eliminated. Finally, it must address the *prescriptive challenge* of providing an account of how, specifically, we ought to revise our thinking about moral responsibility.² If successful, revisionism stands to recast the landscape of the philosophical debate about free will and moral responsibility in a significant way. In this paper, I focus on the potential for revisionism to meet the prescriptive challenge. I take this challenge to raise the most serious worries for revisionism, while a successful response to it lies at the heart of the overall revisionist project.

Here I will distinguish between two related problems for revisionism as it has been recently developed and defended by Manuel Vargas (2005, 2009, 2011, 2013).³ I call these problems the *reference-anchoring* and *normativity-anchoring* problems. Michael McKenna (2009) and Derk Pereboom (2009) have each raised versions of the former, and I will argue that revisionists can successfully respond to their objections. However, it is not clear that revisionists can provide a tenable response to the normativity-anchoring problem, and focus on the reference-anchoring problem has thus far allowed this more serious worry for revisionism to go largely unnoticed.

In section 1, I provide some background on revisionism, in particular the best available account of the methodology for revisionist prescriptive-

¹ It is important to stress that the revisionist need not insist that our folk concept is univocal, nor that it entails that it is *impossible* for human beings like us to be responsible. For further discussion of why revisionists need not insist that our folk concept is univocal, see Vargas (2013). For a defense of the claim that free will and responsibility are impossible, see Strawson (1993) and Hurley (2000) for a series of objections.

² There is of course some ambiguity about what “our thinking” should be taken to refer to here. Is it our shared concept of moral responsibility, or the wide range of beliefs that philosophers and ordinary folks alike hold about moral responsibility? What revisionists are interested in is the former, but it is notoriously difficult to distinguish between concepts and sets of beliefs, and where to draw this line will depend at least in part on one’s preferred view of concept individuation. I will discuss these issues in further detail in section 2.2.

³ Hereafter I will use “revisionism” to refer to Vargas’ brand of revisionism in particular, as he thus far stands alone in his attempts to systematically develop and defend the view.

theory construction (Vargas 2009, 2011, 2013). In section 2.1, I discuss the reference-anchoring problem, particular versions of it put forth by McKenna and Pereboom and what I take to be a successful revisionist response. In section 2.2, I discuss some potential worries regarding this response and whether or not it saddles revisionists with any particular burdens regarding concept individuation. In section 3, I introduce the normativity-anchoring problem, and argue that this problem poses a unique and serious worry for revisionism.

1. Theory Construction and Methodology: Systematic Revision⁴

As the “new kid on the free-will block,” revisionism does not benefit from the same degree of familiarity as its main competitors (Vargas 2009: 46). In order to see how both the reference-anchoring and normativity-anchoring problems arise for revisionism, it is important to first have a clear picture of how revisionists might go about responding to the prescriptive challenge in the first place. In recent work, Vargas has rigorously motivated, developed and defended the most comprehensive revisionist account of responsibility.⁵ Here I will use Vargas’ formulation as the model by which to judge the view more generally.⁶

First, a successful revisionist prescriptive account of moral responsibility must meet at least the following criteria⁷:

- (C1) It must provide justification for the responsibility system.⁸
- (C2) It must be naturalistically plausible.
- (C3) It must be normatively adequate.

While an in-depth discussion of (C1)–(C3) is well beyond my current purposes, it is helpful to say a bit more about (C1) in particular, as the normativity-anchoring problem bears directly on the issue of whether or not revisionism can satisfy this criterion. Meeting (C1) is best understood as requiring an answer to two questions. The first question involves the justification of the

⁴ For discussion of the distinction between this kind of revisionism and what Vargas calls “revisionism on the cheap” or “repurposing revisionism,” see Fischer et al. (2007: 152) and Vargas (2013: 92).

⁵ See Vargas (2005, 2009, 2011, 2013) and Fischer et al. (2007).

⁶ Strands of revisionism can be found in the work of many participants in the contemporary free-will debate. For examples of some of these earlier strands, see Smart (1961), Heller (1996) and Hurley (2000).

⁷ These criteria are first laid out explicitly in Fischer et al. (2007: 153-55), and are developed in much greater detail in Vargas (2013, ch. 4).

⁸ Vargas defines this system as follows:

Taken as a whole, the responsibility norms and their attendant social practices, characteristic attitudes, and paradigmatic judgments constitute what we can call *the responsibility system* (Fischer et al. 2007: 154).

responsibility system external to facts about the system itself. I will call this question the *external question*:

Is there anything that would, in general, justify our participation in practices of moral praising and blaming? (Vargas 2013: 128).

The second question regards the practices, attitudes and judgments that make up the responsibility system as we find it. I will call this question the *internal question*:

Can we explain our customary patterns of assessment in ways that make it plausible that they are tracking normatively relevant features of agents in the world? (Vargas 2013: 128).

In order to meet (C1), a revisionist prescriptive account of responsibility must provide affirmative answers to both of these questions. In doing so, the account will thus license our ongoing participation in the responsibility system, as well as tie the system to whatever normatively relevant features of agents ground our interests in it in the first place.

With these guidelines in mind, how might one go about actually constructing a revisionist prescriptive account of responsibility? Here Vargas is also instructive. First, we must identify what the responsibility system actually is. We must get clear about what it is we are looking to provide an account of. Second, we must look to the “internal logic and structure” of that system (Vargas 2013: 94). In Vargas’ terms, this requires identifying the *work of the concept* of responsibility, characterized in terms of a commonsense notion of concepts:

As I understand them, concepts carve up and categorize parts of the world. Concepts do a kind of work for us: they demarcate one thing from another. Relatedly, they identify a collection of (we suppose) interrelated inferences we can make about things (2013: 99).

