

Epistemic Injustice at the Conceptual Level: Are We Entitled to Our Own Concepts?

Cathal O'Madagain

Département d'Etudes Cognitives, Paris Sciences et Lettres

Epistemic injustice is the phenomenon whereby we commit an injustice against someone in their capacity as a knower. If we ignore someone's knowledge due to their membership in a particular group against which we are prejudiced, a kind of harm arises that is uniquely epistemic. The building blocks of knowledge and belief, of course, are concepts—and so conceptual injustice, if there were such a thing, might be a more serious problem again. Is there such a thing as conceptual injustice? Injustice at the conceptual level has been discussed before, but only with the assumption of “subjectivism” about concepts—that there is no right or wrong way to use a concept—and hence such that its consequences are restricted to problems of misinterpretation or an impoverished ability to express ourselves. However, recent accounts of concepts have shown that some conceptual schemes can be measurably better than others, and this carries the implication that the conceptual schemes recommended by those more knowledgeable than ourselves should be deferred to. Once we recognize that concepts can be *objectively* better or worse, the nature of the kind of epistemic injustice that can arise at the conceptual level is revealed to be very different. It becomes clear that we will often have a duty to defer to another's conceptual scheme, not merely for the sake of understanding what they are saying, but in our own subsequent thinking on the subject.

Keywords: concept utility, epistemic injustice, objectivity, conceptual schemes, Fricker

1. Some familiar kinds of epistemic injustice

Suppose you are an excellent mechanic, and you are having a conversation about the maintenance of sparkplugs at a cocktail party. After presenting your insights regarding their proper care with the right kind of oil, a celebrated doctor leans over and contradicts everything you have said. Something has clearly gone wrong here. As good at doctoring as this celebre may

Corresponding author's address: Cathal O'Madagain, 111 Rue Mouffetard, Paris 75005, France.
Email: cathalcom@gmail.com.

be, her knowledge need not translate into knowledge of sparkplugs, and her disregard of your higher epistemic standing in the domain of mechanics is clearly a failing on her part, and an unpleasant way to treat you.

Cases like this one are perhaps not too grave, albeit annoying. However, if the party who ignores or discredits the speaker does so because of the speaker's membership of a marginalized group, the offense becomes much more serious. If a white administrator in the Canadian government ignores the testimony of a First Nations member with respect to the harms committed against her community as part of the Residential Schools program, and does so precisely because of a prejudiced stance against First Nations people, then the discrediting of testimony has risen to a systematic level. Prejudice against this group has extended into a systematic discrediting of their knowledge, and has become an "epistemic injustice". The harms incurred by the victim of epistemic injustice can be indirect, in this case by lowering the chances of crimes committed in the Residential Schools program being redressed. They can also be direct: our knowledge is hard-won, and ought to command respect of others. Whenever we ignore someone else's knowledge, we inevitably place ourselves on a pedestal above them as more mature, wiser, or more insightful. Arbitrarily denying these "virtues" of another in itself amounts to committing a harm against them—doing so systematically on the basis of their membership in a particular group is an epistemic injustice.

There are likely many forms of epistemic injustice. Miranda Fricker's (2007) ground-breaking study of this phenomenon focuses on two clear cases. The first is what she calls testimonial injustice, where we fail to sufficiently heed the knowledge imparted in another's testimony. In the case of testimonial injustice, the harm comes about due to failing to properly heed the knowledge of another. Epistemic injustice typically discussed at this level concerns failing to defer to the beliefs of another party—propositions they have expressed for us in conversation. But at a lower level again in the architecture of propositions, beliefs and knowledge are to be found our concepts. If the mechanic's beliefs about how to take care of a sparkplug using oil constitute knowledge, they do so only by dint of his grasp of concepts including SPARKPLUG, CARE, and OIL, which concepts are the component parts of those beliefs. If injustice can be committed against someone at the level of their beliefs or knowledge, then surely an injustice committed against someone at the level of their concepts would be graver still, since our concepts are the component parts of our beliefs. If we have failed to appropriately defer to others in our use of concepts, then all beliefs involving those concepts

will be affected; whereas we could fail to adopt one specific belief without it affecting any others (*pace* Quine 1951).¹

