Having a Concept Has a Cost

to appear in *Synthese*

Michael Deigan

Freie Universität Berlin

> . . . all this vivid sympathetic experience returned to her now as a power: it asserted itself as acquired knowledge asserts itself and will not let us see as we saw in the day of our ignorance.
>
> *Middlemarch*
> George Eliot

Having a concept usually has some epistemic benefits. To begin with, concept possession enables knowledge. Without concepts like CAR and NEW YORK CITY, I could not know that there are cars in New York City and so would miss out on the value of knowing facts like this one. On a relatively narrow conception of epistemic value, this may be all there is to it, but more expansive conceptions of epistemic value will complicate the picture. If we take epistemic value to be sensitive to naturalness, for instance, possession of concepts may differ in epistemic value depending on how joint-carving they are. Even if the GRUE/BLEEN and GREEN/BLUE pairs in principle allow us to know the same facts, perhaps there is still more value to possessing the latter.[1]

On another more expansive view, there is epistemic value in an empathetic

[*]For very helpful discussions and comments, thanks to Keith DeRose, Daniel Ferguson, Daniel Greco, Laurie Paul, Zoltán Gendler Szabó, to my commentators at the 2020 Eastern APA, Jens Kipper and Annette Martín, to our audience, Karen Bennett, and to several anonymous reviewers. I am also grateful for funding from the Alexander von Humboldt Foundation, which supported some of the work on this article.

[1]See, for example, Sider (2012) for a view along these lines, as well as Pérez Carballo (2020) for criticism and Egré and O'Madagain (2019) for an alternative that tries to make sense of some of the same intuitions.

understanding of other people. That is, there is epistemic value in a kind of understanding of people 'from the inside', on their own terms, which would not be captured by a third-personal causal explanation, even one that included all the facts under some description or other.[2] My aim in this paper is to explore what consequences this would have for the epistemic value of concept possession. Let us, then, assume such a view is true.

One consequence is an additional benefit from concept possession. Possessing concepts can allow for a particularly direct way of understanding the thoughts of others. You too have the concepts CAR, NEW YORK CITY, and so on, so you can know that I believe that there are cars in New York City through grasping the same content in at least approximately the same way that I do. Even if an alien super-scientist could somehow know all the facts about me at some level without such concepts, they wouldn't be understanding my belief in an empathetic way. On the view we are assuming, then, they would be missing out on some epistemic value. This benefit of concept possession is an interesting one partly because, as Jennifer Carr observes, it will extend even to concepts that are otherwise defective and useless (Carr 2015, p. 224). What I want to explore, though, is in the opposite direction: not the benefits of concept possession, but the costs.

Can possessing a concept have an epistemic cost? At least when we idealize away from our computational limitations, it is tempting to think that it cannot. Carr, in defending (though not officially committing to) the idea that it is always rationally impermissible to lose conceptual resources, says of defective concepts associated with slurs that "[w]hat's really wrong with objectionable concepts is not possessing them, but rather (in some sense) applying them" (Carr 2015, p. 224). And here is Hofweber with a similar line of thought: "Getting rid of a concept from one's repertoire of

[2]Recently Grimm (2016), Ismael (2018), Hannon (2020), and Bailey (2022) have each developed views along these lines.

concepts would always limit one's representational capacities, and thus is arguably always irrational. What matters . . . is whether we should stop applying one concept in certain situations, and apply a different one in those situations instead" (Hofweber 2023, p. 9). So concepts can be misapplied, but on this view there are no drawbacks to merely possessing them, contingent space and processing limitations aside. In the worst case, a concept is just useless and should never be applied. For ideal agents, there would be no epistemic cost to possessing any concept. Let us call this view *Costlessness*.

If Costlessness were true, a normative epistemology of concepts will be trivial in a way normative epistemology of belief is not. No matter how ideal the agent, it is epistemically bad in some respect for them to have a false (or unjustified) belief. In their pursuit of epistemic goods, they must take care to avoid some epistemic bads, whether on broadly epistemic consequentialist, deontological, virtue theoretic, or some other grounds; not all beliefs are permitted. But, given Costlessness, it would seem that an ideal agent would be permitted to have any concept whatsoever, no matter how gerrymandered, confused, or pointless. After all, why not? For non-ideal agents like ourselves there will be considerations of clutter avoidance to take into account (Harman 1986; Friedman 2018). But these do not apply to the ideal agents who lack our contingent limitations. So if Costlessness is true, then an ideal epistemology of concepts looks rather uninteresting, in contrast with an ideal epistemology of beliefs.[3]

I will be arguing that given the value of epistemic empathy, Costlessness is false. The argument applies not just in some fringe cases, but quite generally. Possession of concepts prevents us from seeing as those who lack those concepts see. This is an

[3]There are more general questions about why ideal normative theory is supposed to matter. I will take for granted that it does. For two recent views on the role of ideal epistemology, see Staffel (2020) and Carr (2021).

epistemic cost that would apply to agents no matter how ideal.

I begin in §1 with an illustrative example to give a sense of where the argument will be going. §2 lays the foundations for the argument by developing some ideas about concept possession and identity, answering an argument from Ball (2009) that would make concept sharing relatively trivial. §3 presents the argument against Costlessness and defends its premises. §4 raises and replies to a few objections. §5 concludes by suggesting some potential upshots of the argument.

1. Example: The Trouble with Understanding Kids

Children think of the world in ways we adults have trouble fully understanding. We distinguish weight and density, for example. Suppose two blocks are placed on a scale and the scale is balanced. Why is this? Because the blocks weigh the same. Now suppose they're placed in water, and only one sinks. Why? Because one is more dense. Nothing puzzling here. Children between 6 and 12, on the other hand, often don't have distinct concepts WEIGHT and DENSITY, but rather have an undifferentiated concept which has some features of both, which we might call WEIGHNSITY.[4] When faced with little experiments like this one, children are often puzzled. They want to explain the first observation by saying that neither block is heavier, but also want to explain the second by saying that the block that sank is heavier. But they see that this can't be right and are suitably puzzled.

We can have a scientific account of how children understand the world and we can explain why they feel puzzlement in such cases by appeal to such accounts. But there's a sense in which we don't fully understand them. It's not just that we don't know the phenomenal what-it's-like of being puzzled about the sinking block in the

[4]See Carey (2009a, Ch. 10) for an illuminating discussion of children's undifferentiated concept of weight/density, as well as other properties they attribute to materials.

way they are, though it's true that we don't know this. Rather, it's that we don't have it in us anymore to think of things in the ways they do. When we ascribe thoughts about heaviness to them, it's not through using the same concepts they use, but instead from a perspective that is very different from theirs.

Kids can understand each other, however, in a more direct way, through grasping the same content in the same way as each other, and attributing that content. There are various respects in which their understanding of each other is much worse than the developmental psychologists' understanding of them, but there's at least one respect in which it is better. They can really get each other's puzzlement from the inside in a way we no longer can. This is, I will claim, a necessary consequence of our having some concepts which they lack. If that's right, then given our assumption of the epistemic value of an empathetic understanding of others, Costlessness is false.