While Vargas hopes to avoid commitment to any particular view of concepts, it is not clear that he can do so. I will return to this question in section 2.2. He defines the work of the concept of responsibility in particular as:

the characteristic roles played by the collection of beliefs, commitments, and distinction-making characteristic of moral responsibility (2013: 99).

With an idea of the work of the concept in mind, we then have two choices. First, we might decide to overturn the framework of our responsibility system. Vargas argues that we would only have good reason to choose this option if the work of the concept itself looks problematic for some very powerful, independent reasons (if, for example, one has good reason to think that the correct metaethical theory stands in direct conflict with it). In the absence of such reasons, there is a second option. We can provisionally accept the

framework of our responsibility system and our identification of the work of the concept, and provide a prescriptive account of moral responsibility capable of grounding and explaining it. And we do this by providing answers to the internal and external questions discussed above, answers that are naturalistically plausible and normatively adequate.

Finally, it is important to note that there is at least one major methodological difference between revisionists and conventional⁹ responsibility theorists. Revisionists draw an explicit distinction between *descriptive* accounts of responsibility (what we think about responsibility) and *prescriptive* accounts of responsibility (what we ought to think about it). And they motivate this distinction via the following skeptical claim about what our intuitions can be said to reveal about moral responsibility:

The skeptical claim: It is possible that our intuitions fail to inform us about what responsibility is, and furthermore we lack good epistemic reasons for thinking that they ever do.

For most conventional responsibility theorists, providing the best account of free will and moral responsibility just *is* getting clear about our concept. The best account, or the one that we should ultimately endorse, is the one that best aligns with our theoretically refined and systematized intuitions. But Vargas makes the interesting and somewhat radical claim that there are few reasons (and even less argumentation) in support of the claim that these intuitions tell us much, if anything, about the nature of responsibility. He puts the point as follows¹⁰:

Revisionists are not bound by intuitions in the same way as compatibilists; revisionists are prepared to acknowledge a difference between what we believe and what we should believe and traditional compatibilists are not. For traditional compatibilists, if the theory gets the intuitions right, and if the theory provides some guidance on handling new or borderline cases, then it has done its work. ... Revisionists, however, cannot always appeal to intuitions, for revisionists disavow those intuitions rooted in our putatively error-ridden folk concepts (Fischer et al. 2007: 216).

Commitment to the skeptical claim above is what best motivates Vargas' revisionist claim that what we think about responsibility does not necessarily tell us much, if anything, about what we ought to think about responsibility, and that we should treat descriptive and prescriptive questions about responsibility separately. Not only is it possible that these intuitions fail to get things right, but we also lack any good epistemic considerations in favor of thinking that sometimes they *do* get things right. And this means that they simply do

⁹ I borrow this term from Vargas (2011: 460), who uses it to refer to theories of responsibility that entail consistency between their prescriptive and descriptive accounts of responsibility (and also those that make no explicit distinction between the two projects).

¹⁰ The quote below refers explicitly to "traditional compatibilists," but can be extended to conventional responsibility theorists more generally.

not have adequate epistemic standing to play the evidential role in our theorizing about responsibility that they do for conventional theorists.

With this picture of how a prescriptive account of moral responsibility is to be constructed, I now turn to an oft-articulated objection to revisionism: the reference-anchoring problem.

2.1 The Reference-Anchoring Problem

Versions of what I call the reference-anchoring problem have been raised in the literature by both Michael McKenna (2009) and Derk Pereboom (2009). Broadly, this problem can be understood as the worry that there is no guarantee that a revisionist prescriptive account of “moral responsibility” will in fact be a genuine account of responsibility. There are two different ways of characterizing this worry. First, one might argue that the reference of “responsibility” is and can only be picked out by our folk concept.¹¹ As such, any alteration of the concept that revisionism is likely to prescribe will entail that we are no longer talking about genuine responsibility. So, in providing an account of revisionist-responsibility and calling it “responsibility,” the revisionist is changing the subject. Second, even if the revisionist is clear about precisely what they mean by “responsibility,” there is a further worry that what the revisionist means is not what we actually care about when we talk about responsibility. So, the revisionist is not only changing the subject, but also is not justified in doing so.

McKenna (2009) points out that the comparisons that revisionists attempt to draw between responsibility and other, more widely accepted examples of conceptual revision are not analogous. Unlike these widely accepted historical examples (for example, our conceptual revision of water), it is not clear that we can *discover* facts about responsibility in the same way that our discovery of the atomic structure of water led us to revise that concept.¹² To suppose that we can revise our concept of responsibility is, McKenna claims, to assume a realist position about responsibility and suppose that there is something that responsibility *is* beyond our concept of it. But it is not clear that responsibility is like this. And even if we were to grant that there is some real nature of moral responsibility beyond our concept, the fact that this real nature does not appear to be discoverable in the same way as the atomic

¹¹ This claim need not depend on an assumption about the correct theory of reference more generally, but only on the assumption that an internal theory is the only one appropriate to “responsibility.”

¹² Here one might object to the claim that we should categorize the latter as a genuine case of conceptual revision, and that this is an unfair comparison. If this is correct, then so much the better for revisionism. However, there is a related, deeper worry that any examples of conceptual revision that revisionists use to help motivate their view might best be understood as examples of changes in our widely shared beliefs about these things, not genuine revision of our shared concepts. Here I again provide the reader with a promissory note that I will discuss these issues in further detail in the next section.

structure of water makes it difficult for revisionists to explain why we are licensed to call this thing (rather than properties and agential features picked out by our folk concept) “responsibility.” In changing the reference of “responsibility,” it looks as though the revisionist is changing the subject, and it is not clear that they are licensed to do so.¹³

Pereboom, on the other hand, grants that it might be possible to revise our conception of responsibility while preserving the same overall concept. And, if revisionists can accomplish this, then they need not be accused of changing the subject. However, even if it makes sense to talk about revising our folk concept of responsibility, Pereboom presses the question of whether or not the revised conception will be:

near enough to the folk’s to count as a natural extension of it, one that can do enough of the work the folk conception does in adjudicating questions of moral responsibility and punishment, and in governing our attitudes to the actions of those around us (Pereboom 2009: 25).