Injustice that arises at the level of concepts has so far been explored under what Fricker calls “hermeneutic injustice”. This kind of injustice can arise when we lack the concepts required to articulate an injustice committed against us—for example when there was no widespread concept of POST-NATAL DEPRESSION, so that women experiencing this syndrome lacked an appropriate concept to articulate the cause of their problem, or even understand what they were going through, compounding the difficulty they are faced with. Since the lack of a medical concept to capture these symptoms was due itself to the marginalization of women, the lack of a concept and the wrongs that result has itself become an epistemic injustice. It can also occur when we fail to acknowledge the way another is using a concept. Anderson (2017) describes a black female graduate student using a concept “natural kind” in one way, while a white male graduate student has a different grasp of the concept. Rather than assuming that he has a mistaken understanding of the concept himself, or adopting her use of the expression for the sake of clearly understanding what she is saying, he rejects her as incompetent, and that on the basis of his prejudice against women and black people. In this case, the injustice is committed as a result of failing to defer for the sake of a clear interpretation to the other person’s use of an expression.

“Hermeneutic” injustices, or injustices of interpretation, are generally discussed under the assumption that concepts are the kinds of things whose meaning is really a matter of personal use. If there is no objectively right or wrong way to use a concept, and any harm that can arise here arises due to misunderstanding or an inability to express ourselves. We could call this an assumption of “subjectivism” about concepts or conceptual schemes.² But what if there is an objectively better or worse way to use a conceptual

¹ Of course, failing to defer to someone who is telling us about a better way to use a concept will be failing to defer to their “testimony”. As a result, the kind of epistemic injustice I am going to explore could be subsumed under testimonial injustice. But the difference between failing to defer to a belief and failing to defer to a conceptual scheme seems substantial enough to me to warrant a separate treatment.

² See for example Anderson (2017, 3): “There is no way to determine a person’s credibility regarding purely conceptual or linguistic matters on the basis of objective empirical evidence. When someone says ‘Natural kind terms are not rigid designators,’ or ‘Privacy is a right,’ or ‘Racism entails the existence of systematic oppression,’ there is no presently known way to use objective empirical evidence to check whether that person is credible”, and later “for present purposes, a conceptual claim is simply a claim for which there is no agreed upon method for deciding its truth”. In some places Dotson (2012) seems to suggest that failing to defer to another’s conceptual scheme seems to amount to failing to acquire knowledge, and hence that there are better or worse conceptual schemes; if this is Dotson’s view then her position is closer to the view I am advocating here, although she does not

scheme? If concepts are the kinds of things that we can be objectively wrong about, then we would have a duty to defer to the more knowledgeable party's use of a concept—but not just for the sake of temporarily interpreting what they are saying. Rather, in our own subsequent thinking about the topic the concept pertains to, we ought to adopt the conceptual framework of the more knowledgeable party. And failing to do this because of a prejudiced stance against the group the knowledgeable party belongs to would form a distinctive kind of epistemic injustice.

To put this another way, if we assume subjectivism about concepts, we leave a very specific kind of *excuse* on the table for the epistemic wrongdoer. Even having recognized the conceptual scheme through which another views the world and expresses their testimony, the epistemic wrongdoer can say “I have heard what you said, and I understand what you mean by those words; but, that is just not how I think about things”. This can result in various harms to the person or group whose conceptual scheme I refuse to embrace. For example, suppose someone tells me that the concept RACISM covers positive generalizations about a disenfranchised minority just as much as negative generalizations. For the sake of simply understanding what they are saying, I might interpret their utterances with this in mind. But now suppose that once I have left the context of this particular conversation, I refuse to consider such generalizations as racist. If I have a general belief like “racism is bad”, this will not necessarily stop me from making these positive generalizations I have been warned about—which are clearly harmful to the community I make those generalizations about. But the harms incurred are also direct—if conceptual schemes can be better or worse, then by ignoring another's conceptual scheme, we harm the other as a knower, just as we do if we ignore their testimony in other ways.

In this essay I argue that some conceptual schemes are, indeed, objectively better than others, and that failing to fully adopt the conceptual scheme of another can amount to a serious form of epistemic injustice. But how could concepts be the kind of thing we can get “wrong”? To see how this could be the case we first need to explore the nature of concepts in more depth.

2. Can one conceptual scheme be better than another?

Depending on your background, it may seem either intuitive, or absurd, that one person could have a better conceptual scheme than another. To begin with the intuitive approach, consider the following scenario adapted from

explain how a conceptual scheme could in fact be better or worse, which is part of my current focus.