That gives us a rough idea of how the argument against Costlessness will go. Now we will work through it more carefully, starting by fixing some terminology and making some assumptions about concepts.

2. Some Preliminaries Concerning Concepts

A *concept*, as I will use the term, is an ability to think some content (or contents). *Thinking*, as I will use the term, is just a very general content-directed attitude, one required for having any other attitude—like knowing, desiring, intending, entertaining, wondering—towards a content.[5]

I will remain neutral here about other ways 'concept' has been used, avoiding controversial commitments concerning the nature of minds, contents, and their

[5]I take there to be at least propositional contents and question contents. On the latter, see Friedman (2013). I take them to be relatively fine-grained, though defenders of propositions as sets of possible worlds (and sets of sets of worlds for questions) can make the usual maneuvers and make sense of everything I will say in their own terms.

relation. It will not matter to the argument whether or not concepts in our sense correspond in any direct way to mental entities (or types of mental entities) or to abstract constituents of contents. Presumably there's something about one's mind and its relation to the world together with something about a content which determines whether one is able to think that content. But we need not assume any particular account of how this works.[6]

Some examples will help make this relatively neutral use of the term clearer. The dog doesn't believe that it just ate the child's homework, not because it disbelieves it or is withholding judgment, but because it can't even think that content. The dog doesn't have the concept HOMEWORK. The child doesn't believe that whenever you divide a positive rational number by 2, you get another positive rational number, because they can't think that content. The child doesn't have the concept RATIONAL NUMBER (Carey 2009b). And Mary, not having had visual experiences of redness and so not being able to fully understand experiences of red, lacks the concept RED_{ph}, despite knowing all sorts of things about redness (Jackson 1982, 1986).

There's an objection even to what we've said so far which is worth addressing in some detail, since it will lead us to some points about concept possession which will play a role in the argument against Costlessness. That there's a phenomenal concept, RED_{ph}, which Mary doesn't have, is indeed a standard line in the literature on phenomenal concepts.[7] But there are reasons to doubt that it's true. Some, like Derek Ball (2009), have argued that any concept can be possessed through interaction with a linguistic community which uses the concept or through a properly constrained stipulation. So if there's a concept RED_{ph}, Mary already possesses it before experiencing red.

[6]For an overview of the debate on the ontology of concepts which I am hoping to sidestep, see Margolis and Laurence (2014, §1).

[7]See Sundström (2011) and references therein.

To see the problem, put aside, for the moment, this talk of a phenomenal concept RED$_{ph}$ and instead ask: does Mary have the plain old subscriptless concept RED? Can she know that tomatoes are red? That red is a primary color? It seems plausible to say that she can, which is why we feel the need to add the subscript, saying that what she lacks is the *phenomenal* concept.

Ball argues that this position is untenable. There are not distinct concepts RED and RED$_{ph}$, the latter of which Mary can only have after having red experiences. If there is a phenomenal concept RED$_{ph}$, there should be some content that it allows Mary to think that she couldn't have thought before. But such contents are not easy to find. Consider Mary at the end of her big, red-revelatory day thinking over what has happened. She might truly believe (1).

(1) I used to wonder what it's like to see red, but now I know.

The latter part of this seems to express her new knowledge of what it's like to see red—it's the sort of thing we might want to say having the concept RED$_{ph}$ is required for thinking. But it seems to be directed towards the very same content that she used to wonder. So it seems that Mary could think this content even before having the experience of red. Moreover, Mary seems to continue to believe that, for example, tomatoes are red. And it's not that she now has two beliefs, one using the concept RED and the other with her new phenomenal concept RED$_{ph}$, which she could not have had prior to her red exposure. It seems that there are no contents which acquiring the supposed phenomenal concept RED$_{ph}$ has allowed Mary to access. So if we take concepts to be abilities to think contents, it seems there must not be any concept that Mary has acquired.

I think we should respond to these and related arguments by admitting that yes, Mary could think the content expressed by 'what it's like to see red' both before and

after her red experiences. And we should think her belief that tomatoes are red has the same content it always did. But we should also hold that she now thinks them in different ways. She had the concept RED all along, but now she has the concept in a way that's different from the way she had it before.[8]

We could develop this kind of view by holding that while both Mary and those who have seen red *possess* the concept RED, only the latter have *mastery* of the concept.[9] This strikes me as rather unfair to Mary, given her extensive scientific understanding of redness and how humans perceive it. Even if we were confident that a reasonably neat scale ordering degree of grasp could be constructed for each concept and a line separating the masters from mere possessors could be drawn, placing Mary on the wrong side of this line would make mastery too restrictive a notion to be of much use.

Alternatively, we could follow Crimmins (1989) and distinguish ways of having a concept in terms of whether that concept figures into *normal* beliefs and recognitional capacities.[10] What counts as normal will vary from context to context, but in many contexts Mary's way of possessing RED indeed won't be normal. But I would like a characterization that is more flexible. Neither the mastery/non-mastery nor normal/abnormal divisions suffice for drawing all the interesting distinctions between ways of having a concept.

Recall that we're thinking of concepts as abilities. It can be helpful here to consider the distinctions we draw among familiar practical abilities, like swimming, knitting,

[8]Note that these ways are not modes of presentation in the standard Fregean sense, at least not if we take their content to be the same, since for the Fregean, differences in modes of presentation are differences in content. It might be reasonable, though, to modify the theory of content and modes of presentation to allow for at least a close correspondence between ways of thinking a content and modes of presentation.

[9]Rabin (2011) and Alter (2013) pursued this line in response to Ball. See Ball (2013) for a response.

[10]Crimmins doesn't describe his view this way, but instead distinguishes between *having an idea* of red and having the concept RED, where having a concept requires that one's idea figure into normal beliefs and recognitional capacities. I take there to be merely a verbal difference between his view and the view that we should distinguish having the concept in the normal way and having the concept in an abnormal way.

or playing guitar. Take Django Reinhardt, the great jazz guitarist. He could play the guitar as a teenager, but burns from an accident caused the fourth and fifth fingers on his left hand to be permanently paralyzed. So he relearned the guitar with a new technique for fretting, one relying almost entirely on his second and third fingers. Few would consider the way Reinhardt played the guitar 'normal', but clearly he could play the guitar, even before he had reached a level that could be considered mastery. He just could play the guitar in a way that is different from how most other guitarists can.

We can also draw further, finer distinctions among the specific ways of playing guitar. Reinhardt's style developed over the years after the accident, and we may want to think of the way he played in the fifth year after the accident as different from the way he played in tenth year after, though we could subsume both under what we might call the post-accident Reinhardt way. What counts in a context as doing something in the same way will depend on what sorts of contrasts are relevant. On some occasions the relevant contrast may be between normal and abnormal or between mastery and lesser competence, but often there are relevant distinctions among ways of having abilities that don't match either of these divisions. Different ways of dancing need not involve differences in degree of competence or normality. The tango master and amateur might count as dancing in the same way, a way which differs from the way salsa dancers dance.

Abilities, then, can be had in different *ways*, not limited to the normal and abnormal way, nor the masterful or merely competent way. One way of being able to play the guitar is to be able to play the guitar in the post-accident Reinhardt way. More usual ways involve dextrous control over the third and fourth fingers.