If the answer to this question is negative, then it looks as though the revisionist has gone beyond altering merely our conception, and the resulting view calls for us to use the same term to stand for distinct concepts. But this threatens to result in serious confusion and miscommunication, and the same worries that McKenna raises about changing the subject become salient.

However, if the answer to Pereboom’s question is affirmative, then it looks as though revisionists have succeeded in altering only our conception while maintaining the overall concept, giving them license to use the term “responsibility.” So, a revisionist response to the reference-anchoring problem depends on their ability to answer this question affirmatively, and tie their prescriptive account of responsibility near enough to the folk’s. This would not only address Pereboom’s worries, but likely McKenna’s as well. McKenna’s primary concern is, again, that it seems that:

what moral responsibility is cannot come apart from the concept in such a way that there is, so to speak, something for moral responsibility to be beyond our concept of it (2009: 11).

But if revisionists can successfully tie their prescriptive account of responsibility near enough to the folk’s then it looks as though revisionism need not

¹³ Here I use the example of the purported conceptual revision of water only because it is the example that McKenna himself uses to motivate his arguments against revisionism. But not much hangs on this particular example or the fact that water has empirically “discoverable” features that responsibility seems clearly to lack. Vargas uses many other examples (such as the apparent revision of our concept of marriage), and one might run a version of McKenna’s argument against revisionism using an example that has a more *a priori* flavor. For example, the concept of continuity was revised due to *a priori* mathematical considerations, and it is not immediately obvious what the analogous *a priori* considerations in favor of revising our concept of responsibility might be either. Thanks to Kris McDaniel for suggesting this example.

require that responsibility *come apart* from our concept after all. In jettisoning problematic features of the concept, we need not change the subject entirely. We might still retain enough of the important features of that concept to warrant the claim that we are still talking about responsibility – that is, whatever it is that grounds our attributions of moral praise and blame.

The essential question, then, that the reference-anchoring problem raises for revisionists is the following: Is their prescriptive account of moral responsibility tied, in some way, closely enough to the folk concept to license their use of the term “responsibility”?

There is obviously some ambiguity here regarding what will count as “closely enough.” However, the contours of revisionist-theory construction outlined in section 1 provide strong reason for thinking that a revisionist prescriptive account of responsibility will in fact be tied closely enough to the folk concept. Recall that a prescriptive account of responsibility must, according to the criteria for constructing such a theory, explain how the theory is relevant to what we care about and take to be at stake in our attributions of responsibility. In constructing a prescriptive account of responsibility, revisionists begin by identifying the responsibility system as we find it, and the work of the concept.

In earlier work, Vargas provides some further detail in characterizing the work of the concept of responsibility:

The most useful initial characterization of the conceptual role for moral responsibility is as something that plays an important role in our organization, coordination, and justification of differential treatment of one another. In particular, it is connected to praiseworthiness and blameworthiness. In turn, judgments of praiseworthiness and blameworthiness underwrite a web of emotional reactions, judgments, and social practices that can include (but are not limited to) reward and punishment (Fischer et al. 2007: 154).

This account of the work of the concept sounds quite similar, if not identical, to the features of the folk concept that Pereboom claims revisionists must tie their view closely to. The work of the concept is, at least in part, to regulate our judgments about when agents deserve moral praise and blame.¹⁴ So, it looks as though revisionists who accept Vargas’ model for prescriptive theory construction have a response to the reference-anchoring problem built into their methodology. Not only will revisionists who accept the above characterization of the work of the concept tie their prescriptive account of responsibility to our folk thinking about moral responsibility, they will also

¹⁴ Here one might object that our actual practices of praising and blaming do not always involve the assumption of desert. Some tokens of praising and blaming (especially those directed at children) are solely forward-looking (Pereboom 2013). However, the claim here is merely that the particular kind of praise and blame at issue in the debate between revisionists and their critics is distinctly *moral*, and that the features of practices that are at issue in the discussion of moral responsibility are those features that involve the assumption that praise and blame are deserved.

accept it as one of the most basic features of and starting point for generating their account.¹⁵

Vargas, for example, provisionally accepts the framework of the responsibility system that the work of the concept grounds, and takes the overall revisionist-prescriptive project to be that of justifying our participation in that system and explaining the customary patterns of inference that make up the system itself. The project of justifying and explaining may result in jettisoning some particular features of our folk concept as it currently stands, but not those that are essential or constitutive. If this is the correct picture of the best revisionist methodology for prescriptive-theory construction then it is unclear how one might uphold the charge that revisionists are changing the subject, when it seems that an affirmative answer to Pereboom's question is actually built into their methodology.

The point here is not that Vargas' methodological commitments guarantee that what he is providing a prescriptive account of cannot fail to be moral responsibility.¹⁶ Rather, it is that his methodological commitments guarantee that, if we accept the above characterization of the work of the concept, then what he is providing a prescriptive account of cannot fail to be the same thing (whatever it may be) that his *critics* are calling "moral responsibility." It is an open possibility that all of the interested parties are mistaken, and that none are providing an account of genuine responsibility. However, the main thrust of this response to different versions of the reference-anchoring problem is that Vargas-style revisionists are no worse off than their conventional critics in arguing that they are providing such an account.

