

10

Variance Theses in Ontology and Metaethics

Matti Eklund

Conceptual engineering concerns questions about what concepts we should employ, for various purposes. I would like to place conceptual engineering in a more general theoretical setting. Much of analytic philosophy has been concerned with analysis of concepts we do have, and with investigation into the properties and relations they ascribe. But these concepts, and these properties and relations, are of course just some among all the concepts, properties, and relations there are. There are our ordinary concepts TRUTH, EXISTENCE, KNOWLEDGE, and FREEDOM, and the properties and relations they ascribe. But there are also other possible concepts, including other possible concepts that, while different from our actual concepts, are like them in certain respects—and there are the properties and relations they ascribe. One theoretical project is that of mapping out what kinds of possible concepts there are, and what properties and relations these concepts ascribe. A related project is that of comparing these concepts, properties, and relations along different dimensions of evaluation. These projects are arguably of greater philosophical significance than the one of getting clear on ordinary concepts and the properties and relations they ascribe. What is so philosophically significant about the concepts we happen to have, and about the properties and relations we happen to have concepts of? Engagement in the two projects just described can in turn issue in practical recommendations, and then we are doing conceptual engineering. But conceptual engineering is just one special case of the broader project of investigating what possible concepts there are, and the properties and relations they ascribe.

In this chapter I will discuss some aspects of this broader project. I will begin by discussing a relatively worked example: the so-called quantifier variance thesis discussed in metaontology, according to which there are different existence concepts and none is privileged. I will then make the (obvious) point that the quantifier variance thesis is just a special case of a more general kind of claim. A variance thesis is, generally stated, a thesis to the effect that there is a multitude of different concepts of some particular kind and none of them is privileged. (Naturally, there is much to unpack here.) One can put forward variance theses in other areas. One other area where one can put forward this kind of thesis is in metaethics. I have investigated this case at length elsewhere, and will here primarily focus on two aspects of it.

I will discuss interesting similarities between the metaontology case and the metaethics case, thus attempting to illustrate the benefits of taking a general approach to the issue of variance theses. Then, lastly, I will discuss what general theoretical obstacles there may be to evaluating variance theses: the metaethics case illustrates some such obstacles, and a natural question to ask is to what extent general lessons can be drawn.

1. Ontology

Whereas plain ontology is concerned with what exists, metaontology concerns the nature of ontological questions. Much metaontological discussion concerns the status of ontology as an enterprise: is ontology an enterprise in good standing, or is it somehow misbegotten? Much of the current interest in metaontology is due to Eli Hirsch's writings.¹ Hirsch's writings have introduced two different (but not always carefully distinguished) kinds of theses into the literature:

Quantifier variance: there are different existence concepts, and none of them is privileged over all others.

Verbalism about ontology: (many) ontological disputes are purely verbal, due to the disputants meaning different things by 'there exists' and cognates.²

The theses are different. Even if there are different existence concepts and none is privileged, it can be that the disputants in ontological disputes tend to use 'there exists' with the same meaning. And it can be that while would-be disputants in ontological disputes use 'there exists' with different meanings, one existence concept is privileged over other existence concepts.

These two theses have then been used for criticism of the enterprise of ontology. In the case of quantifier variance the idea is: if quantifier variance is true, what is so interesting about questions about what exists (in the ordinary sense of 'exists')—given as there are other, equally good existence concepts? In the case of verbalism the idea is: if verbalism is true, then the disputes ontologists engage in are merely verbal; the disputants simply talk past each other.

Let me make a few remarks on the less straightforward of these theses, the quantifier variance thesis. As stated, the thesis of quantifier variance immediately invites the question: what is it for something to be an existence concept? This is a hard question, and one that actually has not been much discussed in the literature, but an answer commonly gestured toward is: having the right sort of *inferential role*; more specifically, being governed by the same inference rules as the ordinary existential quantifier. The thesis, famously associated with Quine, that so-called existential quantification expresses existence is presupposed as background. Hence the name 'quantifier variance' for the thesis. (I will keep dropping the 'existential' and

¹ See, for example, the essays collected in Hirsch (2011).

² Note the cautious 'introduced into the literature'. It is doubtful that Hirsch has ever subscribed to quantifier variance as I go on to explain it. And although Hirsch does defend a form of verbalism, the exact verbalist thesis he defends is carefully circumscribed. For some discussion of these matters, see Eklund (2011).

speak of quantifiers and quantifier meanings even though it is specifically existential quantification we will be concerned with.) If instead one thought of so-called existential quantification as not having ontological import and of existence as being expressed by a predicate, the relevant variance thesis for ontology would be that there are different existence predicates, and none of these predicates is privileged. I have formulated the quantifier variance thesis in terms of concepts, but the thesis could equally well be stated in terms of possible existence meanings, and occasionally I will talk that way.

It may be worth comparing a variance thesis which doesn't give rise to this kind of questions:

Liberalized variance: there are different languages within which to state our overall theories of the world, not all employ the ordinary concept EXISTENCE, and no such language is privileged over all others.

Questions about what it is for something to be an existence concept are irrelevant to the liberalized variance thesis. Still the liberalized variance thesis promises to pack the same metaphilosophical punch as the original quantifier variance thesis. If liberalized variance is true, then what is so significant about questions about what strictly *exists* (in the ordinary sense of 'exists')? There are other concepts we equally well could have used to state our overall theory of the world.

The variance theses—both the quantifier variance thesis and the liberalized version—speak of concepts and languages, devices for representing the world. Some may wish to protest that what they are concerned with is *existence*, not the concept EXISTENCE, or the word 'exists', and hence variance theses are irrelevant. But such a protest is beside the point. Even if what we as a matter of fact are concerned with existence and not its representations, one can ask, for example: why focus on existence—that thing which as it happens is ascribed by our concept EXISTENCE—and not what is ascribed by some alternative existence concept?