We can treat different ways of having an ability as themselves different abilities, abilities which are relative determinates of a determinable ability. A way of having

an ability, on this approach, is just another ability, though of a more specific kind, as scarlet is a more specific color than red. Reinhardt could play the guitar because he could play the guitar in the post-accident Reinhardt way, and he could play in the post-accident Reinhardt way at a given time because he could play in some more specific way than that. Most other guitarists can play the guitar because they can play the guitar in some other specific, more familiar way of playing guitar. So while Reinhardt and the others share one ability—the determinable guitar-playing ability—there are other, more determinate abilities which they do not all have. After the accident, Reinhardt lacked the ability to play in any of the usual specific ways and most guitarists lack the ability to play in the post-accident Reinhardt way. None of this involves having to say anything strange about a sentence like (2), like that it ascribes two different abilities.

(2) Reinhardt could play the guitar both before and after his accident.

We'd just want to note that the way he could play changed in the time after the accident. The sentence ascribes the same determinable ability, which Reinhardt has at different times through having different more determinate abilities.

All of this goes just as well for abilities to think contents. People can share the ability to think some content, but have the ability in different ways; there are different ways of having a concept. What counts as a different way of having a concept, or thinking a content, may vary from context to context.[11] Sometimes the relevant differences come down to a difference in mastery or normality, but not always. And we can have a more fine-grained division of concepts which treats different ways of

[11]This contextualism isn't due to anything special about concepts, abilities more generally, or ways of having them. It just follows from the fact that what counts as the same *N* in a context is generally a matter of contextually variable conversational standards. See Lewis (1979) on setting conversational standards, and Nunberg (1984) and Lasersohn (2000) for evidence that 'same' and 'different' are sensitive to such standards.

having some concept as possession of different, more specific concepts.[12]

Before seeing red, Mary can think that tomatoes are red and be thinking the very same content that she thinks after having seen red and that others who have seen red think. And she could wonder the very same content she later comes to know. Thus throughout she has the determinable concepts RED and even WHAT-IT'S-LIKE-TO-SEE-RED. However, her ways of having these concepts before seeing red differ from the ways that those of us who have seen red have them, and they differ from how she has them after seeing red. And so we can say that there's some specific concept, RED_{ph}, which sighted people usually have but Mary does not have until she has seen something red. None of this would require us to say anything strange about (1), like that it ascribes attitudes to two different contents. We'd just want to note that the way she can think that content changed after her experience of redness. Mary can now think this content through use of RED_{ph}, rather than some other more specific version of RED.

So we'll be working with a picture on which concepts are rather fine-grained—any way of having a concept can itself be another, more determinate concept—even if there are not similarly fine-grained contents.[13] Now we can officially state the argument against Costlessness.

3. The Argument Against Costlessness

3.1. The Argument, Briefly

(i) Many concepts, even fine-grained determinates of the same determinable concept, have independent epistemic value.

[12] We may want to revise our definition to make this more explicit, saying that concepts are abilities to think contents *in certain ways*.

[13] I think this picture of concepts is correct, and it makes the argument that follows go more smoothly, but it is not strictly necessary to establish the conclusion. To see how the argument can be recast without appeal to such fine-grained concepts, see §4.3.

(ii) Some pairs of these independently epistemically valuable concepts are incompatible, in the sense that one cannot possess both.

So, (iii) in having a concept from one of these pairs, one would necessarily miss out on some epistemic value.

So, (iv) there is an epistemic opportunity cost to having that concept.

So, (v) Costlessness is false.

The next two subsections defend (i) and (ii), respectively. This is enough to get to (v).

3.2. The Value of Concepts Redux: Epistemic Empathy

It is a fact that I believe there are cars in New York City. An alien super-scientist might understand me entirely on a chemical level, knowing this fact under some extremely long and complicated chemical description. There is a way in which they could understand this belief of mine quite well, just as they might understand well the ductility of some particular chunk of metal. They might know all sorts of details about my belief and its causal explanation, how it interacts with various other phenomena, how it would change in response to interventions, and so on. But we are assuming that there is an epistemically valuable kind of understanding which they lack. I do not feel "really understood" by this alien, since they don't have a perspective on this belief and its content that is anything like my own. Someone who knows about my belief through using using CAR and NEW YORK will have an understanding of it that comes closer to capturing my own perspective and its intelligibility. In general, I will say the more relevantly similar the concepts an agent uses to attribute an attitude to another are to the concepts the agent uses in having the attitude, the more *epistemically empathetic* their attitude attribution is.

Just as epistemic empathy and its value is not exhausted at the level of knowing facts about one's attitudes, it is not exhausted at the level of knowing facts through shared coarse-grained concepts, either. Someone might use the same coarse-grained concepts CAR and NEW YORK that I have in attributing this belief to me, but nevertheless understand my belief in a less epistemically empathetic way than someone else who shares finer-grained concepts with me. Their way of thinking about cars might go entirely through detailed thoughts about their many mechanical parts and processes or their role in the world manufacturing economy, for example, whereas I access car thoughts primarily through vague mental imagery of certain stereotypical exteriors and cabins and their roles in transporting people places for work and leisure. And suppose this interpreter of mine has only just heard of New York and knows practically nothing about it, thinking of it only in a bare, linguistically mediated way, whereas I think of it through a variety of richer means. The finer-grained determinates of CAR and NEW YORK we possess, then, are rather different. They might understand my belief more empathetically than the alien, but there is room for a still more empathetic understanding by sharing not just coarse-grained concepts but also finer-grained ones.

Not just any imaginable fine-graining of concepts should be taken to makes a difference to degree of epistemic empathy, however. Consider some extrinsic property unconnected to the content of the relevant agents' attitudes. You and I might have very similar ways of thinking thoughts about cars, even though you happened to have been within 50 feet of the artist Banksy at some point without your realizing it, whereas I have not. One may draw a distinction between the way people who have been near Banksy have the concept CAR, $\text{CAR}_{\text{near-Banksy}}$, and the way of those who have not, $\text{CAR}_{\text{not-near-Banksy}}$. But even allowing these to count as different ways of having the concept CAR, the fact that you use the former and I use the latter should make no

difference in the epistemic value of attitude attributions.

What kind of difference among fine-grained concepts is relevant, then? One natural thought is that what is important is phenomenal similarity. My experience in entertaining the proposition about there being cars in New York is likely much more similar those of the more epistemically empathetic interpreters of me than to that of the alien. And having unwittingly been near Banksy at some point makes no difference to one's experiences, thoughts of cars included.

The role of phenomenal experience seems even clearer for the case of Mary. She can know that I believe tomatoes are red, that I like my red couch, that I know what red is like, and so on. We may share RED, but the more specific concept she uses, call it RED_{Mary}, is rather different from mine, which is a determinate of RED_{ph}. Her understanding of me is thus that of a relative outsider, and is in some respect worse than it could be. And when Mary does acquire RED_{ph}, this will not only be an epistemic improvement with respect to her understanding of color experiences, but also with respect to her understanding of the attitudes of people like me whose red-related thoughts are had through RED_{ph}. And this seems largely because of the different phenomenal experience of thinking of redness which it involves.