Finally, Vargas at least is willing to *grant* that his brand of revisionism might in fact shift the reference of "moral responsibility," but argues that if this is so then the resulting dispute is merely terminological. He distinguishes between two possible kinds of revisionism: *connotational* and *denotational* (2011: 462). According to the former, the beliefs associated with moral responsibility that revisionism prescribes jettisoning "do no substantive work in designating some property in the world," and so do not affect the reference of the term (2011: 462). According to the latter, revisionism does prescribe giving up "some reference-fixing content" (2011: 462). However, so long as there is "some nearby property" very much like responsibility that "preserves the

¹⁵ However, the fact that revisionists take the work of the concept as their starting point for prescriptive-theory construction does not entail that they will ultimately succeed in generating a prescriptive account of moral responsibility that preserves it. This starting point may ultimately need to be abandoned. In this case, the revisionist should grant that the account on offer is *denotational* (I will discuss this in further detail shortly) and that they are in fact changing the subject. However, unlike the brand of denotational revision I discuss below, a denotational revisionist view that fails to preserve the work of the concept will be difficult to defend against the charge that changing the subject is justified. As such, any revisionist view that takes the work of the concept as its starting point but ultimately fails to preserve it looks untenable. Rather, such views would likely recommend something more akin to elimination, not genuine revision. Thanks to an anonymous referee for pressing this point.

¹⁶ Thanks to an anonymous referee for suggesting this line of objection.

primary inferential roles we take to organize our beliefs ... regiments our practices and characteristic attitudes in familiar ways ... weighs in our deliberation in just the same way” and preserves the same “normative import,” then this shift in reference is warranted (Vargas 2011: 462). So, even if one turns out to be a denotational revisionist, the charge of changing the subject will be merely terminological. Whatever the denotational revisionist is referring to after this shift will respect the most fundamental work of the concept. While the denotational revisionist must bite the bullet and grant that she is in fact changing the subject, she has grounds to claim that she is justified in doing so. In light of this, the denotational revisionist might even counter this kind of objection by challenging those who raise the reference-anchoring problem to explain why we ought to care about what we *have* been referring to, rather than the nearby *moral responsibility**, at all.

So, it seems that revisionists can easily respond to versions of the reference-anchoring problem like those raised by McKenna and Pereboom. As long as they adopt Vargas’ methodology for constructing a prescriptive account of moral responsibility and accept his characterization of the work of the concept, it seems they get this response for free.

2.2 Revisionism and Concept Individuation

By now it should be apparent that the reference-anchoring problem raises a series of issues regarding how concepts are to be individuated. The heart of the problem is the charge that revisionists are changing the subject and that what they are talking about is not genuine moral responsibility but something else entirely. Definitively settling the question of whether or not this charge is warranted will likely depend on a worked-out theory of concepts. Unfortunately, there is little to no consensus on what such a theory might look like in the contemporary concepts literature.¹⁷ The *classical theory* – that the structure of concepts is definitional – has been more or less abandoned, due largely to the fact that the theory looks to be inconsistent with a large body of empirical data in psychology.¹⁸ In its place, at least three competing theories have emerged. First, the *prototype theory* is the view that lexical concepts have probabilistic structure. Second, according to the *theory theory*, concepts stand in relation to one another in the same way that terms in scientific theory do, and so are inter-defined and are individuated at least in part by their conceptual role.¹⁹ Finally, *conceptual atomism* stands in stark contrast to all of these views in that concepts are individuated by their relation to the world alone (not to other concepts).²⁰ This is by no means an exhaustive list of the views

¹⁷ For a helpful survey of this literature, see Laurence and Margolis (1999).

¹⁸ In particular, it conflicts with the data on typicality. For a survey of this data, see Murphy (2002).

¹⁹ For further discussion of a recent version of the theory theory view, see Carey (2009).

²⁰ For example, according to Fodor (1990, 1998) a concept stands in a lawful relation to the property it expresses where other lawful relations involving the concept are symmetrically

one might hold on the nature and structure of concepts (for example, one might also be an eliminativist or pluralist), but is a helpful if coarse-grained way of dividing up the terrain.

There are at least three relatively uncontroversial desiderata for any theory of concepts, and none of the above views does a particularly good job of meeting them all. Ideally, a view of concepts should be able to account for the following phenomena:

- (1) Concepts are sharable.
- (2) Concepts are appropriately fine grained.
- (3) Concepts allow for compositionality.

First, a successful view of concepts should have the tools to explain how different individuals (and the same individual over time) are able to share the same concepts. It must make sense of the fact that when I use the concept CAT²¹ and a distinct individual uses CAT we are interested in the same thing. Second, a successful view of concepts should be appropriately fine grained in the sense that it can account for the strong intuition that HESPERUS and PHOSPHORUS or WATER and H₂O are distinct concepts, even though they refer to the same thing. So, it must allow for intensional differences between terms with the same referent. Finally, a successful theory of concepts should allow for and explain how we compose complex concepts out of simpler constituent concepts. If it does not, then the theory will face serious difficulties explaining how concept acquisition and conceptual development are possible.²² While the prototype, theory theory and conceptual atomist views described above all do a particularly good job of meeting one or more of these desiderata, they each do a notoriously bad job of meeting them all.²³ Given

dependent on the relation between the concept and the property it expresses. Other lawful relations *depend* on the lawful relation between the concept and the property it expresses, and would not hold without that relation.

²¹ Here I make use of the common convention of using small caps to indicate reference to a particular concept.

²² One might object that (3) makes conceptual atomism a nonstarter. However, (1)–(3) are intended only as *desiderata* for a theory of concepts, not necessary requirements for a successful theory. In general, we want a theory of concepts to allow for compositionality because we want the theory to account for concept acquisition and learning. That the conceptual atomist is in a difficult position regarding this particular desideratum does not rule out the view automatically, given that its main competitors are in an equally difficult position with regard to either (1) or (2). Furthermore, the conceptual atomist will likely appeal to a combination of some degree of nativism about concepts and an alternative story about concept acquisition involving particular sustaining mechanisms and mind-world relations in order to explain why they need not allow for compositionality. For an example of a particular positive conceptual atomist account of concept acquisition, see Margolis (1999: 559–64).

