

Having a Concept Has a Cost

Michael Deigan

Rutgers University

... 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. It gives one access to contents which may allow for a better understanding of the world (Egré and O'Madagain 2019). It also allows one to interpret others who use that concept in an especially direct way, by using in one's attribution the very concept which they used to think (Carr 2015, p. 224). But does having a concept ever have an epistemic cost?

At least when we idealize away from our computational limitations, it is tempting to think that it cannot. Concepts can be misapplied, of course, 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 for them to have a false (or unjustified) belief. In their pursuit of epistemic goods, they must take care to avoid some epistemic bads; 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 gerrymanderd, 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 would

seem not apply to the ideal agents who lack our contingent limitations. So if Costlessness is true, then an ideal epistemology of concepts seems rather uninteresting, in contrast with an ideal epistemology of beliefs.¹

I will be arguing that Costlessness is false, and not just in some fringe cases, but quite generally. Possession of concepts prevents us from seeing as those who lack those concepts see, be they others or our past or future selves. This, I will argue, is an epistemic cost, and it is a cost that would apply to agents no matter how ideal.

The paper has two parts. The argument against Costlessness occupies §1. In §2 I present some of this argument's interesting implications. One is that there may be 'bad' concepts, which are epistemically worse to possess than to lack, and so an ideal epistemology of concepts is not trivially permissive. Another is that there cannot be a single perspective from which everything worth understanding can be understood—even maximally ideal agents will be missing out on some epistemic value.

1. The Argument Against Costlessness

Let's start with some examples.

First 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*.² When faced with little experiments like this one, children are often puzzled. They want to explain the first observation

¹There 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).

²See 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.

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—children tend not to be dialetheists—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 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 concept they use, but instead from a perspective that is not at all like 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 it seems to me that there's at least one respect in which it is better. They can really get each other's puzzlement in a way we no longer can.

Second Example: What Mary *Did* Know.

Mary, the color scientist introduced by Jackson (1982, 1986), is famous for her ignorance. But I want to draw attention to something that she *did* know, something of which most of us are ignorant.

Mary is a super-scientist who has learned all that one can learn from black and white media about the physics and neurophysiology of color and color perception. Having been confined to a black and white room, however, she has never had an experience of seeing red. So while she may know a lot about redness, Mary doesn't know what it's like to see red. Those of us who have experienced redness—and who have the phenomenal concept RED_{ph} in the way that is normal for people with color vision—do know what it's like to see red. So despite her vast knowledge, there's something those of us with less thorough color educations know about seeing red that Mary doesn't know, at least

until she emerges from her room.

Our concern here is not with how to make sense of Mary's ignorance or to determine what implications it might have for physicalism. Instead, we will take interest in the fact that Mary knows what it's like to *not* know what it's like to see red. She knows what it's like to wonder what red looks like and not only not know the answer, but have no idea of what the answer could be.

I, on the other hand, don't know what this is like. And you probably don't either, if you've had visual experiences of redness. Go ahead: try to imagine what it could be like to wonder as Mary does. Imagine you know all sorts of facts about red and how it's perceived, but still wonder what it could be like to experience seeing red, and can't manage to think of any answer that is at all plausible, no matter how hard you try.³

I know what seeing red is like, and it seems I can't fully block out this knowledge when I'm trying to understand the state Mary is in. Moreover, my state of ignorance seems more difficult to eliminate than Mary's. In order to learn what experiencing red is like, all Mary needs to do is have some experiences of seeing red. But how can I ever know what it's like to be ignorant of this, given my knowledge? Short of inducing some very thorough amnesia, knowing what Mary knows seems forever off-limits to me. Mary's friend Barry, on the other hand, who also has never seen red, can know what it's like to wonder what red is like. There's a sense in which he can understand Mary better than I can.

I'll be using cases like these to argue that having a concept often has an epistemic cost. To make the argument, I'll first need to fix some terminology and make some assumptions about concepts.

³This aspect of Mary's situation differs from that of Nida-Rümelin (1995)'s Marianna. Marianna has seen the colors but hasn't been taught their names, so can think of possible answers to the question of what seeing red is like, but doesn't know which is the right one. Unlike Marianna, Mary is in what Bromberger (1965) calls a b-predicament—the answer to her question is something she cannot conceive.

1.1. Some Preliminaries Concerning Concepts

A *concept*, as I will use the term, is an ability to think some content (or contents).⁴ *Thinking* (or what I will sometimes call *grasping*), 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 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 what these somethings are, nor even that there is some account involving anything that could reasonably be called ‘concepts’ appearing in the contents themselves or as elements of cognition.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.