Though I do think phenomenal similarity can matter for epistemic empathy, it does not seem to me to be the only feature that matters. We should allow similarities and differences across some broader notion of what we might call a concept's 'cognitive role' to count as relevant: the attitudes that are taken to support it and that it is taken to support, the questions it raises and answers (and answers it could have, if the content itself is a question), the associations it tends to bring, and so on. Why think this? First, generally, because many of our attitudes do not seem to have characteristic phenomenal accompaniments, particularly our non-occurrent attitudes, yet it is possible for such attitudes to be understood in more or less epistemically empathetic

ways. Second, because there are specific cases where factors other than phenomenal similarity seem to count towards a better 'from the inside' understanding.

Consider, for example, two mathematicians who have worked on the same specialized topic in isolation for years. They meet by chance and find their minds race almost identically along the same lines, both on the questions they have already considered and on any new ideas that occur to them on the topic. "Finally", each thinks, "someone who really understands how I think about all of this!". Here we have a case of relatively high degree of epistemic empathy, even if it turns out that their phenomenology is rather different. Suppose one has chronic pain, so that any time they are entertaining some conjecture, their phenomenology includes pain experience. At least on one plausible way of comparing phenomenal states, someone else in chronic pain who attributes a thought about this conjecture—even though they themselves would think very differently about it—would do so from a more similar phenomenal state than would the new pain-free friend. Nevertheless, the pain-free friend would have a better understanding of the belief in question. This would be so even if it is the thought of the conjecture itself which brings with it some pain for one of the mathematicians but not the other.

And recall the case of understanding children from §1. One may be able to conjure a feeling of puzzlement about some other phenomenon (the Liar Paradox, say), that turns out to feel much like the feeling children have when they see the heavier block floating while the lighter one sinks. This would go some way towards understanding their perspective, perhaps, but one would have much a better understanding of it if one could use the WEIGHNSITY concept to see for oneself the behavior of the blocks as unintelligible.

Epistemic empathy requires having concepts in the relevant ways that other people have them, where this includes both phenomenal similarity and similarity of

some broader cognitive role. But people have all sorts of concepts in all sorts of ways. And we might take there to be value in being able to understand in an epistemically empathetic way the thoughts of past, future, and even merely possible people. Even limiting ourselves to the actual and present, very many ways of having concepts will have some independent value, since they each allow for greater degree epistemic empathy with someone or other.

So $\textsc{red}_{\text{ph}}$ is useful not only for thinking about tomatoes, but for a better understanding my beliefs about tomatoes. And $\textsc{red}_{\text{Mary}}$ also has some epistemic value, in making possible epistemic empathy with people like Mary. Lacking this concept, I understand her beliefs about redness, if anything, even less empathetically than she understands mine. Having DENSITY and WEIGHT not only has the value of making sense of sinking blocks, but also of understanding DENSITY- and WEIGHT-users' attitudes in an epistemically empathetic way. But this also goes for (more) confused concepts like WEIGHNSITY. Even if it's relatively deficient for making sense of the sinking blocks, it does have independent value in allowing a more epistemically empathetic understanding of the children who use that concept. Given the fine-grained account of concepts we're working with, this means that many very fine-grained determinate concepts have some independent value, even relative to determinates of the same determinables.

3.3. Incompatible Concepts

Could I learn Mary's way of thinking of red or the children's way of thinking of heaviness while preserving my own current ways of thinking of redness and heaviness? Could I acquire $\textsc{red}_{\text{Mary}}$ without losing access to $\textsc{red}_{\text{ph}}$? Or WEIGHNSITY without losing DENSITY? I think not.

Some of the significant differences in the ways Mary and I think about red things

are owed to the fact that Mary is a leading color scientist and I could tell you only rudimentary things about the physics of color and the neurophysiology of its perception. These differences, perhaps, I could reduce or even eliminate by dedicating my life to the study of color under Mary's tutelage. I could develop a scientific way of thinking about red which is more in line with the way Mary thinks of it. This may bring me closer to Mary's way of thinking red-related contents, and for that reason (among others) it would be valuable. I could better understand Mary's thoughts than I currently would be able to. But there will still be a crucial difference, at least for certain of these contents. I still won't be able to think the content of 'what it's like to see red' in anything like the way Mary does—thinking this content will have both a different phenomenology and cognitive role for me. When Mary considers the question of what it's like to see red, try as hard as she might, no plausible answers come to her mind. It's not merely that she doesn't know which possible answer is the correct one, it's that no plausible full answers are within her grasp.[14] When I consider this content, however, I have easily within reach a variety of full answers that I think with my phenomenal color concepts, including, of course, the correct answer, which I think with $\textsc{red}_{\text{ph}}$. I won't really get Mary's wondering what it's like to see red, except in a distanced, impersonal way.

What's getting in the way of epistemic empathy with Mary is my having the concept $\textsc{red}_{\text{ph}}$. And this is something which I could not change by book-learning or experience. So if I am to fully understand Mary, I'll have to *forget* what red is like, and even what it could be like. I'll have to *lose* my $\textsc{red}_{\text{ph}}$ concept, and perhaps also $\textsc{blue}_{\text{ph}}$, $\textsc{green}_{\text{ph}}$, and so on. There is no addition to my conceptual scheme which can fully subsume both the ways I can think red-thoughts and the ways Mary can think

[14] By 'full answer' I just mean the kind of answer someone with the phenomenal concept knows, rather than negative answers like 'not like experiencing a sound' or very general ones like 'it's a visual experience'.

red-thoughts. Our determinates of RED are incompatible, in the sense that one cannot possess both.

The same holds for WEIGHNSITY, at least in the way children have this concept. Like in the case of RED_{Mary}, there will be WEIGHNSITY-related questions which I will not be able to think of as children do. When I think about the question why the block sank, I will think of it in a way that includes as a possible answer that it is dense, whereas it is an important fact about the children's thinking that they do not have such an answer available—their lacking this option and only having the WEIGHNSITY answer available is what's behind their puzzlement. So even if I acquire some version of WEIGHNSITY, it will not suffice for epistemic empathy with the children. I will not understand their puzzlement from the inside.

It may be helpful to again draw an analogy with practical abilities. The point is not that distinctions among ways of having abilities should always be so fine-grained that no two people can do something in the same way unless they can do everything in the same way. We want to allow, for example, that a switch hitter in baseball can bat in the way a lefty can and the way a righty can. We're not assuming that one must have at most one way of being able to think some content.

The point is that some abilities are incompatible. Nobody could have both the ability to make something so heavy that nobody could lift it *and* have the ability to lift anything that anybody could make, though it's possible for someone to have one or the other of these abilities.[15] Or consider the ability to make a table made only by Alvin Plantinga and the ability to make a table made only by someone

[15]There has been much discussion of this case and its implications for the possibility of an omnipotent god. See Mackie (1955), Keene (1960), Mayo (1961), Mavrodes (1963), Frankfurt (1964), Cowan (1965), and Swinburne (1973). Some say that the ability to make something so heavy that nobody could lift it is impossible, because there's a necessary being which can lift any possible thing. Regardless of how one comes down on the stone-creating ability, I do think there are some possible abilities which are logically impossible for a single agent to jointly possess. Nevertheless I do not wish here to take any stance on whether this should trouble those who think there could be a being reasonably described as omnipotent.

other than Alvin Plantinga. Nobody could have both of these abilities, since Alvin Plantinga is the only person who could have the former ability, but he is of course barred from having the latter.[16] And these incompatibilities are not the result of contingent limitations on humans. Rather, the impossibility of having these pairs of abilities together comes from the nature of the abilities themselves. If each ability had independent ethical value, then the practical analogue of Costlessness would be false. Having one of these abilities would impose a cost on even a fully ideal agent.