²³ Versions of the prototype theory can account for sharability and fine grainedness but fail at compositionality. Versions of the theory theory can handle fine grainedness and compositionality well, but face serious worries regarding sharability because they are often committed to *holism* – the idea that concepts must be individuated by attending to their inferential role in its entirety. In light of worries regarding holism, proponents of the theory theory might opt

these tradeoffs, it is not at all clear which theory of concepts might win the day, and so there is little hope of settling whether or not the above response to versions of the reference-anchoring problem can succeed via appeal to any consensus in the concepts literature.

Vargas himself is aware of this difficulty and attempts to avoid taking a stance on the nature of concepts altogether. In order to see how, we might extend his appeal to the distinction between connotational and denotational revision. If one's particular view of concept individuation entails that what the revisionist is really interested in is RESPONSIBILITY* rather than RESPONSIBILITY, then it will be true that the revisionist and the conventional theorist are not talking about the same shared concept. If this is the case, then the revisionist must take great care not to overlook it and it may also be misleading to call the view in question *revisionist*. If the prescriptive account provided by the revisionist is actually an account of RESPONSIBILITY* then the view might more accurately be categorized as a version of *eliminativism*. However, this result would not trivialize revisionism. The view would still stand in sharp contrast to other eliminativist views, such as Pereboom's (2001, 2009) and Strawson's (1993), in its prescriptive recommendation that we *replace* our concept of RESPONSIBILITY rather than simply give it up.²⁴ So, again, the question of whether or not revisionism prescribes conceptual change or conceptual revision according to one's particular theory of concepts amounts to merely a terminological one.

However, there is a further question about whether or not revisionists are really entitled to sidestep any commitment to a particular view of concepts. In various discussions of the work of the concept, Vargas seems to use this term and the "conceptual role" of moral responsibility interchangeably.²⁵ For Vargas, it looks as though the conceptual role of responsibility is what determines the widely shared semantic content of the concept RESPONSIBILITY, and what ensures that we are talking about the same thing, or something relevantly similar, when we use the term "responsibility." If preserving the work of the concept is primarily what revisionists are beholden to in constructing their prescriptive account and this is how the work of the concept is to be characterized, then it looks as though they are in fact committed to a view about concepts that takes the inferential role to be at least partially con-

instead for some form of *molecularism*, the view that only some *part* of the inferential role of a concept is relevant to its individuation. However, theory theorists who go this route must make appeal to some version of an analytic/synthetic distinction, which most readers will immediately reject as hopeless. Finally, conceptual atomists can account nicely for sharability, but are saddled with providing an account of concept acquisition that does not appeal to compositionality, and seem committed to an unsatisfactorily coarse-grained account of concept individuation. For example, this kind of view looks to yield the counterintuitive result that the concepts of HESPERUS and PHOSPHORUS are the same. In short, it is not clear that conceptual atomism leaves room for intensionality.

²⁴ In this case, perhaps *replacementism* would be a better way to distinguish the view terminologically.

²⁵ For example, see Fischer et al. (2007: 154) and Vargas (2013: 99-101).

stitutive. In other words, it looks as though revisionists who accept this characterization of the work of the concept are committed to some version of a theory theory view of concepts.

At first glance this may seem like a cost to revisionism. There is little to no consensus that the theory theory is correct, and it rules out a number of competing views that do not take the inferential role to be constitutive. However, at least four considerations count against thinking that there is a significant cost to this commitment.

First, if commitment to some version of the theory theory is a cost to revisionism, then it counts equally against conventional theories of responsibility. The inferences that Vargas identifies as constitutive of the concept of moral responsibility are the same inferences that conventional theorists pick out. Contemporary compatibilists, incompatibilists and even hard incompatibilists point to precisely the same basic features of the inferential role of the concept of RESPONSIBILITY (whatever does the work of warranting our attributions of moral praise and blame) to ground the claim that they are not talking past one another.

Second, use of the term “inferential role” here is misleading. Talk of the inferential role often is associated with a commitment to holism, the view that the massive web of *all* of the causal and inferential relations that the concept is embedded in determines how that concept is to be individuated. But this is not what Vargas nor his conventional competitors are appealing to. What they are identifying is instead what they take to be the most basic or essential role of the concept, that of regulating our judgments about when an agent deserves moral praise and blame.

Third, while commitment to a version of the theory theory rules out other views according to which inferential role plays no constitutive role in concept individuation, such as straightforward classical theories, prototype theories or conceptual atomism, each of these views faces its own serious difficulties. As such it is not clear that there is any burden on revisionists to justify an appeal to some version of the theory theory beyond the mere desire to avoid unnecessary theoretical commitments wherever possible. Given the relatively controversial status of all of these views, the theory theory does not saddle revisionists with any comparatively heavy theoretical burdens.

Finally, even if revisionists do assume a theory theory view of concepts and it turns out that the theory theory is wrong, this need not constitute a serious objection to revisionism. In this case, the view in question might again best be characterized as a version of eliminativism, not revisionism. But again, this point is merely terminological, and revisionism would still be a unique and interesting position in its prescriptive recommendation that the eliminated concept can successfully be replaced.

Here I wish only to point out that I am skeptical that the reference-anchoring problem can be settled by appeal to a particular view of concept individuation. Nor am I persuaded by the claim that revisionism can avoid any specific commitments regarding the nature of concepts and how they are

individuated. Revisionists who accept Vargas' characterization of the work of the concept appear tacitly committed to some version of the theory theory, but this commitment should not be considered a significant cost to revisionism.