Burge (1979). A mechanic arrives at the doctor with a headache, and tells his doctor “I have terrible rheumatism in my head”. His doctor reassures him that she probably can help with the pain, but advises him all the same that whatever the cause, it cannot be rheumatism. “That is simply not what the concept RHEUMATISM means”, she insists, “rheumatism is a condition that affects your joints”. Under ordinary circumstances, the patient will accept the doctor’s authority here. He might complain that he intended to refer to pain in his head, and he does not really care what “rheumatism” means; but he will likely accept the doctor’s authority on the concept RHEUMATISM all the same (displaying greater epistemic humility than the doctor did before). In this scenario, we might say, the mechanic has deferred to the doctor on the concept rheumatism. We might also say, the mechanic has recognize that he failed to grasp the concept, while the doctor properly grasped it. The mechanic got it wrong, the doctor got it right.

On some views of concepts, however, such a transaction is mysterious. On some views, after all, concepts are merely labels that we use for objects or collections of objects in our environment. The concepts themselves carry no claims about the nature of the objects they pick out, it is only our beliefs *about* the objects that fall under those concepts that do the work of telling us something about them. On this view, it does not make much sense to suppose that we could conceptualize things in the wrong way, any more than the names in one language could be considered more correct than those in another. As a result, it does not make sense to suppose that the doctor’s concept of RHEUMATISM could be more correct than the mechanic’s—rather, the two are simply speaking different languages. Fodor (1994), for example, defending such an account of concepts, argues that it is a fundamental confusion of semantics (the domain of meaning) with epistemology (the domain of beliefs and knowledge) to suppose that anyone could have a better set of concepts than anyone else. On this view, possessing a particular set of concepts amounts to assigning a semantics to the particular “language of thought” that you use to think about the world, and no one language can be objectively better than any other. But this implies that some heated contemporary debates, such as the debate over what counts as falling under the concept RACISM (see Sanneh 2019), are vacuous. There is no right or wrong answer to this question, on this view of concepts. It is as a result not possible to commit an injustice by failing to defer to another’s authority on the question what counts as RACISM, since no-one could have more well-informed opinion than anyone else. Similarly, such an approach to concepts leaves the “bad excuse” considered earlier on the table. If any one conceptual scheme or “language of thought” (Fodor 1994) is as good as any other, then although I can wrong someone by failing to interpret their utterances in accordance

with the conceptual scheme they use to express themselves (a hermeneutic injustice), I cannot wrong someone by failing to adopt their conceptual scheme in my own subsequent thinking—even if they know far more about the domain under consideration than I do.

Apart from rendering the debate over a concept like RACISM vacuous, this Fodorian position renders many scientific debates about classification vacuous. Whales used to be categorized as fish, and this was convincing for many expert in the nature of whales at that time—“a whale is a fish with a horizontal tail”, insists Ishmael (Melville 1851). But we have since revised the concept WHALE as falling under the concept MAMMAL rather than FISH, and thereby revised both of these latter concepts at the same time (in the case of MAMMAL, now extending to cover WHALE where it did not before; and in the case of FISH, now failing to extend to WHALE, where before it did). And it is surely the case that it is not just a change in fashion that has led to this reclassification, but an improvement in our understanding—surely, we think, we were wrong to classify the whale as a fish, and are right to now classify it as a whale. How can we make sense of this? What did we get wrong before that now we have right?

Égré and O'Madagain (2019) provide an account of the “epistemic utility” of concepts that may make sense of these claims. The account there focuses on the recent dispute over the term “planet”, or we can say, over the proper meaning of the concept PLANET. I will explore this case in some detail, and as we shall see, once we have such an account in place, it becomes clear what is wrong with the “bad excuse”, and the duty to defer more completely to another's conceptual scheme will become apparent.

Properly grasping this dispute involves a little historical groundwork. In 2006, Michael Brown discovered what was widely hailed as “the 10th planet”, which he called Eris. But he shortly thereafter decided that in fact what he had discovered ought not to be considered a planet at all. At the same time, he came to believe that if Eris ought not to be considered a planet, then Pluto, one of the standing family of planets, ought not to be considered a planet either—for Pluto was much more like Eris than either of these were like the other bodies referred to as planets. In essence, Brown started a debate over the proper extension of the concept PLANET.