For a less gimmicky example closer to our target case, consider the following. Bob is a mediocre chess player. So am I, to put it generously. One thing I am able to do is lose to Bob in chess. And not only that, but I can lose to Bob in chess while being sober, focused, and trying as hard as I can to win. Judit Polgár, who has been a Grand Master since she was 15 years old, cannot lose to Bob unless she is either drugged or not trying to win. I cannot play chess in the ways Polgár can, of course, but neither can she play in all of my ways. Her ways involve seeing various moves and recognizing which are better, whether or not she decides to play one of the better moves. My ways of playing, though, involve *not* seeing many of those very moves, or not being able to evaluate them appropriately. I can make certain bad moves while thinking they're good. Polgár couldn't make such moves without seeing that they're bad ones. After observing me, Polgár could probably imitate my way of playing chess. She could play in ways that appear to an outside observer to be the same ways I can play. But she'll be doing so by seeing better moves and choosing the worse, whereas I am doing nothing of the kind.[17]

[16] See Plantinga (1967, pp. 169–170) for an example like this.

[17] *Objection*: haven't I shifted use of 'ability' here? Being able to play the guitar is an intentional ability, something one is able to do by trying to do it. But losing while trying to win is not something one does intentionally. So now we're talking about abilities in some other sense, more like the ability water has to dissolve salt.

Reply: (a) we should not restrict ourselves to intentional abilities. We're interested in abilities to think certain contents, and I don't think these can be abilities in any straightforward intentional sense (otherwise we get a regress: if you think the content that tomatoes are red by trying to think the

Not being able to play chess in just the mediocre way I play is not in itself of any disvalue. Not being able to think in the ways others do, though, is epistemically costly, even if their ways of thinking are inadequate or defective. It prevents one from fully understanding how those others think. Since having the concept $\textsc{red}_{\text{ph}}$ prevents one from thinking in the way that Mary and others do, merely having this concept has an epistemic cost. Similarly with possessing DENSITY, which prevents one from empathetically empathizing with children. The benefits from having these concepts may outweigh their costs, but there are costs nevertheless. Costlessness is false, given the value of epistemic empathy.

4. Objections & Replies

4.1. How Quickly Does Mary Forget?

Part of what was involved in Mary's original way of having RED was not being able to visually imagine what it is like, despite trying to do so. Part of what is crucial to having $\textsc{red}_{\text{ph}}$ is being able to visually imagine what red is like. These are incompatible, not as a matter of contingent human limitations, but in principle, due to the nature of the concepts. But do I really want to say that Mary would lose access to her old way of thinking in an instant? That seeing red would somehow make her forget how she was thinking just moments before? This seems implausible.

I grant to the objector that this is unlikely. But I think it's unlikely not because the concepts are compatible after all, but because I doubt Mary would acquire $\textsc{red}_{\text{ph}}$ instantaneously.

We don't usually acquire concepts in a flash. With children acquiring WEIGHT and

thought that tomatoes are red, didn't you already have to think that content in order to try?); (b) it doesn't matter so much whether we say I have an ability to lose to Bob that is incompatible with some ability of Polgár's. Rather, the case is meant to show that being able to play chess in the ways Polgár does is incompatible with being able to play chess in the ways I do.

DENSITY, it's a gradual process, with various stages on the way to full adult-level competence with these concepts. It's not implausible that by the time they've reached that competence, their old ways of thinking will not be accessible to them, even in memory.

Similarly, I think, for Mary and RED_{ph}. From the literature on Jackson's knowledge argument, one sometimes gets the impression that all it would take for Mary to acquire RED_{ph} in a full-blown way is a glimpse of a tomato. But I find this doubtful. Imagine, for instance, Mary gets to look at a motionless tomato in normal lighting for 5 seconds, then must go back to seeing only black and white. Would this be sufficient for acquiring RED_{ph}? If later she sees the same color, will she know that it's red? Will she be able to clearly imagine it in the days (or even hours) after the event? It's an empirical matter, of course, but I would guess not.[18]

For visual concepts, it can be tempting to think that they would somehow come fully 'given' in someone's first visual sensation of the relevant kind of scene, and that 'learning' them is simply a matter of having had such sensations. We should resist this temptation. Visual perception is a complicated affair, normal development of which depends on experiences in complicated ways. Color perception is not merely a matter of detecting wavelengths of light, as various optical illusions make vivid. And going beyond perceiving—imagining, expecting, preferring, etc.—involves even more complication. How quickly Mary would acquire the abilities required for having RED_{ph} is an interesting, open, and largely empirical question.

Vision scientists have in fact tried depriving non-human primates of normal color

[18]It's worth drawing a comparison here to Molyneux's Problem. There the question is whether a newly sighted person could recognize a shape of which they had previously had only tactile experience. Though many philosophers have speculated about the answer to this, it's generally been recognized that it is a substantive empirical question, and there has been much interest in the actual experiences of those gaining sight after cataract surgery. For a discussion of some empirical work on the problem and some of the complications in interpreting it, see Schwenkler (2013). For a history of the debate, see Davis (1960) and Degenaar and Lokhorst (2017).

experiences. For example, the monkey Femke was born and reared in an environment illuminated only by red lamps, preventing stimulation of the photoreceptive cone cells for perception of blue and green. And while this had surprisingly little long term effect on abilities to discriminate by color,[19] even after two months of life in normal lighting, Femke still reacted to color tasks in somewhat abnormal ways.[20] There has also been some study of congenitally blind humans who gain sight relatively late in life through surgery. On some tasks, such as the Ishihara test for color blindness (the one with letters and numbers made up of little circles of different colors), normal-sighted-level abilities develop very quickly.[21] Performance even on relatively low-level color perception tasks, though, like those for discerning hue discrimination thresholds, improve gradually after the surgery, tending to reach normal levels only after about one year.[22]

I don't take these results to be conclusive, but it seems the relevant empirical work suggests Mary would not acquire $\textsc{red}_{\text{ph}}$ immediately, so we shouldn't expect an immediate loss of $\textsc{red}_{\text{Mary}}$.[23]

Suppose, however, that it really did happen in an instant, either in the way it would happen 'naturally' or through some sci-fi intervention. In such a case the general point made above would apply, and so yes, I think Mary would have lost her old way of thinking in an instant. There would be a respect in which she would not

[19] In contrast, effects of deprivation on acuity and depth perception are typically large and long-lasting.

[20] See Brenner, Schelvis, and Nuboer (1985) and Brenner, Cornelissen, and Nuboer (1990). Unfortunately the authors do not report any tests prior to two months after normal illumination.