This concludes discussion of the reference-anchoring problem. I turn now to what I take to be a far more serious problem for revisionism: the normativity-anchoring problem.

3. The Normativity-Anchoring Problem

I have argued that revisionism can overcome the reference-anchoring problem, but this problem has a far more troubling normative cousin. The heart of the problem is this: Even if it is built into revisionists' methodology that what they are providing a prescriptive account of is something that we *do* deeply care about, there is a further question about whether or not we *should*. We might grant that their prescriptive account captures the work of our concept of responsibility and the features of responsibility that we actually take to be worth wanting (while departing from our folk concept in some significant way), but why think that this is an account of responsibility that we ought to endorse? What grounds the normativity of this prescriptive account, and does the actual *prescribing*? This question is closely tied to the external question discussed in section 1, of whether or not a prescriptive account can justify our continued participation in the responsibility system and general practice of moral praising and blaming. It does not look as though Vargas' own prescriptive account can justify it, and acceptance of the skeptical claim discussed in section 1 makes it uniquely difficult for revisionists to provide such justification.

According to Vargas, the prescriptive account of moral responsibility he offers satisfies the external aspect of (C1) because, on this account, moral responsibility contributes to the cultivation of a particularly valuable kind of agency: moral-considerations responsive agency (2009: 52-53). In particular, what justifies our continued participation in the responsibility system is that:

these practices foster a distinctive form of agency in us, a kind of agency sensitive to and governed by moral considerations (2013: 177).

But if this is correct then it looks as though the value of our actual responsibility system depends on the value of this kind of agency, and Vargas' own prescriptive account might best be interpreted as a kind of *buck-passing* account of moral responsibility. While his view tracks normatively relevant features of agents in the world (as it must if it is to satisfy the internal aspect of (C1)), it is not *facts about responsibility* that provide us with reason to continue to participate in the practice of attributing moral praise and blame, but *facts about whether or not this practice promotes something that is independently valuable*: moral-considerations responsive agency.

But what is wrong with a buck-passing prescriptive account of responsibility? Moral-considerations responsive agency is clearly a unique form of agency that we do value greatly. However, the fact that we do value this particular kind of agency does not entail that we should continue to participate in any particular practices that promote it, including the practice of moral praising and blaming. In order to see why, it is helpful to distinguish between three distinct claims that might mistakenly be run together:

- (1) *Psychological*: We value moral-considerations responsive agency.
- (2) *Axiological*: Moral-considerations responsive agency is valuable.
- (3) *Normative*: We should continue to participate in any practice that promotes moral-considerations responsive agency.

The psychological claim is obviously true. We do indeed value moral-considerations responsive agency. However, Vargas seems to assume that the psychological claim entails the axiological claim, and that the axiological claim entails the normative claim. And because the practice of praising and blaming (and the responsibility system more generally) promotes moral-considerations responsive agency, our continued participation in this practice is justified.

But it is not obvious that *any* of these entailments hold. Take first the connection between the psychological and axiological claims. It is clearly possible to value something that is not in fact valuable. Human beings make mistakes about what is valuable all the time, and not just on an individual level. Might we be mistaken about the value of moral-considerations responsive agency? Perhaps, for example, the correct metaethical view is some kind of error theory, and so the kinds of considerations that our moral-considerations responsive agency is responsive to are systematically false. It is not at all clear that under these circumstances this kind of agency would be genuinely valuable, regardless of whether or not we actually value it. At the very least, the move from the psychological claim to the axiological claim requires further argument.

In light of this, Vargas might argue that reasons independent of the psychological claim justify our acceptance of the axiological claim. But what sort of reasons? Vargas cannot appeal to the fact that the axiological claim is intuitive or obvious, since his commitment to the skeptical claim blocks him from reading normative claims off of descriptive facts about our intuitions.²⁶ For Vargas, the fact that it is possible that our intuitions get things wrong,

²⁶ At least in regard to moral responsibility, and it is not at all clear that there is any principled difference between responsibility and other normative concepts capable of providing revisionists with a reason to maintain the skeptical claim in regard to the former but not the latter. For further discussion of the motivation for Vargas' skepticism about intuition in this particular domain, see Vargas (2011: 459-60). It may also be helpful to note that my claims about revisionists' potential to offer independent support for the axiological claim are not essential to the overall force of the normativity-anchoring problem. As I discuss shortly, even if we grant the axiological claim, it does not entail the requisite normative claim.

and that we have no good epistemic reason for thinking that they get things right, is enough to block appeal to them as evidence, at least when it comes to matters normative. But then it is not clear what alternative options are available for supporting the claim that we are epistemically justified in accepting the axiological claim that moral-considerations responsive agency is genuinely valuable.

However, even if we grant that Vargas can offer successful arguments to support the axiological claim, support for the normative claim is required in order to meet the external aspect of (C1). The truth of the axiological claim does provide *pro tanto* reason to accept the normative claim, but that is not enough for the task at hand. What Vargas needs to meet the external requirement of (C1) is to show that our participation in the responsibility system should, *all things considered*, continue. The fact that moral-considerations responsive agency is genuinely valuable and that the responsibility system promotes this kind of agency fails to establish this. One might think that we could devise a new system that could promote the cultivation of moral-considerations responsive agency that made no reference to praise and blame at all. Perhaps, for example, we simply reward behavior that promotes this kind of agency and leave praise and blame out of the picture entirely. Pereboom, for example, suggests something along these lines (Kane 2011: 417), arguing that moral admonition and encouragement are sufficient to communicate a sense of right and wrong and achieve effective moral education and improvement.²⁷ If this is correct and there are in fact alternative systems capable of promoting moral-considerations responsive agency, then our continued participation in the responsibility system is not the *only* way to promote this valuable kind of agency.²⁸ And if the responsibility system is not the only way to promote this kind of agency, then appeal to its value fails to show that we should, *all things considered*, continue to participate in the responsibility system.²⁹

So, appeal to the value of moral-considerations responsive agency fails to do the normative work necessary to successfully meet the external aspect of (C1). The psychological fact that we do value this kind of agency does not entail the axiological claim that this kind of agency is actually valuable. And Vargas' commitment to the skeptical claim places him in a particularly diffi-

²⁷ However, Pereboom's arguments assume that giving up on praiseworthiness and blameworthiness need not entail giving up on rightness and wrongness, which has been challenged at length by Haji (1998, 2002).