Although some commentators held that the dispute over the meaning of planet was a mere terminological switch, arbitrarily decided by convenience, such that the old classification could not be “less correct” than the new one (Chalmers 2011), Brown himself insisted that this was no mere terminological dispute:

The debate about whether or not Pluto is a planet is critical to our understanding of the solar system. It is not semantics. It is fundamental classification. (Brown 2010, 232)

Brown's opinion is that one classification of the objects in the solar system turned out to be objectively incorrect, and that its flaws were made salient upon the discovery of new information about the solar system. Brown and many others had already suspected the classification of Pluto as a planet to be dubious. Pluto has an orbit around the sun of 17 degrees relative to the other planets, which all orbit on the same plane. Pluto is not largely spherical, but is shaped more like a potato. Pluto does not have a gaseous atmosphere, but is instead cloaked with ice. And finally, Pluto has not cleared its orbit of other objects—Pluto orbits the sun in a ring that includes discernible large rocks and debris. With the discovery of Eris, a second body that shared all of these notable properties that distinguish them from the other planets was discovered, and this had the effect of intuitively creating a new category of objects distinct from the original planets. This intuitive distinction was made more compelling still by the discovery shortly after of a whole range of bodies in the same region of the solar system, that all share these distinctive features. Now, rather than a category with a single “oddball” member, being Pluto, it seemed to many that we had instead two categories—the original 8 planets, and the new objects which for want of a neater title came to be called “Trans-Neptunian Objects”. In a decision of the International Astronomical Union, a new definition of PLANET was adopted, including one particularly salient feature that set the original 8 planets clearly apart from this growing category, namely that of having cleared its orbit of smaller objects and debris:

A planet is a celestial body that:
 (a) is in orbit around the Sun,
 (b) has sufficient mass for its self-gravity to overcome rigid body forces so that it assumes a hydrostatic equilibrium (nearly round) shape, and
 (c) has cleared the neighborhood around its orbit. (General Assembly of the International Astronomical Union 2006)

This had the effect of officially restricting the concept planet to those 8 objects in the solar system that had cleared their orbit of smaller rocks and debris, and placing Pluto, Eris, and several other objects into the category of TNO. But what is it, if anything, that makes the new classification *more correct* than the old one?

This is where an account of concept utility becomes useful. Égré and O'Madagain argue that a better conceptual scheme yields better beliefs—not as a matter of subjective preference, then, but in measurable objective terms. The account appeals to two features of beliefs that have been widely recognized as contributing to their goodness or “utility”. On the one hand,

we want the beliefs we form to be plausible, and ideally true. If we see a bluebird lay a blue egg, we might infer “bluebirds lay blue eggs”. This seems like a sound inference. We would not on the other hand be tempted to infer from the same observation that “all birds lay blue eggs”. We may after all have seen eggs of other colours lying about, so that we are pretty sure that such a general claim would be false. And false beliefs are unhelpful for obvious reasons.

We might be inclined to think from this that our goal is always to adopt only those beliefs that are maximally likely to be true, or are maximally plausible. But as James (2016), Popper (2005) and others have observed, this cannot be the case. Consider the judgment we made about bluebirds above. If our goal was always to maximize truth in our beliefs, we could restrict our claim to a more narrow conjecture than “bluebirds lay blue eggs”—perhaps to “bluebirds in my garden lay blue eggs”, or even, to be absolutely sure, simply “the bluebird I just saw laying a blue egg laid a blue egg”. Given the observation, such a belief is guaranteed to be true. Of course, we do not exercise such extreme caution in our beliefs, because more than just reliability is valuable to us in our beliefs. We also want our beliefs to be informative. If reliability were the only thing we cared about, then we could reduce the strength of our beliefs so that they were guaranteed to be true—perhaps by believing only tautologies. But such a set of beliefs would not tell us anything about the world. What we are after are beliefs that somehow strike a balance between informativeness and plausibility, since it is important to us that both of these features are present in our beliefs.

With these considerations in hand, we may note that concepts, the component parts of our beliefs, have two features that systematically covary with the plausibility and informativeness of beliefs. First, our concepts can be more or less *homogeneous*, so that the objects falling under a concept are similar to one another. This will affect the plausibility of inductive generalizations we form using those concepts. Consider the concept BLUEBIRD. If the observation of one bluebird laying a blue egg leads us to make the generalization “bluebirds lay blue eggs”, the plausibility of this generalization will depend in part on how homogeneous the concept BLUEBIRD is. Notice that if we make the generalization over a more diverse group, for example “birds in my garden lay blue eggs”, where “BIRDS-IN-MY-GARDEN” will now include robins, crows, thrushes and so on, the plausibility of the inference immediately drops. When a generalization is made over a bunch of things from an observation of one member of the bunch, the generalization gets less plausible as the diversity of the bunch goes up. The homogeneity of a concept will therefore clearly effect the plausibility of the generalizations the concept is

used to make.³ One property we would like our concepts to have, therefore, is homogeneity.