[21] In one well-known early case study of the patient S.B., Gregory and Wallace (1963) register surprise that S.B. got every item on the Ishihara test correct. It's worth noting, however, that they only did this test 48 days after S.B. gained sight, and that his likely having worse than normal visual acuity may have improved his performance, as discussed in Gordon and Field (1978).

[22] See McKyton et al. (2015).

[23] But even supposing the process of acquiring $\textsc{red}_{\text{ph}}$ is gradual, wouldn't the intermediate stages also involve different ways of thinking about red, ways incompatible with Mary's original way of thinking about it? Yes, but only as described in contexts with very high standards for 'same way of thinking'. Her interpretation of herself from these intermediate stages will be more epistemically empathetic than they are at the end of the process, but not maximally so.

be able to understand herself from just moments before as well as she used to.

Should we call what happens to Mary once she acquires $\textsc{red}_{\text{ph}}$ 'forgetting'? Perhaps, but we should note that it doesn't necessarily involve forgetting any particular episodes, it just means certain aspects of those episodes are not fully accessible. Compare: I have some memories (or so it seems to me) of not being able to read. But I can't recall in a detailed way what it was like to see English words without being able to read them. And I can't bring to mind any precise enough image of English text without also interpreting that text in the way I could not do as a young child. I did not forget these events, but I can no longer understand them in the way I used to. I'm glad I learned to read, but do think I lost something of value in the process. I think of what happens to Mary as she acquires $\textsc{red}_{\text{ph}}$ as similar to this.

4.2. Self-restraint?

Having the concepts $\textsc{red}_{\text{ph}}$ or DENSITY, I've claimed, blocks epistemic empathy with someone who thinks with $\textsc{red}_{\text{Mary}}$ or WEIGHNSITY, since having the latter involve lacking abilities that the former require. But why couldn't someone with $\textsc{red}_{\text{ph}}$ or DENSITY just exercise some self-restraint and refrain from using those abilities?

For example, after seeing red, couldn't Mary just decide not to use the new abilities that allow her to know what phenomenal red is like when she is interpreting her past self? That is, in recalling her state of mind, couldn't she simulate her former ignorance by just refusing to use her new concept? And couldn't adults just refrain from using their DENSITY concept, and be able to interpret children's puzzlement about sinking blocks in an epistemically empathetic way? Or better yet, couldn't one temporarily block one's access to certain concepts, and so be in a state just like those one is trying to understand? Even if this is psychologically impossible for humans, it seems like something a cognitively ideal agent could do if it needed to.

I have some sympathy for this proposal. Restricting access to certain concepts for the purpose of interpreting others, I think, can give one a better, more epistemically empathetic understanding of their thoughts.[24] Grasping a content with some self-imposed restrictions does make one's way of thinking that content more similar to the one who lacks the concepts one is restricting access to, but I don't think it gets one all the way there—there will still be a significant difference in how one thinks the content in question. The way pre-revelation Mary thinks about red means that when she is trying *in an unrestricted way* to imagine what red is like, she still has no idea what the answer could be. So while someone who has the concept $\textsc{red}_{\text{ph}}$ may be able to think 'what is red like'-thoughts in a way very similar to Mary's old way by restricting access to it, there will still be a way in which it differs. And the way that children think means that when they try *to the best of their abilities, without any kind of self-imposed restriction*, they cannot come up with answers that would make sense of why the block that sank doesn't outweigh on the scale the block that floated. One cannot duplicate this with self-restriction. Even allowing self-imposed restrictions, it looks like the presence of a concept like DENSITY will prevent the possession of a concept of WEIGHNSITY with the same kind of cognitive role that it has in the children's minds.

But, we might rejoin, is there really a relevant difference here in cognitive role here or will this amount to some irrelevant difference like the one between $\textsc{car}_{\text{near-Banksy}}$ and $\textsc{car}_{\text{not-near-Banksy}}$? Lacking a proper theory of epistemic value and cognitive role, I am left only with the option of reporting my judgements about cases and inviting others to share them. And it seems to me that someone who can restrict their access to $\textsc{red}_{\text{ph}}$ can understand pre-revelation Mary from the inside better than someone

[24] Indeed, elsewhere I argue on this basis that having fragmented mental states can be epistemically ideal, contrary to the usual assumption that fragmentation is always a non-ideal state (Deigan 2020, Ch. 5).

who has $\textsc{red}_{\text{ph}}$ but cannot restrict access in this way, but still not quite as well as someone who, like Mary, also lacks $\textsc{red}_{\text{ph}}$. Suppose that like Mary, Barry has grown up in a room without being visually exposed to redness, and so lacks $\textsc{red}_{\text{ph}}$, whereas Cary normally thinks with $\textsc{red}_{\text{ph}}$, but is able to temporarily restrict access to it for the purpose of interpreting people like Mary. There's a real sense, it seems to me, in which Barry shares Mary's way of thinking in a way that Cary doesn't, one that affords him an understanding of Mary that is at least in some respect better than Cary's.

For those who don't share this judgement, I'm afraid I don't have an argument that they should. But drawing out a difference between Barry and Cary's understanding of Mary might help one see how they could differ in epistemic value. Consider how Barry and Cary's judgements about Mary can be integrated into their respective cognitive economies. Barry can attribute RED-involving thoughts to Mary while still having access to all the concepts he uses for other purposes. So integrating his Mary attributions with his other thoughts will require no filtering or reinterpretation.

Contrast this with Cary. In interpreting Mary, we're supposing, he has restricted access to $\textsc{red}_{\text{ph}}$. For maximal epistemic empathy, we can suppose his simulation of Mary's failure to grasp that concept is so complete that while it is ongoing he doesn't even realize that he has any sort of access to $\textsc{red}_{\text{ph}}$, and that the restriction isn't one that he can easily override, but instead only expires after some set time limit. At the moment of interpretation, as he attributes some red-thought to Mary, things may seem just the same to him as they did to Barry. But now consider how his interpretation is to be integrated with the rest of his attitudes, including his other beliefs about Mary. Many of those attitudes, presumably, involve use of $\textsc{red}_{\text{ph}}$, so his access to it will need to again be unrestricted for integration to take place. Cary may have some memory of the restricted-access period, but now he is in a position similar

to that of Mary recalling her former $\textsc{red}_{\text{ph}}$-impoverished state: he is doing the best he can to make sense of his former state, one in which he lacked access to some of his current resources. This doesn't mean that nothing about his Mary interpretation can filter through, but it does mean that it cannot be integrated without change or loss, in the seamless way that Barry's interpretation can be integrated. He will no longer be able to think the relevant contents from the same limited perspective he had during the restricted period.

There is no knockdown argument here for sharing my judgement about the case, of course. It can be denied that the failure to fully integrate makes Cary's understanding of Mary in any way worse than Barry's. But, to reiterate, it seems to me that there is a way in which it is worse, and given that successful understanding seems to be crucially connected to integration with one's other information,[25] it is plausible that this is at least a partial explanation of why it would be worse. If this is right, the cost of concept possession cannot be fully eliminated by restricting one's concepts, even if, like Cary, one can restrict concepts with much more thoroughness than humans can normally manage.