The quantifier variance thesis speaks of different existence concepts. This invites questions about how concepts are individuated. But the thesis gets its bite from the associated claim that some purported entities may 'exist' in one sense of 'exist' but not another: that the different concepts can differ in extension. So as far as the talk of concepts in a variance thesis are concerned, we can think of concepts as being different exactly when they have different extensions.

Another question regarding quantifier variance concerns what 'privileged' comes to. In some way to be privileged is to be better than the competition. But there are lots of different dimensions along which to evaluate concepts. One concept may be better relative to one aim, another better relative to another aim. One may be more practical to use, another may be more explanatory, a third may in principle have certain aesthetic qualities that the others lack, and so on. A common view in the metaontology literature is that the relevant dimension of evaluation is something like joint-carvingness, naturalness, fundamentality, . . .³ While there are differences between these notions I will treat them under the same heading: *metaphysical eliteness*, or

³ See, for example, Lewis (1983, 1984); Schaffer (2009); and Sider (2011).

eliteness for short. A common idea is that what is elite is what makes for objective similarity, and what is fundamentally explanatory. There is much to say about this notion of eliteness (as well as objections to address regarding this notion). But for present purposes I will treat the notion as perfectly in order, and none of what I will go on to say will depend on finer details regarding what eliteness comes to.

Saying that the relevant dimension of evaluation for the purposes of the variance theses in metaontology is eliteness is not to say that eliteness is the only possible relevant dimension of evaluation for existence concepts. It could be, for example, that the most elite existence concept is impracticable to use for creatures with minds like ours, and then for various practical purposes other existence concepts will be preferable. A different dimension along which existence concepts could be evaluated is then: suitability for these particular practical purposes. In principle, one could consider different quantifier variance theses corresponding to different dimensions along which concepts may be evaluated. But a working assumption has been that to focus on eliteness is to focus on something that is directly germane to standard concerns of ontologists.

A variance thesis seems to me to be of deeper significance than verbalism, as far as criticism of the enterprise of ontology is concerned. Here is one argument for why verbalism is not of principled importance regarding the enterprise of ontology.⁴ It is compatible with verbalism that there is a privileged existence concept. But then suppose, for example, that ontological disputes tend to be purely verbal, as verbalism says, but there is such a privileged existence concept. Then ontology could simply be recast as an enterprise of asking what exists in the sense of *that* existence concept. It could be insisted in response that even if there is a unique, privileged existence concept, it is simply impossible to have a non-verbal dispute over what exists in the sense of that. The idea would be that even if there is a privileged existence concept, when two theorists have a philosophical dispute over, as they would put it, 'Fs exist', it is simply *inevitable* that they do not use 'exist' to express the same existence concept. I take this to be clearly absurd. (Note that to say that this is absurd is not to deny that many ontological disputes as they actually are conducted are merely verbal.)

That verbalism lacks principled significance regarding the enterprise of ontology does not mean that a variance thesis does better in this regard. Why cannot the practicing ontologist respond to the thesis of quantifier variance by saying: I don't care if existence (in the ordinary sense) is not privileged—I will continue focusing on it as before anyway? If someone consistently holds on to this stance, there may be no arguments that could rationally sway her. To each her own. The supposed point of quantifier variance rather comes in through the assumption that the significance that ontologists attach to their enterprise is due to their thinking of questions of existence as deep, in a way properly spelled out in terms of existence being privileged.⁵

⁴ For a different but related argument to the same effect, see Eklund (2016).

⁵ Sider (2011: 62) holds that it is a correctness condition on belief that it track the world's "structure" (what is elite). Given Sider's view, to not care about getting structure right is akin to not caring about getting at the truth.

A variance thesis compares existence concepts with each other and concerns whether one existence concept is privileged over all others. One may think a better question to ask concerns not about how privileged some existence concept is compared to alternative existence concepts, but how privileged it is full stop. So long as it is privileged enough it, or what it ascribes, is a worthy object of investigation. While the question has not been the focus in metaontology it is worth keeping it in mind, also for when, later, parallel issues in metaethics are considered.⁶

Largely in response to what he sees as the threat posed by the thesis of quantifier variance, Ted Sider has developed and defended what he calls *ontological realism*: the thesis that there is a privileged existential quantifier meaning.⁷ This is in opposition to the quantifier variance claim that no existence concept is privileged. It is worth stressing that Sider's ontological realism is not simply the thesis that ontological sentences—sentences about what exists—have objective truth-values, and that some atomic ontological sentences are true. That much is fully compatible with quantifier variance. What sets Sider apart from the friend of quantifier variance is Sider's claim that one existential quantifier meaning is *privileged*. Given that eliteness is the relevant dimension of evaluation, the claim is that one meaning is more elite than the others.