Secondly, the number of things a concept extends over—which we call the *inclusiveness* of the concept—will determine the *informativeness* of beliefs the concept is used to make.⁴ If we expand the domain covered by the concept (in the subject position) in the generalization, as before, to include not just bluebirds but “all birds”, then since ALL-BIRDS includes bluebirds and many more birds beside, the generalization is more informative than the generalization extended to just bluebirds. But, of course, since increasing the inclusiveness in this case decreases the homogeneity of the concept, the generalization has thereby become less plausible. What is becoming clear is that we want our concepts to strike some sort of a balance between inclusiveness and homogeneity. If they include too many objects (e.g. “birds”), the concept will become highly heterogeneous, and not very reliable for forming generalizations; but if we focus on maximizing the homogeneity, for example by only operating with extremely specific concepts like “the bird I saw in my garden yesterday”, we reduce the inclusiveness of our concepts so that we cannot make *informative* generalizations or predictions about our environment. When categorizing objects in our environment, we want to find a conceptual scheme that strikes a balance between the two.

We can illustrate how to strike this balance easily. First consider how the homogeneity of a conceptual scheme might change. The way we conceptualize our solar system (leaving aside for now the complicating case of Pluto), is largely into four concepts: PLANET, MOON, ASTEROID, and STAR. These concepts are reasonably homogeneous, in the sense that the objects that fall under each one are more similar to each other than they are to any

³ We assign a value of homogeneity to a concept as the number of properties the objects falling under the concept have in common, divided by the number of properties they can be distinguished by, assuming a fixed set of discernible properties in the domain which are determined by the perceptual mechanisms of the observer (e.g. has wings, a beak, lays blue eggs). This assumes a basic set of concepts that an observer is likely to operate with, largely focusing on perceptual properties. We ignore the “philosophical-puzzle” type properties such as Cambridge properties or grue-like properties (cf. Goodman 1983).

⁴ At least, when the concept is in the subject position of an ampliative inference or (informative) universal generalization in the form “All As are Bs”. When the concept is in the predicate position in such a claim, the claim becomes more informative when the concept is *less* inclusive—or, more *exclusive*. Compare “All those things are birds”, with “All those things are Grey Wagtails”. The latter tells you more than the former. Nevertheless, a conceptual scheme that maximizes the inclusiveness and homogeneity of concepts overall will also feature very narrow and very homogeneous concepts like “grey wagtail” at lower levels. A conceptual scheme that maximizes homogeneity of concepts throughout will therefore permit both maximally informative and reliable generalizations both when concepts taken from that scheme are in the subject and predicate position.

objects falling under any other concept. But we could of course reclassify things, adopting instead four alternative concepts: Moon-or-Large Planet, Asteroid-or-Star, Small Planet-other-than Earth, and Earth. This overall conceptual scheme has just the same inclusiveness as the concepts we use, since the four concepts between them cover all the same objects. However, the concepts themselves are highly heterogeneous. An observation about one member of the concept Asteroid-or-Star, will clearly not support a plausible generalization to all members of the category. If we discover that asteroids are covered in ice, we should not expect to generalize this inference to all members of the category (the only other member being the Sun). And if we discover that one member of the category "Moon-or-Large Planet" is inclined to have its own moons, we should not expect this of other members of the same category. Clearly, simply for the sake of supporting plausible generalizations, the standard conceptualization of the objects in the solar system is better than this one.

Inclusiveness needs also to be taken into consideration in determining a conceptual scheme. When Pluto was discovered first, as noted above, it was already recognized as being an "oddball" among the planets. Nevertheless, it was grouped initially with the others. Now consider that if all we were concerned with in forming conceptual groups was maximizing their homogeneity, and thereby raising the chances of making plausible generalizations over the category based on observations of their members, then surely we should have immediately assigned Pluto to its own special category at the outset. We would then have had as our conceptual scheme of the solar system, the concepts STAR, PLANET, MOON, ASTEROID, and PLUTO. Why was there no temptation to do this? The reason is that not only do we value homogeneity in our concepts, but we also value inclusiveness. A concept with just one member does not support inductive inferences from observations of that member to other objects—since there is only one object in the category. This is not to say that there can never be any use for concepts with just a single object, there can—but if an object is not too dissimilar to other objects in an environment, then grouping it with them allows us to retain a high degree of inclusiveness in our concepts. Once Eris was discovered some 70 years later, another object in the same region of the solar system quite like Pluto, and then many other objects besides in that region, we now found ourselves with a reason to distinguish Pluto from the other planets. We would no longer be assigning Pluto to a category with very low inclusiveness, but instead we would be assigning it to a category that would soon have as many, indeed many more members than the category PLANET. And this was why (according to Égré and O'Madagain), with the discovery of these objects, a new category called "TRANS-NEPTUNIAN-OBJECT" was invented, and Pluto

was moved from the category PLANET to this new category. The category PLANET now increased in homogeneity, since the remaining members were more similar to one another than any was to Pluto, and Pluto joined the category of TNO, where it had more in common with its new fellows than it had with the Planets.