4.3. Costless Coarser Concepts?

Here's another worry about my argument. I've been assuming that for each way of having a concept, there corresponds a more determinate concept. Perhaps this assumption had seemed innocuous enough, but now having seen where it leads, a defender of Costlessness may wish to deny it. Concepts, they might hold, are less fine-grained than I've been assuming. There may be different ways of having RED, but these ways of having concepts should not themselves be taken to be concepts. So

[25] See Kvanvig (2003, p. 192), Elgin (2007), Gardiner (2012), Grimm (2016, pp. 215–216), and Bengson (2017, pp. 37–39), among others. It should be noted, however, that not all of these philosophers are talking about understanding people, and not all of their points will uncontroversially carry over.

the argument that having certain concepts is incompatible with having others doesn't go through, since it relied on incompatibility of different ways of having a concept showing that there are incompatible concepts. So the argument against Costlessness fails.

Let's grant to the objector that there are only relatively coarse concepts and that different ways of having these won't count as separate concepts. My argument against Costlessness can be reformulated in a way that accommodates this. All we need to do is revise it as follows:

(i*) Many different ways of having concepts have independent epistemic value.

(ii*) Some of these independently epistemically valuable ways of having concepts are incompatible with having certain coarse-grained concepts.

So, (iii*) in having one of these coarse-grained concepts, one would necessarily miss out on some epistemic value.

From there the argument proceeds as before.

We can defend (i*) as we defended (i)—epistemic empathy requires having the same concepts *in the same ways* as the person to whom one is attributing the thought. And then we just need to argue that having certain concepts (in any way) are incompatible with having some concepts in a particular way.

Having the concept DENSITY prevents you from having the WEIGHNSITY *in the way that children have it*. You can't think about question of the sinking block in the same way that children can. So you can't attribute puzzlement about why the block sank in an epistemically empathetic way. So given that this way of having WEIGHNSITY has some value—it allows for epistemic empathy with children—simply having DENSITY will have an epistemic cost. Even limiting ourselves to coarse-grained concepts, then, Costlessness is false, given the value of epistemic empathy.

5. Conclusion

I have been assuming that a certain kind of empathetic understanding has epistemic value, without doing much to defend that assumption. Officially, then, we should take the conclusion here to be the conditional one: if this kind of epistemic empathy is valuable, then having a concept has a cost. The implications of this conclusion will depend on whether this kind of epistemic empathy really is valuable, as well as how valuable it is, and whether and how its value varies from case to case. I am uncertain about the answers to these questions, but even if I did have settled views, it would be too great a task to adequately defend them here. Nevertheless I will conclude by pointing towards some potential implications of this paper's central argument, suggesting lines of further research that may be worth pursuing.

First, the most direct implications, I suspect, will be for the normative epistemology of concept possession and change. The argument against Costlessness suggests that it may sometimes be epistemically permissible to refrain from acquiring or even to lose a concept, even idealizing away contingent limitations.[26] The route to this is clearest on telic theories, according to which the expected epistemic value of concept possession directly determines whether it epistemically ought to be possessed. But there are likely to be consequences even for more deontically or virtue-theoretically flavored approaches to epistemology, since epistemic value often plays a central role even for these other kinds of theories, just not as a target of direct maximization.[27]

Second, there will also likely be implications for various projects in conceptual ethics and engineering, at least those that aim to target concepts (as opposed to language).[28] There is a wide diversity of aims and approaches among these projects;

[26]It also suggests an interesting possible descriptive phenomena to be on the lookout for: do agents or communities ever in fact forgo potential concept acquisition in order to maintain an empathetic understanding with those who lack the concepts in question?

[27]See, for instance, Sylvan (2020).

[28]For overviews, see Burgess and Plunkett (2013a,b), Burgess, Cappelen, and Plunkett (2020), Isaac,

what implications can be drawn from my argument will likely depend on which of these one has in mind. Nevertheless, I see two main lines of thought that should apply relatively widely. One goes through the epistemology of concepts. Some work in conceptual engineering is meant to be responsive to or constrained by epistemic considerations, so insofar as the argument developed here has consequences for the epistemology of concepts, it will have consequences for conceptual engineering as well. The other goes more directly through the possibility of incompatible concepts. If possessing one concept prevents one from possessing others, conceptual engineers should take into account not only the relatively direct merits of possessing a given concept, but also the opportunity cost from excluding possession of certain others.

Third, assuming that epistemic reasons have at least some all-things-considered authority or influence on practical reasons, there will be downstream consequences for practical decision-making. In principle this applies generally, but is likely to be particularly relevant to those thinking about transformative decisions (Paul 2014), which will often involve significant conceptual change.[29]

Finally, there may be an implication similar to one which Thomas Nagel has argued for in various places.[30] On Nagel's view, there is no way that one could come to know what it is like to be a bat through purely objective inquiry. Nothing about what Nagel claims, though, rules out coming to understand the experience of bats through other means. For all he says, an agent might come to know all there is to know about subjective experiences, they just can't accomplish this with the tools of objective inquiry. Indeed, he sometimes claims that a "being of total imaginative flexibility could project himself directly into every possible subjective point of view, and would not need such an objective method to think about he full range of possible

Koch, and Nefdt (2022), and Koch, Löhr, and Pinder (2023).

[29]See Bailey (2023) for an insightful discussion of this issue.

[30]See Nagel (1974), Nagel (1986), and Nagel (1997, Ch. 2).

inner lives" (Nagel 1986, p. 17).[31]

If my argument from against Costlessness works, it pushes us beyond Nagel's skepticism about the reach of an objective point of view. Not only is it impossible to fully understand everything about subjective experience from an objective point of view, it is impossible to fully understand it from *any* point of view, subjective or objective, no matter how imaginatively flexible one is. There is no view from anywhere from which everything worth understanding can be fully understood. Once one has the concepts required to see from one perspective, one thereby loses what's required to see from others.

[31]This idea is similar to the phenomenal simulation part of Chalmers's Cosmoscope, see Chalmers (2012, p. 115).

References

Alter, Torin (2013). "Social Externalism and the Knowledge Argument". In: *Mind* 122, pp. 481–496.

Bailey, Olivia (2022). "Empathy and the value of humane understanding". In: *Philosophy and Phenomenological Research* 104.1, pp. 50–65.

— (2023). "What must be lost: on retrospection, authenticity, and some neglected costs of transformation". In: *Synthese* 201.6, p. 189.

Ball, Derek (2009). "There Are No Phenomenal Concepts". In: *Mind* 118, pp. 935–962.

— (2013). "Consciousness and Conceptual Mastery". In: *Mind* 122, pp. 497–508.

Bengson, John (2017). "The Unity of Understanding". In: *Making Sense of the World: New Essays on the Philosophy of Understanding*. Ed. by Stephen R. Grimm, pp. 14–53.

Brenner, Eli, Frans Cornelissen, and J. F. W. Nuboer (1990). "Striking absence of long-lasting effects of early color deprivation on monkey vision". In: *Developmental Psychobiology: The Journal of the International Society for Developmental Psychobiology* 23.5, pp. 441–448.