²⁸ Of course, we could *call* the behavior in question "responsible" behavior, but to do so would again raise versions of the reference-anchoring problem discussed in section 2, and in this case I am inclined to say that we *would* here be changing the subject.

²⁹ Here one might argue that even if there are alternative ways of promoting moral-considerations responsive agency the practice of moral praising and blaming might still be the *best available* way of promoting it. However, given worries about the negative features of moral *blaming* in particular (especially in regard to the negative reactive attitudes that play an important role in many tokens of blame), this does not look like a promising line of argument for revisionists. Thanks to an anonymous referee for this suggestion.

cult methodological position when it comes to providing independent support for the axiological claim. Furthermore, even if we grant for the sake of argument the truth of the axiological claim, this claim does not entail the normative claim. In particular, the fact that moral-considerations responsive agency is genuinely valuable provides only *pro tanto* support for the normative claim that we should continue to participate in any practice that promotes this kind of agency. If the practice of moral praising and blaming is not the only way to promote this kind of agency (and it looks as though it is not) then appeal to the value of moral-considerations responsive agency fails to establish that we should, all things considered, continue to participate in this practice.

Finally, little hangs on Vargas' appeal to the value of moral-considerations responsive agency in particular. The normativity-anchoring problem poses a serious worry for any brand of revisionism that accepts the skeptical claim and offers a buck-passing prescriptive account of responsibility. Commitment to the skeptical claim will make it uniquely difficult for revisionists to offer independent support for the relevant axiological claim, whatever that particular claim may be. Appeals to the intuitiveness of such claims will leave revisionists in, at best, the uncomfortable position of requiring an explanation for why we ought to be skeptical about the epistemic status of our intuitions about moral responsibility but not other normative concepts. And, even if we are willing to grant the truth of the relevant axiological claim, this, again, will not do the work necessary to meet the external requirement of (C1). Revisionists must show that the practice of moral praising and blaming is the *only* way to promote the kind of value referenced in the preferred axiological claim they appeal to. Vargas' own attempt to meet (C1) fails because there are in fact other ways to promote moral-considerations responsive agency. And it is not at all clear what alternative facts about independent sources of value revisionists might appeal to such that the practice of moral praising and blaming will be the only way to promote this kind of value. It looks as though the methodological commitments used to motivate revisionism therefore block revisionists from providing a tenable prescriptive account of moral responsibility capable of satisfying the external requirement of (C1).

Finally, it is important to note that the normativity-anchoring problem is a unique problem for revisionism. Conventional responsibility theorists sidestep this problem entirely by way of their own methodological commitments. According to the conventional theorist, the best account of responsibility, the one that we ought to endorse, just is the one that best captures what we actually think about moral responsibility and what we take to be most important and valuable about it. If one accepts this assumption then there is no reason to ask why we ought to endorse an account of responsibility beyond the fact that it best aligns with our intuitions about particular cases and our systematically refined beliefs and theoretical commitments. In granting this assumption, one grants that these things are in some sense tracking the truth, and so

a theory of moral responsibility generated by this methodology will not need to tell a further story about why it is normatively anchored. The theory will be normatively anchored because, given our standing methodological commitments, it is the account that we have the best reason to think is *correct*.³⁰

Revisionists committed to the skeptical claim cannot appeal to these considerations. But commitment to the skeptical claim is precisely what makes revisionism so interesting. It is what motivates the distinction between descriptive and prescriptive accounts of responsibility and distinguishes revisionism from conventional theorizing about moral responsibility. So, it seems that revisionists are in a uniquely difficult position in regard to the normativity-anchoring problem. They must find some alternative way to ground the normativity of their prescriptive account of responsibility, and explain why the account they generate is one that we ought to endorse. However, if the above arguments are correct then it is not at all clear that they can.

4. Why the “No Real Difference” Objection Makes No Difference

Before concluding, I would like to acknowledge and address one potential objection to the overall goals of this paper. The main thrust of the objection is that I (and perhaps revisionists more generally) am assuming an uncharitable picture of the conventional theorist’s methodology. On a more charitable view, there is no real difference between the methodological approaches of revisionists and conventional theorists such that the normativity-anchoring problem poses a unique and serious problem for the former. More specifically, one might argue that conventional theorists, like revisionists, are aware of and sensitive to the distinction between diagnostic and prescriptive accounts of responsibility, and would not endorse my claim above that the best account, or one that we ought ultimately to endorse, is the one that best captures what we think about responsibility. The conventional theorist, too, is willing and able to acknowledge that what we think about responsibility might be mistaken.³¹

Providing a detailed account of the going methodology for contemporary conventional-responsibility theorizing and explaining how this approach differs from revisionism is itself an interesting and formidable project, and one that I cannot do adequate justice to here.³² However, it may be helpful to note that I am not claiming that what conventional theorists are doing is mere concept mapping as things stand. Rather, I take it to be something more akin to a process of reflective equilibrium, where there is at least some give and take between our intuitions about cases and standing theoretical commitments, which are themselves up for debate. One important difference

³⁰ For those who deny that this is an accurate picture of the *conventional* theorist’s methodology, I address this complaint in the final section.