It has become more and more common in recent years to hold that it is not sufficient for "realism" about discourse D that D-statements are capable of objective truth and some D-statements are true—realism requires something more *beefed-up*. Apart from Sider on ontological realism one might compare Kit Fine on realism, Crispin Wright on marks of realism, and the discussion in metaethics of creeping minimalism. Kit Fine (2001) operates with a primitive notion of what is real, such that it can be that F without it being the case that in reality, F. Statements of a discourse can then be true without them holding in reality. Crispin Wright (1992) distinguishes between different marks of realism. Many of the marks are held to gauge objectivity of a discourse. But one of them, what Wright calls wide cosmological role, has, in brief, to do with the explanatory power of the truth of a class of statements. If statements of discourse D are true but lack wide cosmological role, then discourse D is not fully realist. The problem of creeping minimalism in metaethics (see James Dreier 2004) has to do with the fact that even the non-cognitivist can, as it is often put, earn the right to speak of normative statements as true and mind-independently so: and this raises the issue of what the would-be realist can say to distinguish herself from this non-cognitivist. In each of these cases, there are pressures to say that realism demands something more than mind-independent truth. In this way, these views are like Sider's ontological realism. But one central thing that sets Sider's ontological realism apart is that the fate of Sider's ontological realism is explicitly not bound up with how *actual* ontological discourse works. It is no part of

⁶ Another complication regarding quantifier variance is the following. How threatening quantifier variance is to the enterprise of ontology may depend on details not captured by the above statement of the thesis. Quantifier variance as stated could be true for the reason that there are exactly two existence concepts that are maximally, and equally, elite (and they are, moreover, nearly coextensive); or it can be true for the reason that there is a wide variety of significantly different best existence concepts.

⁷ See Sider (2009) and (2011: chapter 9).

Sider's ontological realism that the *actual* existential quantifier has the privileged meaning: the extra demand that Sider imposes is only that some possible quantifier has a privileged meaning. Whether or not one accepts Sider's view on what makes for privilege, something seems right about the focus on possible instead of actual discourse. It could, for example, be that non-cognitivism, or an error theory according to which no atomic statements of the discourse are true, is correct regarding actual normative discourse but one could engage in some possible normative discourse which is cognitivist and where some atomic statements are true—and, generally, satisfies any demands a realist may wish to impose. This could be sufficient for the realist's demands. The important thing for the realist about the normative is that there are genuine normative aspects of reality; not that our actual languages or conceptual scheme contains the means to pick them out.

Someone like Fine or Wright could certainly take on board this aspect of Sider's view. For example, Fine could say that what matters for normative realism is not whether actual normative sentences express truths which hold in reality, but instead whether some normative truths hold in reality. A would-be normative realist concerned with creeping minimalism could be more concerned that some possible normative discourse has features by virtue of which it cannot be understood as non-cognitivist than that actual normative discourse does so. What I wish to emphasize is just that the fact that the point that the details about actual language do not matter is something that is stressed by Sider but is not a theme in other discussions of beefed-up realism.⁸

2. Variance

Hirsch focuses on ontology. But the same themes can rather obviously crop up elsewhere:

X-variance: there are different X concepts, and none of them is privileged.

Verbalism about X: (many) X disputes are purely verbal.

For the same reasons as given earlier, I believe that a variance thesis is of greater potential significance than a verbalism thesis. Let me then focus on variance theses.

Many philosophical debates—over knowledge, free will, meaning, . . . —concern (or, as I will turn to shortly, are conducted *as if* they concern), our actual concepts, or, better, the properties and relations they stand for.

The subject matter, say knowledge, is taken as given and theorists concerned to, by for example eliciting judgments about cases, figure out the nature of knowledge. It is not asked whether there may be some other epistemic relation—call it knowledge*—such that knowledge* is of greater epistemic significance than knowledge. (To take a tired example: it is as if natural scientists took the notion of weight for granted and sought to illuminate what weight is, without considering whether there are other

⁸ Sider was of course not the first to defend a realist thesis while formulating his preferred view not in terms of ordinary notions but in terms of reformed counterparts. For example, in metaethics, Railton (1986) does the same, relating to Brandt's (1979) earlier talk of reforming definitions. What is new in Sider is the explicit focus on this way of conceiving of realist theses.

notions, like various notions of mass, which are capable of doing a better job as far as physical theory is concerned.)

In each case one can wonder whether our actual concept is the best concept in the relevant class. Take again knowledge. The post-Gettier literature has seen many different suggested analyses of the concept KNOWLEDGE. Whether or not these analyses succeed as such, they characterize various possible knowledge concepts. It is even possible to hold that even if, as per Gettier, knowledge is not justified true belief, the best knowledge concept is a justified true belief concept.

In response to what I have just urged, it may be objected (and here the *as if* from above comes in) that philosophers who are concerned with, say, knowledge are not in fact concerned with the ordinary folk concept and what it ascribes. If one considers how epistemological inquiry is actually conducted, one will find that philosophers are already concerned with improving concepts, and are using a somewhat technical knowledge concept, one perceived as meeting theoretical needs better than the folk concept KNOWLEDGE does. I actually think there is a lot to this objection. But even if what the objection alleges regarding the current state of philosophy is correct, two points deserve stressing. One is that the substantive point remains, regarding the justification for not merely focusing on our actual concepts and what they ascribe—it is just that the advice given may already be followed. Second, even if in fact philosophers tend to use somewhat technical concepts, the strategy of doing so has not been always followed in a self-conscious way, and it has not often been explicitly considered.

I believe the above points about how we ought to consider alternatives to our concepts are all rather intuitive. But there are obviously questions to be raised, analogous to questions raised regarding quantifier variance.