⁴Depending on what contents we take there to be, we may also want to allow that concepts may be *parts* of abilities to think certain contents or abstractions over abilities to think contents. We could take the concept CAT, for example, not to be an ability to grasp any particular content, but instead to be an ability that one exercises in thinking that the cat is on the mat, or that cats are mammals, or that there are at least two cats in the house. We can think of CAT, then, either as a sub-ability involved in thinking these propositional contents, or else an abstraction over the abilities to think contents with cat-related subject matters. On subject matters, see Lewis (1988a,b), Yablo (2014), and Fine (2017). If one accepts that propositions have constituents, then the concept can be understood as an ability to grasp the constituent part corresponding to an unsaturated cathood-predicate (or Fregean concept).

⁵I 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.

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

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 grasp that content. The child doesn't have the concept RATIONAL NUMBER.⁷ 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}.

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.⁸ 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 could possess it before experiencing red.

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.

⁷Development of number concepts has been one of the most thoroughly studied aspects of conceptual development. See Carey (2009b) for a philosophically oriented overview. For more recent work, see Leibovich et al. (2017) and the associated commentaries.

⁸See Sundström (2011) and references therein.

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 grasping. 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 grasp 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.⁹ 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.¹⁰

We could develop this kind of view by holding that while both

⁹An alternative response to the argument is to deny that this is so and explain the apparent truth of Mary's utterance of (1) in another way. We could point out that this case involves an argument from copredication much as does the argument from Chomsky (2000) against referential approaches to semantics, based on sentences like (2).

- (2) The book that he is planning will weigh at least five pounds if he ever writes it.

In such sentences, it seems the very same thing—the book—is being treated as an abstract, informational entity and a concrete entity. But, Chomsky thinks, a serious scientific theory should not countenance strange entities which seem to be both abstract and concrete.

There are various responses to this kind of argument. For some recent ones, see Liebesman and Magidor (2017), Collins (2017), and Gotham (2017).

¹⁰Note 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.

Mary and those who have seen red *possess* the concept RED, only the latter have *mastery* of the concept.¹¹ 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.¹² 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. I would like a characterization, though, 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, 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

¹¹Rabin (2011) and Alter (2013) pursued this line in response to Ball. See Ball (2013) for a response.

¹²Crimmins 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.

accident, and we may want to think of the way he played in the 1936 as different from the way he played in 1949, though we could subsume both under what we might call the 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 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 Reinhardt way, and he could play in the Reinhardt way 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 Reinhardt way. None of this involves having to say anything strange about a sentence like (3), like that it ascribes two different abilities.

(3) 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 grasp 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.¹³ 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 having some concept as possession of different, more specific concepts.¹⁴

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 grasp that content changed after her experience of redness. Mary can now grasp 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

¹³This 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.

¹⁴We may want to revise our definition to make this more explicit, saying that concepts are abilities to grasp contents *in certain ways*.

determinate concept—even if there are not similarly fine-grained contents.¹⁵ Now we can state the argument against Costlessness.

1.2. The Argument, Briefly

(i) There are some pairs of concepts which are incompatible, in the sense that one cannot possess both.

But (ii) these concepts have independent epistemic value.

So, (iii) whichever one someone has, they'll necessarily be missing out on some epistemic value, simply because they have the concept they do.

So, (iv) there is an epistemic opportunity cost to having the concept they have.

So, (v) Costlessness is false.

I'll start by defending (ii), arguing that many concepts have independent value. Then I'll argue that some of these independently valuable concepts are incompatible, so (i) is true. This is enough to get to (v). The remainder of the section addresses three objections.

1.3. The Value of Concepts Redux: Epistemic Empathy

Having a concept, I began by saying, usually has some epistemic benefits. This may have seemed unobjectionable, but we should reconsider now that we're allowing such fine-grained individuation of concepts. Is there any value to having a determinate concept through having some particular more determinate concept? Or are all ways of having some concept equal, and once you have it in one way, you've exhausted all the value that can be wrung from it? I will argue here that, at least in many cases, different more determinate ways of having a concept are independently epistemically valuable.

¹⁵I 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 §1.7.

Let's start with RED_{ph} , a determinate of RED . Does having it provide any epistemic value, beyond what is made available by having RED in some way or other?