What we can see from these considerations is that there is, after all, a straightforward sense in which it is possible to have an objectively better or worse conceptual framework (for a full treatment see Égré and O'Madagain 2019). Even if concepts do not individually do anything other than to label a set of objects in our environment, when viewed as part of a conceptual scheme it becomes clear that some sets of concepts are epistemically better than others. As a result, it becomes clear that if someone else knows more about a given domain than we do ourselves, then their classification of that environment will likely be better than our own. If the astronomers tell us that Pluto should not be grouped with the other planets, then we should defer to their judgment—it is not simply a matter of their convenience, which could of course be different for our own convenience. Rather it is likely that, as a matter of fact, they have found a conceptual scheme with a higher epistemic utility—one that is likely to yield more plausible and informative generalizations overall.

Similarly, if the doctor tells the patient that headaches should not be grouped together with joint aches under the concept RHEUMATISM, this is probably because the doctor knows enough about the overall landscape of ailments that shows her that a conceptual scheme that groups these two aches together will have a lower epistemic utility (although she did not likely think of it in these terms) than one that classifies them distinctly. And if the mechanic tells the doctor that sparkplugs should not be thought of under the concept engine-part, this is likely because the mechanic knows something about the landscape of artefacts that make up the assembly of a car such that classifying spark plug in this way will lead to poor generalizations. It is not a matter of convenience, and neither is it a consequence of our simply speaking different “languages of thought”: some people in a community can know more than others about how best to conceptualize a given domain.

Finally, although I would not try to decide here what will optimize a conceptual scheme that includes concepts like RACISM and SEXISM, since I am no expert on these issues, I wager it is a safe bet to suppose they can indeed be measured in the same ways, such that one cannot arbitrarily alter their extension without damaging the utility of the beliefs we form using these concepts. Imagine someone who insists on using the concept “RACISM” to include the kinds of behaviour that we are familiar with under this concept, plus kite-flying. If, when we try to identify the extension of RACISM, we

are trying to identify a category in a taxonomy of human behaviours that will allow us to form a healthy epistemic outlook in this domain (i.e. reliable and informative beliefs), then it is clear that generalizations made using this concept of RACISM will be worse than they would be were *kite flying* excluded from the concept. There is a fact of the matter about what the optimal extension of the concept will be, and getting it right goes beyond personal preference or conventional coordination.

Perhaps it might be worried that I am placing too much emphasis on the epistemic role of concepts here, and that there could be other roles for concepts that could over-ride their optimal use for forming beliefs. For example, the US government at one point tried to change the meaning of the concept TORTURE to exclude water-boarding, in order that it could avoid being accused of violating the UN conventions against torture (Sundell 2011). I am inclined to insist, agreeing with Simion (2017) that the epistemic role of concepts will always be more important than other roles—so that optimizing the epistemic utility of concepts should always over-ride other considerations. Although it is true that concepts play different roles, some of these roles are more fundamental than others, and the epistemic role is the most fundamental. While we can adopt a concept with a particular extension in order to get out of being considered racist or in violation of an international convention, such considerations will not *always* be affected by our concepts—there are lots of other ways we could get out of such difficulties, and there are lots of ways we could use concepts that would have no bearing on these difficulties. Our beliefs, on the other hand, will always be affected by our concepts. Any deployment of a concept will result in the formation of a belief. The epistemic utility of a concept should therefore be prioritized over any other considerations—since it is the “primary purpose” of a concept, insofar as the function of a concept is to allow us to form beliefs, and the goal of belief-formation is to have reliable and informative beliefs. And even if other considerations sometimes do arise in our adjustment of our conceptual schemes, epistemic utility should act as a “limiting consideration” (Simion 2017) for any further choices we might make as to how to deploy a concept.