What, in general, makes something an X concept? When introducing quantifier variance, I made some remarks about what counts as an existence concept. But the issue now arises in a more general setting. If one understands 'X concept' as 'concept of X', one can take this to be a concept that ascribes or refers to X. But that idea is a non-starter: for we wish to be able to regard non-coextensive concepts as X concepts. For example, the actual concept KNOWLEDGE and the concept JUSTIFIED TRUE BELIEF can both be knowledge concepts even if some cases of justified true belief are not cases of knowledge. What one might wish to say is something of the form: an X concept, in the relevant sense, is a concept that could play the X role. The thought would be that even though our concept KNOWLEDGE is not coextensive with the concept JUSTIFIED TRUE BELIEF, the latter concept of justified true belief could play the knowledge role—it could be used for epistemic evaluation in the way the actual concept KNOWLEDGE actually is. In the discussion of existence above, it was said that what makes something an existence concept is its having the right inferential role. This can be made to fit the present mold: having the right inferential role suits a concept to do what the actual existential quantifier does for us. Note that matters here are delicate; not to say slippery. The specific claim in the case of existence concepts, that being governed by some specific inference rules is what makes something an existence concept, arguably does not generalize. It is not plausible that sharing of such structural, inferential features is necessary and sufficient for being a knowledge concept.

I am talking about ‘the’ X role, and I will for the most part talk about ‘the’ role a concept has. But obviously, if the role talk is acceptable in the first place, a given concept can be used to play different roles. Taking this into account would complicate some of my formulations. I will instead continue to, naively, speak about ‘the’ X role.

It is important to be clear on exactly what the talk of roles amounts to in the present context. It is essential to the present notion of role that different, non-coextensive concepts may play the same role. It is if this assumption is met that we can ask significant questions about which one of these concepts best plays this role. Compare, by way of contrast, the notion of role that is employed in David Lewis’s philosophy of language. On Lewis’s descriptivist view, descriptions associated with a term amount to a reference-fixing theory associated with the term, and the term refers to whatever best satisfies this associated theory. This is often expressed as: the theory specifies a *role* for something to play, and the term refers to whatever plays this role or comes closest to doing so.⁹ Assuming this conception of reference-determination works for ‘knows’, it is a given that ‘knows’ ascribes the relation that best plays the knowledge role, in this sense.

What is the knowledge role, in the present sense? At a first stab, one may appeal to the use of knowledge in epistemic evaluation. But by itself, this is rather unhelpful. There are different kinds of epistemic evaluation. The uses of the concept KNOWLEDGE are different from the uses of JUSTIFICATION, RELIABILITY, To play the knowledge role is to play the specific role knowledge plays in epistemic evaluation. But this slogan is problematic in two different ways. First, it sounds rather uninformative. (Concept X plays the X role—duh.) Second, uninformative though it may be, on one natural way of understanding the slogan it may yield unwanted results. Compare our actual concept KNOWLEDGE with the concept JUSTIFIED TRUE BELIEF. Assuming that a lesson from the Gettier cases is that for knowledge that P some anti-luck condition not satisfied by mere justified true belief would have to be met, one may think: so the specific role of knowledge in epistemic evaluation is to rule out epistemic luck of the relevant kind. So justified true belief does not play the knowledge role; not even poorly. This is an unwanted result, not because it is a given that the concept JUSTIFIED TRUE BELIEF must count as a knowledge concept but because the reasoning seems to generalize. For many broadly knowledge-like concepts not coextensive with the actual concept KNOWLEDGE one could construct similar arguments that these concepts are not apt to play the knowledge role: given the differences in extension between such a concept C and the concept KNOWLEDGE, C does not play the same role as KNOWLEDGE. Or so the reasoning runs. For the talk of the knowledge role to do the work it is supposed to do in the context, roles must somehow be individuated differently.

In the case of quantifier variance, I noted that one can sidestep questions about what it is to be an existence concept by appealing to a notion of liberalized variance. One can similarly seek to sidestep questions about what it is to be a knowledge concept by appealing to a corresponding notion of liberalized variance—a notion of

⁹ See, for example, section 2 of Schwarz (2015) for a nice exposition of Lewis’s views on this.

liberalized epistemic variance. The idea would be that one can sidestep questions about what it is for a concept to be a knowledge concept by appealing instead to languages that contain expressions expressing different epistemic concepts but are equally good, in whatever dimension is relevant for epistemic purposes. This avoids the problems regarding how to identify knowledge concepts, but immediately highlights a separate question about the various X variance theses: how is the 'equally good' to be understood? In the metaontological case this is, as above noted, often cashed in terms of eliteness. But even assuming that this is the best way to cash it there, that does not mean that this is always the best way to cash talk of concepts being equally good. A philosopher with her metaphysics hat on can certainly embark on the project of evaluating all sorts of concepts for eliteness, and nothing I wish to say here is meant to suggest that this would not be a worthwhile project. But even if it is a worthwhile project, it is not clear what relevance it has for epistemology, or philosophy of action, or To elaborate: One way that variance-type issues can arise for the epistemologist is that one can reasonably worry that concepts different from our actual ones are better for epistemic purposes. Goodness for these epistemic purposes need not line up with eliteness. There are separate questions about what concepts are best for these epistemic purposes.

Another illustration of how philosophically central concepts may be evaluated along different dimensions is provided by the case of personal identity. One can approach the issue of personal identity with the metaphysician's hat on and wonder which person-like entities exist, which person-like entities are the most fundamental or joint-carving, and so on. Or one can approach the issue with the practical philosopher's hat on, wondering about how praise and blame should be distributed, and how our prudential concerns should be structured. The investigations may line up. Maybe the person-like entities that are metaphysically privileged are the ones that are relevant to the practical philosopher's questions. Maybe the metaphysically privileged person-like entities are bodies and it is also the case that if person A at *t* does something blameworthy then it is the person with A's body at *t** who ought to be blamed for this. But the investigations may also come apart. Maybe a physical criterion of personal identity is correct for the metaphysician's purposes while a psychological criterion of personal identity is correct for the practical philosopher's purposes. More radically, it may be that as far as the practical philosopher's purposes are concerned, the focus on identity is misplaced. Parfit's (1971, 1984) arguments regarding, for example, fission are naturally seen as having this upshot. What "matters" in personal identity is the holding of a psychological relation which does not have the logical characteristics to be an identity relation.