A common response to Jackson's argument against physicalism is that when Mary escapes her room and sees red things, she doesn't learn any new facts, doesn't eliminate any possibilities. Nevertheless, the experience of redness does impact her in various ways, for example by giving her new abilities to imagine red and discriminate the red from the non-red by sight. These are the sorts of things which we are reacting to when we feel pulled to say that Mary has learned something.¹⁶

Suppose with her version of the concept RED and her vast knowledge of color perception, there are no new facts she can know by acquiring RED_{ph} . Does this imply that there's no epistemic value in her having experiences of red? Must we hold that as far as epistemic value is concerned, it's a matter of total indifference whether to have an experience of red or not? No, for we need not—and I think should not—hold that epistemic value comes down simply to which facts one knows. There can be independent value in knowing facts in particular ways. It seems that Mary *has* gained something of epistemic value by coming to experience red, even if she doesn't thereby come to know any new facts. She has come to know what red is like in a new way, one that relies on different cognitive abilities, and the value of knowing this in this way is not something she had when she knew this fact in her old, purely descriptive way.

Having a concept in some particular way can have its own value, then, even if it doesn't provide one with access to new facts. But is this something specific to the odd case of Mary? Or does it hold more generally? And what is the source of this value? I cannot develop a general theory of the epistemic value of concept possession here. However, I think we can get what we need for my argument by taking into consideration one important source of concepts' epistemic value: their role in understanding others.

¹⁶This is the view taken by, among others, Nemirow (1980), Lewis (1988c), and Stalnaker (2008). For alternative ways of accepting that Mary learns something without granting that there is some new fact which she has learned, see Nida-Rümelin (2015, §4.5–7).

Mary, before she leaves her room, may know all sorts of things about the red-related attitudes of those who have seen red. 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. Presumably she attributes these attitudes to me in part through reliance on having the concept RED. But the more specific concept she uses for thinking these and other red-related thoughts, call it RED_{Mary}, is rather different from mine, which is a determinate of RED_{ph}. She can access the contents which I have attitudes towards, but her way of accessing them is very different from my way of doing so. Her understanding of me is thus that of a relative outsider, and is in some respect worse than it could be. There's a respect in which fellow RED_{ph}-users understand my red-related thoughts better than Mary does. 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}. Now she really gets my belief that tomatoes are red and my liking my red couch, or at least she gets these attitudes better than she did before. Having RED_{ph} has epistemic value that goes beyond having other determinates of RED, since it allows for a better understanding of RED_{ph}-users. And this seems so regardless of whether RED_{ph} gives one access to any otherwise inaccessible facts.

This isn't a special feature of RED_{ph}. Rather, it seems that there's something valuable in general about attributing thoughts in what I'll call a more *epistemically empathetic* way. The more similar the concepts that one uses to attribute thoughts are to the concepts used by the attributees to have those attitudes, the better one's understanding of them, other things being equal.

Coming to have RED_{ph} is epistemically valuable for Mary partly because it makes possible a more epistemically empathetic understanding of people like me. But we should also make the parallel observation: I don't fully understand *Mary's* wondering what red is like (or, for that matter, her belief that tomatoes are red), but instead attribute these attitudes as a relative outsider. I don't fully understand her any more than she fully understands me. If I had the concept RED in

a way more like hers, I could better understand her. So having RED in Mary's way—RED_{Mary}—also has some epistemic value. It makes possible epistemic empathy with people like Mary.

Similarly with DENSITY, WEIGHT, and WEIGHNSITY. 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. Perhaps an alien super-scientist could know all the facts about humans' attitudes without having anything like these concepts, but their understanding would be lacking in one respect: they would not understand our thoughts 'from the inside' in the way that we can understand each others' thoughts. 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. Even if an adult super-scientist could know all the facts about this concept and about children's use of it, not having WEIGHNSITY will mean they will not be able to think of things as children do, and will only understand them as outsiders. Their understanding may be overall better than that of the children themselves, but again, there seems to be something valuable missing, something that children can understand about each others' thoughts. Allowing greater epistemic empathy with children is a source of epistemic value for WEIGHNSITY that is not captured by DENSITY and WEIGHNSITY.

Epistemic empathy requires having concepts in the ways that other people have them. But people have all sorts of concepts in all sorts of ways. And not only that, but we might take there to be value to being able to understand in an epistemically empathetic way the thoughts of past, future, and even merely possible people. But even limiting ourselves to the actual and present, very many ways of having concepts will have some independent value, since they are each required for epistemic empathy with someone or other. Given the fine-grained account of concepts we're working with, this means that many determinate concepts have some independent value, even when they are determinates of the same determinable.