Brenner, Eli, J. Schelvis, and J. F. W. Nuboer (1985). "Early colour deprivation in a monkey (Macaca fascicularis)". In: *Vision Research* 25.9, pp. 1337–1339.

Burgess, Alexis, Herman Cappelen, and David Plunkett, eds. (2020). *Conceptual Engineering and Conceptual Ethics*. Oxford University Press.

Burgess, Alexis and David Plunkett (2013a). "Conceptual Ethics I". In: *Philosophy Compass* 8.12, pp. 1091–1101.

— (2013b). "Conceptual Ethics II". In: *Philosophy Compass* 8.12, pp. 1102–1110.

Carey, Susan (2009a). *The Origin of Concepts*. Oxford University Press.

— (2009b). "Where Our Number Concepts Come From". In: *The Journal of Philosophy* 106.4, pp. 220–254.

Carr, Jennifer Rose (2015). "Epistemic Expansions". In: *Res Philosophica* 92.2, pp. 217–236.

— (2021). "Why Ideal Epistemology?" In: *Mind*.

Chalmers, David J. (2012). *Constructing the World*. Oxford University Press.

Cowan, J. L. (1965). "The Paradox of Omnipotence". In: *Analysis* 25.3, pp. 102–108.

Crimmins, Mark (1989). "Having Ideas and Having the Concept". In: *Mind & Language* 4.4.

Davis, John W. (1960). "The Molyneux Problem". In: *Journal of the History of Ideas* 21.3, pp. 392–408.

Degenaar, Marjolein and Gert-Jan Lokhorst (2017). "Molyneux's Problem". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2017. Metaphysics Research Lab, Stanford University.

Deigan, Michael (2020). "On the Epistemology of Concepts". PhD thesis. Yale University.

Egré, Paul and Cathal O'Madagain (2019). "Concept Utility". In: *Journal of Philosophy* 116.10, pp. 525–554.

Elgin, Catherine Z. (2007). "Understanding and the facts". In: *Philosophical Studies* 132, pp. 33–42.

Frankfurt, Harry G. (1964). "The Logic of Omnipotence". In: *The Philosophical Review* 73.2, pp. 262–263.

Friedman, Jane (2013). "Question-Directed Attitudes". In: *Philosophical Perspectives* 27, pp. 145–174.

— (2018). "Junk beliefs and interest-driven epistemology". In: *Philosophy and Phenomenological Research* 97.3, pp. 568–583.

Gardiner, Georgi (2012). "Understanding, Integration, and Epistemic Value". In: *Acta Analytica* 27, pp. 163–181.

Gordon, Ian E. and David Field (1978). "A possible explanation as to why the newly sighted commonly perform well on pseudoisochromatic colour vision tests". In: *Perception* 7.1, pp. 119–122.

Gregory, R. L. and J. G. Wallace (1963). "Recovery from early blindness". In: *Experimental psychology society monograph* 2, pp. 65–129.

Grimm, Stephen R. (2016). "How Understanding People Differs from Understanding the Natural World". In: *Philosophical Issues: A Supplement to Noûs* 26, pp. 209–225.

Hannon, Michael (2020). "Empathetic understanding and deliberative democracy". In: *Philosophy and Phenomenological Research* 101.3, pp. 591–611.

Harman, Gilbert (1986). *Change in View: Principles of Reasoning*. MIT Press.

Hofweber, Thomas (2023). "Inescapable Concepts". In: *Australasian Journal of Philosophy*.

Isaac, Manuel Gustavo, Steffen Koch, and Ryan Nefdt (2022). "Conceptual engineering: A road map to practice". In: *Philosophy Compass* 17.10.

Ismael, Jenann (2018). "Why (study) the Humanities?" In: *Making Sense of the World: New Essays on the Philosophy of Understanding*. Ed. by Stephen R. Grimm. Oxford University Press, pp. 177–193.

Jackson, Frank (1982). "Epiphenomenal Qualia". In: *The Philosophical Quarterly* 32.127, pp. 127–136.

— (1986). "What Mary Didn't Know". In: *The Journal of Philosophy* 83.5, pp. 291–295.

Keene, G. B. (1960). "A Simpler Solution to the Paradox of Omnipotence". In: *Mind* 69.273, pp. 74–75.

Koch, Steffen, Guido Löhr, and Mark Pinder (2023). "Recent work in the theory of conceptual engineering". In: *Analysis* 83.3, pp. 589–603.

Kvanvig, Jonathan L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.

Lasersohn, Peter (2000). "*Same*, Models and Representation". In: *Semantics and Linguistic Theory*. Vol. 10, pp. 83–97.

Lewis, David (1979). "Scorekeeping in a language game". In: *Journal of Philosophical Logic* 8, pp. 339–359. Reprinted in Lewis (1983, pp. 233–249).

— (1983). *Philosophical Papers*. Vol. I. Oxford: Oxford University Press.

Mackie, J. L. (1955). "Evil and Omnipotence". In: *Mind* 64.254, pp. 200–212.

Margolis, Eric and Stephen Laurence (2014). "Concepts". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2014. Metaphysics Research Lab, Stanford University.

Mavrodes, George I. (1963). "Some Puzzles Concerning Omnipotence". In: *The Philosophical Review* 72.2, pp. 221–223.

Mayo, Bernard (1961). "Mr. Keene on Omnipotence". In: *Mind* 70.278, pp. 249–250.

McKyton, Ayelet et al. (2015). "The limits of shape recognition following late emergence from blindness". In: *Current Biology* 25.18, pp. 2373–2378.

Nagel, Thomas (1974). "What is It Like to be a Bat?" In: *Philosophical Review* 83.4, pp. 435–450.

— (1986). *The View From Nowhere*. Oxford University Press.

— (1997). *The Last Word*. Oxford University Press.

Nunberg, Geoffrey (1984). "Individuation in context". In: *Proceedings of the West Coast Conference on Formal Linguistics*. Vol. 2, pp. 203–217.

Paul, L. A. (2014). *Transformative Experience*. Oxford University Press.

Pérez Carballo, Alejandro (2020). "Conceptual Evaluation: Epistemic". In: *Conceptual Engineering and Conceptual Ethics*. Ed. by Herman Cappelen Alexis Burgess and David Plunkett. Oxford University Press, pp. 304–332.

Plantinga, Alvin (1967). *God and Other Minds: A Study of the Rational Justification of Belief in God*. Cornell University Press.

Rabin, Gabriel (2011). "Conceptual mastery and the knowledge argument". In: *Philosophical Studies* 154, pp. 125–147.

Schwenkler, John (2013). "Do things look the way they feel?" In: *Analysis* 73.1, pp. 86–96.

Sider, Theodore (2012). *Writing the Book of the World*. Oxford University Press.

Staffel, Julia (2020). *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford University Press.

Sundström, Pär (2011). "Phenomenal Concepts". In: *Philosophy Compass* 6.4, pp. 267–281.

Swinburne, Richard (1973). "Omnipotence". In: *American Philosophical Quarterly* 10.3, pp. 231–237.

Sylvan, Kurt (2020). "Epistemic Nonconsequentialism". In: *Philosophical Review* 129.1.