Focusing on variance theses is of a piece with seeing philosophy as *conceptual engineering*. Those focusing on conceptual engineering think that rather than resting content with what concepts we actually have, we should think about how these concepts can be improved or replaced. To focus on variance is to shift attention from our actual X concept to what possible X concepts there are and how they are to be ranked along some dimension.

But there are differences, at least differences of emphasis, between conceptual engineering and focus on variance theses. The conceptual engineering project is held to have some practical import: recommendations are made regarding which concepts

to use. One can in principle be skeptical of that project—“is it really the business of philosophy to reform language and thought?”—and still think that there are reasonable variance questions to ask about betterness along some dimension. Using ‘conceptual ethics’ as a label for the enterprise of evaluating concepts, studying variance theses is part of conceptual ethics, even if it is not directly geared to proposals for language reform or conceptual reform.

That said, one should not exaggerate the differences between the overtly practical, activist project of conceptual engineering and the on the face of it more theoretical project of evaluating concepts along different dimensions. Friends of conceptual engineering tend to stress that they are not concerned to reform how we ordinarily think and talk, but only suggest replacement *for particular purposes*.¹⁰ And if one says a concept is better than another along one dimension, one thereby says that the concept is better to use for some associated purpose.

Still, even though one should not exaggerate the difference, there is a difference between the variance-related project of mapping what possible concepts there are and how they are related, and the engineering project of making particular recommendations regarding concept use. (A well-known point from ethics serves to highlight and dramatize the difference: One can think that consequentialist normative concepts are the ones that get at the features that really matter normatively, while at the same time—and on consequentialist grounds—thinking that it would be bad if agents making decisions about how to act deployed consequentialist concepts.)

3. Thin Normative Concepts

One place where one can ask questions similar to those that have come up in metaontology is in the case of normative concepts:

Normative variance: there are different rightness, goodness, ... concept, and none of them is privileged.

Verbalism about normative discourse: (many) normative disputes are purely verbal.

Verbalism about normative discourse does tend to come up as a topic in philosophical discussions, but typically in negative arguments. If a theory of how the reference of normative terms is determined leads to verbalism about normative discourse, that is seen as reason to give up the theory. This is for example a theme in the lively debate over moral disagreement.¹¹

As before, and for the same reasons, I think the variance thesis is the more significant one. So regardless of the plausibility, or not, of verbalism I will set it aside and focus on variance.

¹⁰ So, for example, in Scharp (2013), the main theme of which is that the ordinary concept TRUTH ought to be replaced, Scharp keeps reminding the reader that the replacement is only for certain purposes. For everyday use, the ordinary concept works just fine.

¹¹ See, for example, the Moral Twin Earth argument due to Horgan and Timmons (1992 and 2009).

I have elsewhere discussed at some length various questions relating to normative variance (although I have primarily discussed these matters in different terms).¹² In this section, I will briefly rehearse some main points.

The same questions arise regarding normative variance as regarding other variance theses. What makes something a rightness concept (goodness concept, etc.)? What is it for a rightness concept (goodness concept, etc.) to be privileged? I will soon pause on these crucial questions. But before focusing primarily on these questions, let me first pause on the significance of the normative variance thesis. The discussion of this matter will *inter alia* shed some light on the questions mentioned.

Consider the following scenario (which may or may not be possible):¹³

Tragic There is a linguistic community—the Tragic—speaking a language much like English, except for the following differences (and whatever differences are directly entailed). While their words ‘good’, ‘right’, and ‘ought’ (in their “thinnest” uses) are associated with the same normative roles as our words ‘good’, ‘right’, and ‘ought’ (in their “thinnest” uses) are associated with, their words aren’t coextensive with our ‘good’, ‘right’, and ‘ought’. So even if they are exactly right about what is ‘good’ and ‘right’ and what ‘ought’ to be done, in their sense, and they seek to promote and to do what is ‘good’ and ‘right’ and what ‘ought’ to be done in their sense, they do not seek to promote what is good and right and what ought to be done. Moreover, what their ‘good’, ‘right’, and ‘ought’ are true of are things that really ought not to be valued: their normative language is in that way off.

I am not claiming that the Tragic scenario is in fact possible. But there are prominent views on normative language on which it clearly is. For example, views on the reference of normative terms on which the reference is determined causally as much as the reference of natural kind terms is usually held to be, or views on which widely held beliefs linking these normative terms to the descriptive play a reference-fixing role. Their use of their normative terms may be causally linked in the relevant way to properties that ought not to be valued, or their widely held beliefs link the terms to properties that ought not to be valued.

Now, if Tragic is possible, then it would appear that we could in principle be in the same kind of situation. If the reference of normative terms can be determined in either of the ways characterized, the same can go for the reference of our normative terms. But then our positive (/negative) terms too can be causally linked to properties that do not warrant the positive (/negative) evaluation associated with the term. It is harder to *state* the problem as it arises in our own case. For we use our own normative terms when attempting to state what the supposed problem is regarding our normative terms. (Compare a loose analogy: radical indeterminacy arguments such as those presented by Kripke and Quine are typically not first presented as concerning our language now: it seems clear that, for example, ‘rabbit’ refers to rabbits. The standard strategy is to argue that a language qualitatively like ours is radically indeterminate, and then note that our language cannot be different in that regard.)

¹² Eklund (2017).

¹³ Again, see Horgan and Timmons, for example (1992 and 2009).