1.4. Incompatible Concepts

Could I learn Mary's way of thinking of red or the children's way of thinking of heaviness? Could I acquire RED_{Mary} without losing access to RED_{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. When Mary thinks this content, 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.¹⁷ When I consider this content, however, I have easily within reach a variety of full answers that I grasp with my phenomenal color concepts, including, of course, the correct answer, which I grasp with RED_{ph} . I still won't really get Mary's ignorance, except in a distanced, impersonal way.

What's getting in the way of epistemic empathy with Mary seems to be my having the concept RED_{ph} . And this is something which I could not change by book-learning *or* experience. It seems that 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 RED_{ph} concept, and perhaps also $BLUE_{ph}$, $GREEN_{ph}$, and so on. It seems that my way of having the concept RED is incompatible with the way Mary has it. There is no addition to

¹⁷By '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'.

my conceptual scheme which can fully subsume both the ways I can think red-thoughts and the ways Mary can think red-thoughts. Our determinates of RED seem to be incompatible.

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 childrens' 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.¹⁸ Or consider the ability to make a table made only by Alvin Plantinga and the ability to make a table made only by someone other than Alvin Plantinga. Nobody could have both of these abilities, since Alvin

¹⁸There 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.

Plantinga is the only person who could have the former ability, but he is of course barred from having the latter.¹⁹ 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.²⁰

¹⁹See Plantinga (1967, pp. 169–170) for an example like this.

²⁰*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 salt has to dissolve in water.

Reply: (1) 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 thought that tomatoes are red, didn't you already have to think that content in order to try?); (2) 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.

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 RED_{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.

1.5. Self-restraint?

Having the concepts RED_{ph} or $DENSITY$, I've claimed, block epistemic empathy with someone who thinks with RED_{Mary} or $WEIGHNSITY$, since having the latter involve lacking abilities that the former require. But why couldn't someone with RED_{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.²¹ Grasping a content with some self-imposed restrictions

²¹Indeed, 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.

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 grasps 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 RED_{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.

It isn't enough, though, just to find *some* respect in which the ways the concept restrictors and the concept lackers grasp some thought differs. Not every respect in which concepts might be said to differ matters for epistemic empathy. Having been within 50 feet of Banksy, for example, is not an aspect of one's way of thinking of redness that is relevant for the epistemically significant kind of similarity we have in mind. So even if we can define a determinate of RED that involves this—the way of thinking of red while having had been within 50 feet of Banksy—it will not mark any dimension of similarity that matters for epistemic empathy.

How can we tell which respects of similarity are relevant for epistemic empathy? It seems to me the best we can do is consider various cases and judge whether the respects of similarity or difference in those cases seem to give one a better or worse understanding of someone 'from the inside'. Perhaps with further investigation we can construct an informative theory of epistemic empathy that will have enough virtues to warrant overturning some of our initial judgements. In the meantime, though, there's not much else we can do beyond thinking about the cases one-by-one.

The difference between $RED_{near-Banksy}$ and $RED_{not-near-Banksy}$ seems not to matter for epistemic empathy, whereas the difference between RED_{ph}

and RED_{Mary} does. But what about the one between RED_{Mary} and someone who has RED_{ph} but can restrict access to it for the purposes of interpreting others?

It seems to me that someone who can restrict their access to RED_{ph} can understand pre-revelation Mary from the inside better than someone who has RED_{ph} but cannot restrict access in this way, but still not quite as well as someone who, like Mary, also lacks RED_{ph} . Suppose that like Mary, Barry has grown up in a room without being visually exposed to redness, and so lacks RED_{ph} , whereas Cary normally thinks with RED_{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.

I don't expect to convince everyone to share this judgement, but drawing out a difference between Barry and Cary's understanding of Mary can perhaps help one to see why they might 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 RED_{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 RED_{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 RED_{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 RED_{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. And restricting access again won't help, since it will just leave him in the same position once again.

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 is often taken to be crucially connected to integration with one's other information,²² 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.

1.6. 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 RED_{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.

²²See 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.

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 RED_{ph} instantaneously.

We don't usually acquire concepts in a flash. With children acquiring $WEIGHT$ and $DENSITY$, 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.²³

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.

²³It'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).

Vision scientists have in fact tried depriving non-human primates of normal color 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,²⁴ even after two months of life in normal lighting, Femke still reacted to color tasks in somewhat abnormal ways.²⁵ 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.²⁶ 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.²⁷

I don't take these results to be conclusive, but it seems the relevant empirical work suggests Mary would not acquire RED_{ph} immediately, so we shouldn't expect an immediate loss of RED_{Mary} .²⁸

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

²⁴In contrast, effects of deprivation on acuity and depth perception are typically large and long-lasting.

²⁵See 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.

²⁶In 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).

²⁷See McKyton et al. (2015).