3. A duty to defer

So far the examples considered are relatively trivial. Even if we were to ignore the position of the IAU on the concept PLANET, the upshot might simply be some poor inferences about planets. However, such failures can easily rise to the level of epistemic injustice. One debate that serves as a clear example is the debate over the question what falls under the concept RACISM. If the way we use term “racist” were simply a matter of coordinating on using a

concept in the same way, so that we could communicate fluently, then each party to the debate should surely contribute equally (as Lewis (1979) imagines us coordinating, game-theoretically, on the most widespread meaning for a word in a linguistic community). Hence, groups that have historically contributed to the oppression of others, such as men in the case of sexism, or whites when it comes to racism, should presumably take their own opinion of what counts as “racist” as seriously as the opinion of those who have been victims of these attitudes—since we all use the words. And, it often seems to be the case that exactly this concern motivates many people when confronted with a charge of being racist or sexist, to counter the charge by offering their own account of what falls under those concepts—consider the perhaps familiar refrain “well, what I mean by racism is something like a lynch mob—so I cannot accept that my talking over you is an instance of racism, sorry...”. This is, of course, the “bad excuse” considered at the outset, that subjectivism about concepts leaves on the table. If the way a concept is to be used is simply a matter of convenience or individual preference, and there is no right or wrong answer about what counts as the extension of a concept, then it is hard to say what is wrong with such a reply. Equally, if there is no right or wrong way to use the concept, then the debate over what counts as RACISM is rendered vacuous.

As mentioned at the outset, epistemic injustice that assumes this kind of subjectivism about concepts has already been extensively explored. Anderson (2017) has explored what he calls “conceptual competence injustice”. On Anderson’s view, this kind of injustice is committed against another person by assuming s/he lacks competence with some concepts s/he is working with simply because s/he belongs to a marginalized community. Anderson considers a white male graduate student ignoring the opinion of a black female graduate student, because he assumes she is using the concept “natural kind” in the wrong way. In fact she is using the term following Soames while he is using the term following Kripke. Instead of assuming either that he has misunderstood the concept himself, or that they are simply using the concept in different ways, he assumes that she is incompetent, and that because of her race and gender.

Similarly, Fricker describes “hermeneutic” injustices—where the lack of a concept, or a failure to grasp a concept, will result in deliberate misunderstanding or mischaracterization of someone else’s view, or an inability to articulate your own. And Dotson (2012) considers “contributory injustice”, where we fail to allow someone else to use a concept in the way most suitable for their purposes, in the context of their discussion. The guiding idea through these approaches to conceptual injustice is that it will result

in wrongs concerning misinterpretation or an inability to express ourselves. This is why Fricker characterizes such cases as “hermeneutic”.

Hermeneutical injustice is of an entirely different species, however, to the kind of injustice that can arise once we acknowledge the possibility of objectivism about concepts. If someone else merely speaking a different language of thought by using concepts in a different way to the way we use them ourselves, we certainly have a duty to capitulate to their use of words in hearing them out. But we do not seem to have a duty to adopt their conceptual scheme ourselves. We have every right to continue using our own concepts the way we have been—indeed, it would seem that at, pains of incoherence, we have just as much a right to use concepts the way we are used to, perhaps to articulate and defend our own world-view and grievances, as we have a requirement to acknowledge a different way of using those words, when they are used differently.

If, on the other hand, some people are using a concept in a way that is more correct than others, then it would be wrong to fail to defer to their use of concepts in our own thought and speech—we will have at least an epistemic duty to defer. Not just temporarily, as Dotson considers, for the sake of understanding someone—but also in our own use of the concept, which we will have discovered to be, simply put, incorrect. If we fail to defer to someone because of their membership in a group that we are inclined to disparage, then such a failure will be an epistemic injustice.

The considerations raised in this essay indicate that the view on which everyone's opinion about what counts as racist or sexist is equally valid (*à la* Fodor), or even that it might depend on some kind of implicit convention or coordination between speakers (*à la* Lewis) is indeed wrong.⁵ How to conceptualize a domain is not merely a matter of opinion or convenience, and some conceptualizations of a domain can be objectively better than others. Not only that, but those most knowledgeable of a particular domain should be *deferred to* in their conceptualization. In that case, more weight should be assigned to the opinion of the group with the greater knowledge of the subject matter. This knowledge can be gained through careful study, so that experts will know more than novices; but it can also be acquired through experience, with some knowledge depending on having appropriate experience (Paul 2014). Experience with the subject matter of racism is presumably experience with being subjected to racist behaviour. Those who have been

⁵ Or at least that it is wrong *in general*. There could be specific cases for which nothing more than coordination is required—but for some straightforward cases like natural kinds, I think the view is certainly mistaken, and indeed I am inclined to think that the kind of case where meaning comes down to nothing more than coordination for the sake of communication are marginal (cf. King 2012).

routinely subjected to racist behaviour will have greater sensitivity to the contours of the phenomenon, will spot it more easily, and will have a greater grip on the landscape of its varieties, subcategories, and super-categories. At least if we are to take attempts to explore the concept of RACISM seriously (e.g. Kendi 2016), we should defer to those with the greatest amount of experience in the same way we will defer to the International Astronomical Union when we heed their advice with respect to a concept like PLANET.