Let me now state the issue just brought up in more general terms. Consider a scenario like those encountered in the Moral Twin Earth literature:¹⁴

Alternative There is a linguistic community speaking a language much like English, except for the following differences (and whatever differences are directly entailed). While their words ‘good’, ‘right’, and ‘ought’ (in their “thinnest” uses) are associated with the same normative roles as our words ‘good’, ‘right’, and ‘ought’ (in their “thinnest” uses) are associated with, their words aren’t coextensive with our ‘good’, ‘right’, and ‘ought’. So even if they are exactly right about what is ‘good’ and ‘right’ and what ‘ought’ to be done, in their sense, and they seek to promote and to do what is ‘good’ and ‘right’ and what ‘ought’ to be done in their sense, they do not seek to promote what is good and right and what ought to be done.¹⁵

Typically when scenarios like this are brought up, they are used to evaluate the plausibility of the predictions of different theories of reference-determination. More specifically, we are supposed to have the intuition that the words really are coextensive, and the scenario described not possible—and theories in conflict with this are to be rejected. Thus employed, the scenarios are used to gauge how our actual words and concepts are used. In the context of considering normative variance, a different question arises: *Suppose, provisionally, that scenarios like Alternative are possible. Then what?*

One natural thought is that there then is a question of which of the rightness concepts one ought to employ. Some things are right; others are right*. There is a question of whether to act in accordance with what is ascribed by one concept or what is ascribed by the other.

But there are complications regarding how to understand this supposed further question. When trying to state it just now I used ‘ought to employ’. But if there are different rightness concepts, there are different ought-concepts. And the question of which rightness concept we ought to employ is different from the question of which rightness concept we ought* to employ. If we use one ought-concept we ask one question; if we use another, we ask a different question. Neither of these questions seems to be the one we wanted to ask when we wondered which rightness concept was objectively privileged. For they both are stated *using* some normative vocabulary or other, and what is in question is the propriety of using these pieces of normative vocabulary. Relatedly, one may suspect that it is rather trivial that we ought to care about what is right but ought* to care about what is right*. These two sets of facts—about what we ought to care about and ought* to care about—don’t immediately bring us any closer to the practical question of how to structure our concerns.

The same sort of problem would seem to arise regardless of which normative vocabulary we would use when trying to state the supposed further question. And if instead we tried to state the further question using only descriptive, non-normative vocabulary, our attempts to state the supposed further question would misfire in

¹⁴ Again, see Horgan and Timmons, for example (1992 and 2009).

¹⁵ From Eklund (2017: chapter 2).

another way: we didn't just wish to know which descriptive concepts the various rightness concepts fall under but which one, so to speak, *really ought* to guide action.

Noting the difficulties in stating the supposed further question, one might wish to deny that there is a further issue there. There is what is right, what is right*, etc., and *that is that*. It is right to do what is right, right* to do what is right*, etc. Maybe this in the end is the correct view. But I think many of us are intuitively inclined to take normativity to be *objective*, in such a way as to find this view repugnant. (Suppose, e.g., to dramatize things, that among the actions that are right* are some that we find deeply abhorrent.)

The notion of the objectivity of normativity alluded to in the last paragraph is important, but elusive. Even someone who holds that there is what is right, what is right*, etc., and that is that *can* hold on to the objectivity of normativity in the sense that she can hold that it is an objective matter what is right, an objective matter what is right*, etc. Facts about what is right, right*, etc., are normative facts, it may be said, and it is an objective matter whether they obtain or not. The sense in which the objectivity of normativity is jettisoned on this view is that there does not seem to be a fact of the matter as to whether to go with what is right or what is right* (or ...) in one's choices about how to act.

We can put the above reflections in the form of a dilemma, what I will call the *alternative concepts dilemma*. Either there is a further question of the kind indicated or there is not. The former alternative seems problematic, for the supposed further question would be unstatable. The latter alternative seems problematic for the mere 'that is that' does not capture our sense that the normative is objective.

Return now to the questions about how to understand normative variance. One question was: what is it for a concept to be a rightness concept (etc.)? A natural reply is to appeal to a concept's normative role—its role in action-guiding and deliberation, perhaps its relation to reactive attitudes, etc. Saying that there are different rightness concepts then amounts to saying: there are different concepts associated with the same normative role in this sense but different in other ways, so that they are not coextensive. And the view that there are different rightness concepts in this sense is what gives rise to the alternative concepts dilemma. These remarks on normative role are obviously vague and sketchy: but one can still see that if a concept's reference is determined by what its use is causally related to or by what descriptions the concept is associated with, its reference is not determined by normative role alone.

When it comes to what makes a rightness concept privileged, the problems in cashing this out were in effect displayed through the discussion of the alternative concepts dilemma. If we explain this using normative terms the problem is that there are different normative terms we could use. If we do it using descriptive terms we don't seem to address the right thing.

Let me elaborate on the last point. As earlier stressed, there are many dimensions along which concepts may be evaluated and compared. For some purposes, for example when one is concerned with metaphysics for metaphysics' sake, one may wish to compare rightness concepts in terms of eliteness. If one tries to do so in the present case, one tries to ask the further question in descriptive terms. When it is insisted that we are asking the wrong question about privilege if we attempt to pose this question in descriptive terms, what is claimed is only that if we are approaching

the issue of variance from the perspective of normative theorizing and we are primarily concerned with questions about how to think and act. The metaphysician's question of which rightness concept best carves the world at the metaphysical joints is not immediately such a question.