²⁸But even supposing the process of acquiring RED_{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'.

Should we call what happens to Mary once she acquires RED_{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.²⁹ 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 RED_{ph} as similar to this.

1.7. 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 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*) Having some concepts is incompatible with having other concepts in certain ways.

²⁹For other, rather different examples of losing knowledge without any normal kind of forgetting, see Block’s extension of Harman’s Inverted Earth thought experiment in Block (1990, pp. 63–65) and Two Roads to Shangri La from Arntzenius (2003).

But (ii*) having these other concepts in these ways has independent epistemic value.

So, (iii*) if someone has one of these concepts, they will necessarily be missing out on some epistemic value, simply because they have the concepts they do.

From there the argument proceeds as before.

We can defend (ii*) as we defended (ii)—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 concepts, then, Costlessness is false.

2. Some Implications

I've argued that many concepts have epistemic costs—merely possessing them prevents one from being able to attribute thoughts to certain people in epistemically empathetic ways. Suppose this is right. So what?

In this section I draw out some implications of my argument against Costlessness. The most direct implications are for the epistemology of concept change (§2.1), but there are also potentially consequences for philosophical projects which require that everything can be understood, at least in principle, from a single perspective (§2.2).

2.1. Epistemology of Concept Change

Contemporary normative epistemology has been concerned primarily with determining what one should believe or what credences one should have, *given a particular conceptual scheme*. An agent has various doxastic options—all the propositions they could believe or have credence in—and the problem is to determine which options they should pick, given their situation.

But a similar normative epistemic question can be asked about which concepts one should have. On the face of it, conceptual schemes can be epistemically better or worse and changing from one to another can be more or less epistemically warranted.³⁰ The failure of Costlessness likely has interesting implications for this project.

If we take a consequentialist approach to epistemic evaluation of concepts, the implication is clear. Had we accepted Costlessness, we would have thought that, at least idealizing away physical limitations, it is never epistemically irrational to acquire a given concept, since there may be some epistemic benefit to having it, but never any cost.³¹ Once we've rejected Costlessness, though, we cannot assume this is the case. Possession of many concepts will have benefits which outweigh their costs, but it's not guaranteed that this will be so. In particular, if there are concepts that are defective and don't give one access to any useful thoughts about the world, it seems possible that the costs might outweigh the benefits, and that it is epistemically better to lack many such concepts.

To reach more detailed conclusions, we would need to work out

³⁰There has been little work directly on this issue. Important exceptions include Carr (2015), Egré and O'Madagain (2019), and Pérez Carballo (2020). Much work has been done recently on more general evaluation of concepts, going under the heading of conceptual engineering and conceptual ethics. For overviews, see Burgess and Plunkett (2013a,b), Cappelen (2018), and Burgess, Cappelen, and Plunkett (2020). Most of this, though, seems to me only tangentially related, analogous to the implications for normative epistemology from non-epistemic reasons for belief. See Reisner (2018) and Berker (2019) for recent overviews of the literature on whether there are practical reasons for belief and how they are thought to interact with epistemic reasons.

³¹Carr (2015) favors a view like this, though doesn't officially commit to it. It should be said, though, that her view is intended only to concern *pure* epistemic expansions, through which one gains but does not lose any concepts. If the argument of §1 is right, though, there might not be any pure epistemic expansions, at least if we allow for fine-grained individuation of concepts.

how to weigh the benefits of epistemic empathy against other sorts of epistemic benefits. One thing the above argument calls for, then, is further investigation into the epistemic value of understanding others in this way and how to weigh it against other values.

There are straightforward implications, then, for epistemic consequentialist approaches to evaluating concepts. What if we are not epistemic consequentialists about evaluating concepts? Whether having a concept is epistemically warranted may not be determined by the epistemic costs and benefits it has. Indeed, the same sorts of arguments against epistemic consequentialism for belief also apply to concept possession. If a demon offers you access to many important concepts for understanding the world, so long as you have your concept *LAPTOP* eliminated, this doesn't mean that your having the concept *LAPTOP* is unwarranted, despite its being epistemically better, overall, to lack the concept.³² I do not want to take a stand here on whether such arguments count conclusively against epistemic consequentialism. They do suggest, though, that it's worth thinking about what implication, if any, the argument against Costlessness could have for a non-consequentialist epistemology of concepts.