³¹ Thanks to an anonymous referee for articulating this worry; a similar line of concern about revisionism has also been raised by McKenna (2009).

³² Though for the most comprehensive attempt on offer, see Vargas (2009, 2011, 2013).

between this conventional methodology and Vargas-style revisionism is that the general skepticism about intuition used to motivate the latter's explicit endorsement of the distinction between prescription and diagnosis leaves revisionists without a class of relatively uncontroversial intuitions with which to begin this process. As such they must, like Vargas, appeal directly to justification (in the form of the internal and external components of (C1)), while the conventional theorist assumes that their method of reflective equilibrium will, at the end of the process, generate an account that we have most reason to think is getting things right or best captures the nature of responsibility.

However, for those who wish to maintain what I call the "no real difference" objection despite these considerations, the normativity-anchoring problem will actually be a far *more* serious worry and have much wider scope than I have allowed it here. For conventional theorists willing to disavow that they are in the business of doing some form of conceptual analysis by way of reflective equilibrium, it is far from clear what they are in fact in the business of doing. If such conventional theorists are instead (albeit implicitly) working within a methodological framework similar to revisionists,' then the "no real difference" objection will of course be relevant and revisionism would in large part collapse as a novel view. But, in this case, the normativity-anchoring problem would end up counting equally against the conventional theorist. What is it, if not alignment with our intuitions about cases and principles refined via some process of reflective equilibrium, that prescribes that we endorse the particular conventional view on offer? The revisionist has a story about justification to offer, albeit one that has not yet found much success. The conventional theorist has no such story, and is thus even more poorly suited to respond to the problem.

So, even if there is "no real difference" between these two approaches methodologically speaking, there is an important difference between conventional theorizing and revisionism. In making their methodological commitments explicit, revisionists can tackle the normativity-anchoring problem head on. In failing to clearly explicate their own methodological commitments, conventional theorists sweep this problem under the rug.

In conclusion, I take the normativity-anchoring problem to be a serious worry for revisionism as it has been formulated thus far. This problem has thus far been overshadowed by another family of objections (the reference-anchoring problem), which I argue revisionists can easily overcome. Furthermore, for those proponents of the "no real difference" objection," who remain skeptical that revisionism is itself a novel view worthy of being distinctly situated and defended in the dialectic, I argue that granting this leaves conventional theorists faced with their own version of the normativity-anchoring problem, but with fewer explicit methodological tools to address it than the revisionist.

Kelly Anne McCormick
Washington & Jefferson College
Department of Philosophy
kmcormick@washjeff.edu

References

- Carey, S. (2009) *The Origin of Concepts*, New York: Oxford University Press.
- Fischer, J. M., Kane, R., Pereboom, D. and Vargas, M. (2007) *Four Views on Free Will*, Malden, Mass: Blackwell Publishing Ltd.
- Fischer, J. M. and Ravizza, M. (1998) *Responsibility and Control: A Theory of Moral Responsibility*, Cambridge: Cambridge University Press.
- Fodor, J. (1998) *Concepts: Where Cognitive Science Went Wrong*, New York: Oxford University Press.
- ____ (1990) *A Theory of Content and Other Essays*, Cambridge, Mass: MIT Press.
- Haji, I. (2002) *Deontic Morality and Control*, Cambridge: Cambridge University Press.
- ____ (1998) *Moral Appraisability*, New York: Oxford University Press.
- Heller, M. (1996) "The Mad Scientist Meets the Robot Cats: Compatibilism, Kinds, and Counterexamples," *Philosophy and Phenomenological Research* 56: 333-37.
- Hurley, S. (2000) "Is Responsibility Essentially Impossible?" *Philosophical Studies* 99: 229-68.
- Laurence, S. and Margolis, E. (1999) "Concepts and Cognitive Science," in E. Margolis and S. Laurence, eds., *Concepts: Core Readings*, Cambridge, Mass: MIT Press, pp. 3-82.
- Margolis, E. (1999) "How to Acquire a Concept," in E. Margolis and S. Laurence, eds., *Concepts: Core Readings*, Cambridge, Mass: MIT Press, pp. 549-67.
- McKenna, M. (2009) "Compatibilism and Desert: Critical Comments on *Four Views on Free Will*," *Philosophical Studies* 144: 3-13.
- Murphy, G. (2002) *The Big Book of Concepts*, Cambridge, Mass: MIT Press.
- Pereboom, D. (2013) "Free-Skepticism, Blame, and Obligation," in D. J. Coates and N. Tognazzini, eds., *Blame: Its Nature and Norms*, Oxford: Oxford University Press, pp. 189-206.
- ____ (2011) "Free-Will Skepticism and Meaning in Life," in R. Kane, ed., *The Oxford Handbook of Free Will*, 2nd edition, Oxford: Oxford University Press, pp. 407-24.
- ____ (2009) "Hard Incompatibilism and Its Rivals," *Philosophical Studies* 144: 21-33.
- ____ (2001) *Living Without Free Will*, Cambridge: Cambridge University Press.
- Smart, J. J. C. (1961) "Free Will, Praise, and Blame," *Mind* 70: 291-306.
- Strawson, G. (1993) "The Impossibility of Moral Responsibility," *Philosophical Studies* 75: 5-24.
- Vargas, M. (2013) *Building Better Beings: A Theory of Moral Responsibility*, Oxford: Oxford University Press.
- ____ (2011) "Revisionist Accounts of Free Will: Origins, Varieties, and Challenges," in R. Kane, ed., *The Oxford Handbook of Free Will*, 2nd edition, Oxford: Oxford University Press, pp. 457-74.
- ____ (2009) "Revisionism About Free Will: A Statement and Defense," *Philosophical Studies* 144(1): 45-62.
- ____ (2005) "The Revisionist's Guide to Responsibility," *Philosophical Studies* 125(3): 399-429.