A similar conclusion applies to a concept like SEXIST. Since women have been subjected to sexism while men generally have not, then women have more experience with the landscape of varieties of sexism than men do. Again, they will be more familiar with the phenomenon, and quicker to spot it. As a result, women are in a better position to categorize some behaviour as sexist than men. By dint of unfortunate experience, they know more about it. Although men are often quick to offer their own opinion with respect to what does or does not count as sexist, we should really consider such counsel in the same way that we might consider the council of the doctor to the mechanic when she advises him on what should really be classed as RHEUMATISM, or the mechanic's advice to the doctor on what counts as a SPARKPLUG. When a man refuses to defer to a woman's conception of sexism precisely because he has a dismissive attitude toward gender, he commits a form of epistemic injustice. And again this deference cannot be temporary—it is not simply for the sake of understanding what she is saying. Rather, the deference should be permanent—adopting the concept of SEXISM recommended by the more knowledgeable party for the sake of one's own thinking about the world, rather than merely for interpretation of someone else's beliefs. The “bad excuse” that subjectivism about concepts leaves on the table—“well that is just how I see things, I have my own conceptual framework after all”—is not available once we recognize that conceptual frameworks can be objectively better or worse.

4. Conclusion

The injustice of ignoring someone who knows more than you do, particularly because that person is a member of some marginalized group, is increasingly widely recognized as an important issue at the crossroads of epistemology and ethics. But one form of this kind of injustice, where we fail to defer to someone's *conceptual scheme* even when they understand the landscape that this conceptual scheme carves up better than we do ourselves, is not only underexplored, but will actually be considered incoherent on many standard views of concepts. Where it has been explored, the most that has been thought to be morally required is that we temporarily defer to another's conceptual scheme in order to properly understand what they are saying.

But once we realize that different conceptual schemes can be objectively better or worse, however, and indeed that their competing epistemic values can be measured, we can see how failing to defer to a better informed conceptual scheme—in the sense of adopting that scheme for ourselves rather than simply tolerating it for the sake of a conversation—will constitute a centrally important form of epistemic injustice. An account of concept utility like that discussed here can be used to show how deference to the conceptual scheme of another can be not only useful, but that failing to defer can be a form of injustice. And it shows how the scale of deference required is much more substantial than has been considered explored under the kind of conceptual injustice explored so far.

Bibliography

- Anderson, D. (2017). Conceptual competence injustice, *Social Epistemology* **31**: 210–223.
- Brown, M. (2010). *How I Killed Pluto and Why it Had it Coming*, Spiegel and Grau.
- Burge, T. (1979). Individualism and the mental, *Midwestern Studies in Philosophy* **4**: 73–122.
- Chalmers, D. (2011). Verbal disputes, *Philosophical Review* **120**: 515–566.
- Dotson, K. (2012). A cautionary tale: on limiting epistemic oppression, *Frontiers: A Journal of Women Studies* **33**: 24–27.
- Égré, P. and O'Madagain, C. (2019). Concept utility, *The Journal of Philosophy*.
- Fricker, M. (2007). *Epistemic Injustice*, Oxford University Press.
- General Assembly of the International Astronomical Union (2006). Resolution B5.
https://www.iau.org/static/resolutions/Resolution_GA26-5-6.pdf
- Goodman, N. (1983). *Fact, Fiction, and Forecast*, 4th edn, Harvard University Press, Cambridge, Mass.
- James, W. (2016). *The Principles of Psychology*, Macmillan, London.
- Kendi, I. X. (2016). *Stamped from the Beginning*, Nation Books.
- Lewis, D. (1979). *Convention*, Blackwell.
- Paul, L. A. (2014). *Transformative Experience*, Oxford University Press.
- Popper, K. (2005). *The Logic of Scientific Discovery*, Routledge.

- Quine, W. V. O. (1951). Two dogmas of empiricism, *The Philosophical Review* **60**: 20–43.
- Sanneh, K. (2019). The fight to redefine racism, *The New Yorker* **August**.
- Simion, M. (2017). The ‘should’ in conceptual engineering, *Inquiry* **61**: 914–928.
- Sundell, T. (2011). Disagreement about taste, *Philosophical Studies* **155**: 267–288.