First, note that even if we are not consequentialists about first order questions about epistemic warrant of concepts, we may still take a teleological approach on a foundational level.³³ Costs and benefits in particular cases won't determine whether a concept is warranted, but facts about costs and benefits may still have an important indirect role in determining what concepts are warranted. The fact, then, that Costlessness is false may well be relevant to the epistemology of concepts even without assuming some kind of direct epistemic consequentialism. Exactly what implications it would have would turn on just what indirect role epistemic value plays. One can imagine various options here, paralleling those familiar from ethics and from

³²This kind of trade-off argument is a direct parallel to those made against epistemic consequentialism about belief made by Firth (1978), Berker (2013), and others. For epistemic consequentialist responses to these arguments, see Goldman (2015), Wedgwood (2018), Driver (2018), and Singer (2018).

³³On this distinction in ethics, see Kagan (1998, Ch. 6). It should be noted that this is an issue of which Berker and his respondents are well aware.

epistemology of beliefs. In pretty much any case, though, the fact that most concepts will have some built-in cost will have some relevance.

On various accounts, even non-consequentialist ones, it may turn out that a concept's role in allowing for epistemic empathy is one of the factors that contributes to epistemic warrant for a concept. And so if we accept my claim above that many concepts are incompatible, agents may often face situations in which there is some degree of warrant for possessing multiple, incompatible concepts. We will need to determine, then, what our normative epistemology of concepts should say about such cases. Should it be permissible to have either? Neither? Or must we always possess the concept that would be most warranted?

In one way or another, then, it seems likely that the failure of Costlessness, the benefits of epistemic empathy, and the existence of incompatible concepts will have some role to play in determining what concepts we epistemically ought to have.³⁴ And this will apply even if we idealize our contingent computational limitations.

2.2. Incompatible Concepts and Idealized Thinkers

Philosophers often appeal to idealized thinkers and sometimes it's important to their projects that an idealized thinker would be able to know everything that there is to know about some subject matter. One particularly ambitious such project is that of Chalmers (2012).

Chalmers's view, roughly put, is that every truth is 'scrutable' from certain base truths. That is, there is a relatively compact class of truths from which *all* truths can be *a priori* inferred. If an appropriately idealized agent knew all the base truths, they could know all the truths there are.³⁵

³⁴Besides have implications for individual epistemology of concepts, the argument may well matter for social epistemology of concepts, though this is a topic currently too little explored for me to be able to say much about it. It might be that the value of epistemic empathy within the group counts in favor of group members' having very similar conceptual repertoires. But it might also be that a group of conceptually diverse thinkers can, in aggregate, understand more than any individual or conceptually homogeneous group could understand, which could count in favor conceptual diversity within the group.

³⁵In order to avoid issues related to Fitch's paradox, the more refined version of Chalmers view is one of *conditional* scrutability: there is a compact class of base truths *C* such that an ideal agent could know that *if* the members of *C* are true, then *S* is

But at least if we are taking truths to be relatively fine-grained, such that through having different determinate concepts one can know different truths (like what it is like to see red, or what it is like to wonder what red is like), then the claim—argued for in §1.4—that some concepts are incompatible would mean that it is in principle impossible to know *all* the truths there are.³⁶ No agent, even among very idealized agents, could know all the truths there are. So they couldn't know them through inferring them from a compact class of base truths, perhaps because they couldn't reach the conclusions from the base truths without acquiring concepts which prevent them from reaching other conclusions, or perhaps because they couldn't grasp all the base truths at once.³⁷

true, where S is the set of all sentences. See Chalmers (2012, pp. 53–58). The problem I will raise applies to scrutability theses in either conditional or non-conditional form.

³⁶One might still hold that Chalmers's scrutability theses are true, just vacuously so. I take it, though, that this would be unsatisfactory for Chalmers's project.

³⁷Chalmers does briefly consider and reply to an objection from incompatible concepts:

Suppose that for two incompatible properties (being intelligent and unintelligent, perhaps), one has to have the property in order to possess a phenomenal concept of what it is like to have the property. Then no one can possess both concepts at once, and if both are involved in Q , no one can entertain all of PQI . In response, I think that there are no concepts whose possession is mutually incompatible. The supposition just mentioned cannot obtain: in these cases there is usually no bar to having a phenomenal concept of what it is like to have a property through imagination even without having the property. In any case, this worry will not affect scrutability from PQI when this is understood in terms of a warrant-based idealization (see 2.7 and E4) rather than a modal idealization.

(Chalmers 2012, p. 114)

So Chalmers's reply has two parts. First, that there are no incompatible phenomenal concepts, since one can typically imagine having some property without actually having it. Second, if we define scrutability not with a modal idealization, but with a warrant-based idealization, there will be no problem from incompatible concepts.