When discussing the quantifier variance thesis and its supposed deflationary consequences for ontology, I mentioned that one may consider it sufficient for ontology to be in good standing if some existence concept is privileged: it does not matter if other existence concepts also are privileged. Even if this is a reasonable view in the case of ontology, its counterpart in the present case would be misguided. Assuming one sees the considerations that I have brought up in the present question, by appeal to the scenarios *Tragic* and *Alternative*, as serious, it does not help at all if my rightness concept is privileged, if some alternative rightness concept is equally privileged.

One natural way to attempt to avoid the alternative concepts dilemma is to deny that there are these different rightness concepts to begin with: *Alternative* is not possible. One can insist that normative role *determines reference* so that if two concepts are associated with the same normative role they are guaranteed to have the same reference (and more generally, same intension). On a fine-grained way of individuating concepts one can perhaps insist that there still are different rightness concepts, it is only that they are all coreferential. But since the concepts necessarily apply to the same things, there is no momentous question regarding which one to employ. The dilemma is avoided, for if there are not these different rightness concepts, questions about what to say about one being privileged do not arise.

A variance thesis is a conjunction of two claims: one to the effect that there is a multitude of concepts of such-and-such a kind, and one to the effect that no concept in this multitude is privileged. The present suggestion amounts to denying the first conjunct of the relevant normative variance thesis. The move has an analogue in the original metaontology case: it can be insisted that there is no multitude of existence concepts, for example on the ground that any two concepts governed by the standard inference rules associated with existential quantification must be coreferential.¹⁶

Of course, good questions can be asked about whether normative role can indeed determine reference in such a way that the alternative concepts dilemma can be avoided in the way suggested. But my aim here is not to evaluate this suggestion. I am only concerned with the conditional claim that if normative role determines reference, then problems like the ones brought up in connection with the alternative concepts dilemma can be avoided.

This conditional claim itself can reasonably be resisted. If normative role does determine reference, there can be different possible normative concepts associated with *slightly different normative roles*, and the same issues as before can be brought up by appeal to such possible concepts. Compare again ontology. Even if, among the different existential quantifier meanings there are, where the condition for being an existential quantifier meaning is that of satisfying the classical inference rules, some unique meaning is privileged, there can be other quantifierish meanings, satisfying

¹⁶ For relevant discussion, see, for example, Williamson (1988); McGee (2006); and Turner (2010).

some slightly different rules, and no quantifierish meaning is privileged over all the others.) Compare too how in the ontology case one could without obvious loss consider a liberalized variance thesis instead of the quantifier variance theses focusing specifically on existence concepts. Analogously one can in principle, in the normative case, focus on a liberalized normative variance thesis which does not focus specifically on rightness concepts (or goodness concepts, or ought concepts, or reason concepts), but instead concerns different normative languages as wholes.

Moreover, just to make things really confusing (sorry!): just as one can reasonably think that there are alternative notions of rightness, goodness, etc., and questions about whether any particular ones among these notions are privileged, one can reasonably think that there are alternative notions of reference and questions about whether any notion of reference is privileged. With this complication in mind: which notion of reference should we employ in the thesis that normative role determines reference?

4. General Lessons

I have talked about issues related to normative variance, and compared the issues that come up in this case with parallel issues that come up in the parallel metaontological debate. Let me now ask: are there more general lessons regarding variance and conceptual engineering to be learned here?

Consider first the possibility that there genuinely are different rightness concepts. Then, as stressed, there arises the question of whether one is privileged, and, more fundamentally, what privilege amounts to in the relevant case. In the case of rightness concepts, there were problems regarding what privilege might amount to. There seemed to be no way of getting at the supposed further question.¹⁷

One limitation of variance inquiry is presented by the type of case we may be faced with here: there is no way to make sense of the relevant question of privilege.

Let an *ultimate* concept be an X concept such that there are other X concepts and the question of which X concept is privileged cannot be asked in suitably independent terms. The possibility just described is that thin normative concepts are ultimate concepts.

Another possibility in the case of rightness concepts is that there are not, in the relevant sense, different rightness concepts: normative role determines reference, so any concept with the normative role associated with rightness has the same reference. I mentioned that it can be and has been argued to be so also when it comes to existence: the inferential role associated with being an existence concept is such that no two non-coextensive concepts can be associated with this role.

¹⁷ *Intuitively*, this is related to rightness being in some sense *basic*. When it comes to other, less basic normative concepts, like paradigmatic thick concepts such as COURAGEOUS, LEWD, RUDE, . . . , one can reasonably think that, say, a courageousness concept being privileged simply is a matter of it being the courageousness concept one *ought* to evaluate people and actions in terms of. However, once one has problematized 'ought' and raised to salience the possibility that there are different possible ought-concepts, matters look more complicated here too.

Let a *fixed* concept be an X concept such that there are no other X concepts. The possibility just described is then that the concept rightness is a fixed concept.

Ultimate concepts and fixed concepts can appear to present complications for evaluation of variance theses, and for the project of conceptual engineering. When C is ultimate we cannot get a handle on the relevant question of privilege when it comes to C: that is part of the characterization of what it is for a concept to be ultimate. And when C is ultimate, some questions we wish to ask about whether to replace C cannot really be asked. When it comes to fixed concepts, the problem is that when X is a fixed concept, there just are no other X concepts to replace X with.

However, I do not think that fixed concepts in fact do present serious theoretical problems for inquiry into variance theses. Even if there is only a unique X concept there can be concepts in various ways similar to X concepts, and one can still ask whether the X concept is privileged, in whichever respect is relevant, over these other, similar concepts. And if we have a situation where there not only is a unique X concept but moreover there is only one concept in question that can serve the purpose at hand, so that there is no competition, this is not so much an obstacle to inquiry into variance theses as a result regarding what sort of variance there can and cannot be.