On the view developed above, it is crucial to Mary's concept that she *can't* imagine the kinds of answers which anyone with my concepts can imagine. If we accept this, then there really are incompatible concepts. I can't entertain both the thoughts I have using RED_{ph} and the thoughts that Mary has using her concept RED_{Mary} . Similarly with thoughts grasped with $DENSITY$ and those with $WEIGHNSITY$.

What about replacing the modal idealization with a warrant-based one? The idea behind the warrant-based idealization is that instead of defining scrutability and other notions in terms of what some idealized agent would do (a modal idealization), we define it in terms of what they would have warrant to do. The conditional scrutability thesis understood in terms of warrant says that the relevant conditional—if $PQTI$,

None of this conclusively refutes any of Chalmers's central claims. One can object that we should not take there to be different truths in these cases, just different ways of knowing the same truths. Or one can follow Chalmers's suggestion that we not think of the relevant claims in terms of modal idealizations, but rather in terms of warrant, so that his thesis could hold even if it's not possible for there to be an agent that grasps all the truths that there are. Since the point of this section is not to go after Chalmers, but to illustrate the kinds of implications the argument of §1, we need not get into such issues here. Instead, we can put things irenically: insofar as scrutability theses (or other philosophical theses) rely on the possibility of grasping truths knowable only through incompatible concepts, they will be false if my argument against Costlessness works.

The implication drawn here here is similar to one which Thomas Nagel has argued for in various places.³⁸ What Nagel aimed to establish is that not everything is knowable from an objective 'view from nowhere'. 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 the full range of possible inner lives" (Nagel 1986, p. 17).³⁹

If my argument from §1 works, it pushes us beyond Nagel's skept-

then *M*—has a priori warrant. That is, it has a priori 'ideal' justification, in a sense that does not require that it is possible to actually believe it. Supposing we allow that there can be ideal warrants of this kind, this would avoid the problem from incompatible concepts. That said, it's not clear Chalmers would be home free at this point, since it's not clear that what is warranted to believe will be independent of what concepts one is warranted in having. So if there are truths that are only graspable with unwarranted concepts, it may be that belief in these truths wouldn't be warranted, and so not scrutable on the warrant-based understanding of scrutability.

³⁸See Nagel (1974), Nagel (1986), and Nagel (1997, Ch. 2).

³⁹This idea is similar to the phenomenal simulation part of Chalmers's *Cosmoscope*, see Chalmers (2012, p. 115).

ticism 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 often thereby loses what's required to see from another.

3. Two Concluding Perspectives

I have argued that Costlessness is false and laid out a few implications my argument would have, if successful. Depending on one's temperament, one might find some of these implications disappointing. Someone might take it to be their epistemic duty to understand more and more of the world, and, in the idealized limit, to fully understand everything. But concepts play a dual role; besides understanding the world in an objective way, they also allow us to understand other thinkers in their own terms. The latter role, we've seen, can pull us in different directions. No individual, not even a god, can access all the epistemic value there is.⁴⁰ From one perspective, it is tragic that what it seemed like we should be aiming for is in principle impossible.

Another perspective takes this not to be tragic, but liberating. We need not take it upon ourselves to understand everything, since that task is impossible. What's more, the very limits of our understanding are valuable, for they allow for an important way of understanding other, similarly limited thinkers.

⁴⁰For discussion of other cases of apparent epistemic ideals turning out to be unsatisfiable, see Christensen (2007) and Pasnau (2017).