Ultimate concepts are a different matter. If a concept C is ultimate but not fixed, then there are alternative concepts that can be used, but the question of whether C or some alternative concept is privileged cannot be asked in suitably independent terms. This is a real limitation to variance inquiry.

In the discussion of thin normative concepts we saw that there is a real threat that thin normative concepts are—in the terminology now introduced—ultimate. But even if this is so in the case of thin normative concepts, is it plausible that there are other instances of this phenomenon, and instances that are not immediately bound up with the problems having to do with thin normative concepts?

Compare a toy example. (With possible similarities to actual debates. But I want to discuss a simple made-up case, in abstraction from various complexities.) Suppose that one group of metaphysicians are fundamentally concerned with what is REAL and another group is fundamentally concerned with what is REAL*. The first group of metaphysicians think that in some deep sense there is nothing more to reality than what is REAL, and the other group think the same about what is REAL*. (Lots here is sketchy: what ‘deep sense’? and what does the ‘reality’ in small letters mean—REALITY, REALITY*, or something else? But I believe the sketchiness does not actually matter to the questions I am about to bring up.) Then the groups attempt to ask the question: is what is REAL or what is REAL* privileged for the purposes of metaphysical theorizing? There is a sense that we may be dealing with something ultimate: that when attempting to ask the relevant question of privilege one group will ask whether it is what is REAL or what is REAL* that is REAL, and the other group will ask which it is that is REAL*. Both questions are trivial; neither gets at an interesting underlying question of privilege.

It may be retorted that there is a way to get a suitably independent handle on whether REAL or REAL* is privileged: one need simply consider whether overall theories of the world which employ one concept is more theoretically virtuous—simpler, more explanatory, . . .—than the other. This is no different from choice of

ideology in theory construction quite generally. Of course, it could turn out that one cannot in fact use this method to choose ideology: REAL and REAL* score equally high. But again this would just be an instance of a familiar phenomenon: under-determination of theory by data.

But what if the REAL-users say: So what if a theory instead employing REAL* would be in these ways more theoretically virtuous? That theory, since it does not speak of what is REAL, does not state what the world is REALLY like, and it is a theory that states this we should aim for. If the REAL-users respond this way, REAL functions for them as an ultimate concept.

Acknowledgements

Many thanks to audiences at Cornell University, Trinity College Dublin, the Moral Sciences Club at the University of Cambridge, a workshop at University of Leeds, and two Conceptual engineering workshops, at University of Oslo and in Zermatt, for useful feedback. Special thanks to Tim Button, Andy Egan (who was my commentator in Zermatt), and David Plunkett for helpful comments.

References

- Brandt, Richard. 1979. *A Theory of the Good and the Right*. Oxford: Clarendon Press.
- Dreier, James. 2004. Meta-ethics and the Problem of Creeping Minimalism. *Philosophical Perspectives* 18:23–44.
- Eklund, Matti. 2011. Review of Eli Hirsch, *Quantifier Variance and Realism*. *Notre Dame Philosophical Reviews*. <http://ndpr.nd.edu/news/24764-quantifier-variance-and-realism-essays-in-metaontology/>
- Eklund, Matti. 2016. Carnap's Legacy for the Contemporary Metaontological Debate. In Stephan Blatti and Sandra Lapointe (eds.), *Ontology after Carnap*. Oxford: Oxford University Press.
- Eklund, Matti. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.
- Fine, Kit. 2001. The Question of Realism., *Philosophers' Imprint* 1.
- Hirsch, Eli. 2011. *Quantifier Variance and Realism*. New York: Oxford University Press.
- Horgan, Terence, and Timmons, Mark. 1992. Troubles for New Wave Moral Semantics: The 'Open Question Argument' Revived. *Philosophical Papers* 21:153–75.
- Horgan, Terence, and Timmons, Mark. 2009. Analytical Moral Functionalism Meets Moral Twin Earth. In Ian Ravenscroft (ed.), *Minds, Ethics and Conditionals* (pp. 221–37). Oxford: Oxford University Press.
- Lewis, David. 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy* 61:343–77.
- Lewis, David. 1984. Putnam's Paradox. *Australasian Journal of Philosophy* 62:221–36.
- McGee, Vann. 2006. A Rule for Everything. In Agustín Rayo and Gabriel Uzquiano (eds.), *Absolute Generality* (pp. 179–202). New York: Oxford University Press.
- Parfit, Derek. 1971. Personal Identity. *Philosophical Review* 80:3–27.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Railton, Peter. 1986. Moral Realism. *Philosophical Review* 95:163–207.
- Schaffer, Jonathan. 2009. On What Grounds What. In David Chalmers, David Manley, and Ryan Wasserman (eds.), *Metametaphysics* (pp. 347–83). Oxford: Oxford University Press.
- Scharp, Kevin. 2013. *Replacing Truth*. Oxford: Oxford University Press.

- Schwarz, Wolfgang. 2015. Analytic Functionalism. In Barry Loewer and Jonathan Schaffer (eds.), *A Companion to David Lewis* (pp. 504–18). Oxford: Blackwell.
- Sider, Theodore. 2009. Ontological Realism. In David Chalmers, David Manley, and Ryan Wasserman (eds.), *Metametaphysics* (pp. 384–423). Oxford: Oxford University Press.
- Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- Turner, Jason. 2010. Ontological Pluralism. *Journal of Philosophy* 107:5–34.
- Williamson, Timothy. 1988. Equivocation and Existence. *Proceedings of the Aristotelian Society* 88:109–27.
- Wright, Crispin. 1992. *Truth and Objectivity*. Cambridge, MA: Harvard University Press.