References

- Alter, Torin (2013). "Social Externalism and the Knowledge Argument". In: *Mind* 122, pp. 481–496.
- Arntzenius, Frank (2003). "Some Problems for Conditionalization and Reflection". In: *The Journal of Philosophy* 100.7, pp. 356–370.
- 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.
- Berker, Selim (2013). "Epistemic Teleology and the Separateness of Propositions". In: *Philosophical Review* 122.3, pp. 337–394.
- (2019). "A Combinatorial Argument against Practical Reasons for Belief". In: *Analytic Philosophy* 59.4, pp. 427–470.
- Block, Ned (1990). "Inverted Earth". In: *Philosophical Perspectives* 4, pp. 53–79.
- 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.
- Bromberger, Sylvain (1965). "An Approach to Explanation". In: *Analytical philosophy*. Ed. by R. J. Butler. Vol. 2. Blackwell Oxford, pp. 72–105. Reprinted in Bromberger (1992).
- (1992). *On What We Know We Don't Know: Explanation, Theory, Linguistics, and How Questions Shape Them*. University of Chicago Press.
- 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.
- Cappelen, Herman (2018). *Fixing Language*. Oxford University Press.
- 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.
- Chomsky, Noam (2000). *New Horizons in the Study of Language and Mind*. Cambridge University Press.
- Christensen, David (2007). "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals". In: *Oxford Studies in Epistemology*. Ed. by Tamar Szabó Gendler and John Hawthorne. Vol. 2, pp. 3–31.
- Collins, John (2017). "The copredication argument". In: *Inquiry* 60.7, pp. 675–702.
- 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.
- Driver, Julia (2018). "The "Consequentialism" in "Epistemic Consequentialism"". In: *Epistemic Consequentialism*. Ed. by Kristoffer Ahlstrom-Vij and Jeffrey Dunn. Oxford University Press, pp. 113–122.
- 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.
- Fine, Kit (2017). "A Theory of Truth-Conditional Content II: Subject-matter, Common Content, Remainder and Ground". In: *Journal of Philosophical Logic* 46, pp. 675–702.
- Firth, Roderick (1978). "The Schneck Lectures, Lecture 1: Epistemic Utility". Delivered at the University of Arizona. Printed in Firth (1998).
- (1998). In *Defense of Radical Empiricism: Essays and Lectures*. Ed. by John Troyer. Rowman and Littlefield.
- 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.
- Goldman, Alvin I. (2015). "Reliabilism, Veritism, and Epistemic Consequentialism". In: *Episteme* 12.2, pp. 131–143.
- 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.
- Gotham, Matthew (2017). "Composing Criteria of Individuation in Copredication". In: *Journal of Semantics* 34, pp. 333–371.
- 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.
- Harman, Gilbert (1986). *Change in View: Principles of Reasoning*. MIT Press.
- Jackson, Frank (1982). "Epiphenomenal Qualia". In: *The Philosophical Quarterly* 32.127, pp. 127–136.

- Jackson, Frank (1986). "What Mary Didn't Know". In: *The Journal of Philosophy* 83.5, pp. 291–295.
- Kagan, Shelly (1998). *Normative Ethics*. Westview Press.
- Keene, G. B. (1960). "A Simpler Solution to the Paradox of Omnipotence". In: *Mind* 69.273, pp. 74–75.
- 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.
- Leibovich, Tali et al. (2017). "From "sense of number" to "sense of magnitude": The role of continuous magnitudes in numerical cognition". In: *Behavioral and Brain Sciences* 40.
- 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.
- (1988a). "Relevant Implication". In: *Theoria* 54.3, pp. 161–174.
- (1988b). "Statements partly about observation". In: *Philosophical Papers* 17.1, pp. 1–31.
- (1988c). "What experience teaches". In: *Proceedings of the Russellian Society* 13, pp. 29–57.
- Liebman, David and Ofra Magidor (2017). "Copredication and Property Inheritance". In: *Philosophical Issues: A Supplement to Noûs* 27, pp. 131–167.
- 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.
- Nemirow, Laurence (1980). "Review of *Mortal Questions* by Thomas Nagel". In: *The Philosophical Review* 89.3, pp. 473–477.
- Nida-Rümelin, Martine (1995). "What Mary Couldn't Know: Belief about Phenomenal States". In: *Conscious Experience*. Ed. by Thomas Metzinger. Imprint Academic Exeter, pp. 219–41.
- (2015). "Qualia: The Knowledge Argument". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2015. Metaphysics Research Lab, Stanford University.
- Nunberg, Geoffrey (1984). "Individuation in context". In: *Proceedings of the West Coast Conference on Formal Linguistics*. Vol. 2, pp. 203–217.
- Pasnau, Robert (2017). *After Certainty: A History of Our Epistemic Ideals and Illusions*. 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.
- Reisner, Andrew (2018). "Pragmatic Reasons for Belief". In: *The Oxford Handbook of Reasons and Normativity*. Ed. by Daniel Star. Oxford University Press, pp. 705–728.
- Schwenkler, John (2013). "Do things look the way they feel?" In: *Analysis* 73.1, pp. 86–96.
- Singer, Daniel J. (2018). "Permissible Epistemic Trade-Offs". In: *Australasian Journal of Philosophy*.

- Staffel, Julia (2020). *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford University Press.
- Stalnaker, Robert (2008). *Our Knowledge of the Internal World*. 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.
- Wedgwood, Ralph (2018). "Epistemic Teleology: Synchronic and Diachronic". In: *Epistemic Consequentialism*. Ed. by Kristoffer Ahlstrom-Vij and Jeffrey Dunn. Oxford University Press, pp. 85–112.
- Yablo, Stephen (2014). *Aboutness*. Princeton University Press.