

15

Conceptual Evaluation

Epistemic

Alejandro Pérez Carballo

My topic here is the question of conceptual evaluation along an epistemic dimension. My question is: what makes for good concepts, epistemically?

Thus formulated, this question cries out for clarification in a number of ways. And I will turn to that in due course. But it will help to get some sense of where we're going before getting started.

The distinction between epistemic and non-epistemic evaluations of belief is not only familiar, but rather intuitive. In part, I suppose, that is because we have a reasonable grasp of ways of evaluating beliefs—in terms of their truth, or their warrant—that are plausibly classified as 'epistemic'. There is no similar, pre-theoretically available way of evaluating concepts that is uncontroversially classified as epistemic. Still, it seems uncontroversial that scientific progress often involves introducing new conceptual tools. And this suggests that scientific progress sometimes involves conceptual progress—that better science incorporates better concepts.

What could make some concepts epistemically better than others? On what is perhaps the orthodox answer to this question—the answer that most philosophers I have encountered are inclined to give, once they get past their reluctance to take talk of concepts at face value without some heated preliminary discussion—the issue turns on what properties correspond to the concepts in question. A concept is better than another, epistemically, if and to the extent that the property corresponding to one is 'more natural' than the one corresponding to the other. This line is often accompanied by a quick reference to Goodman—'you mean like how "blue" is a better concept than "grue"?'—or to some discussion of natural kinds.

My proximate aim in this chapter is to argue that this answer cannot be right. More generally, I want to argue against a particular way of addressing this question—an approach that leans heavily on metaphysical, non-epistemic notions like 'structure' or 'naturalness'. But the ultimate goal is to make some progress in clarifying what a theory of conceptual evaluation should hope to accomplish.

As it happens, I am skeptical that there is a theoretically useful, non-epistemic notion of 'naturalness' or 'structure' out there. But here I will simply take for granted that the relevant metaphysical notions are in good order. In short, I want to grant for the purposes of this chapter that there is a theoretically useful notion of 'naturalness'

according to which some properties are natural and some are not—or better yet, that some properties are more natural than others. (I will assume, however, that the relevant notion of ‘naturalness’ is either a primitive notion or can be defined in purely non-epistemic terms.) My claim will be that there is no straightforward connection between a theory of epistemic evaluation of concepts and a theory of ‘natural’ properties thus understood.

1 Preliminaries

A number of authors have recently made explicit claims in what we can call ‘applied conceptual ethics’:¹ claims to the effect that some concept is better—in some sense or other—than another. Sally Haslanger, to take a well-known example, has defended a number of theses about “how we might usefully revise [our concepts of race and gender] for certain political and theoretical purposes” (Haslanger 2000: 34). On her view, there are legitimate purposes that will be better served by our using ‘woman’ so that (very roughly) *x is a woman* conceptually entails *x is subordinated (along some dimension or other)* (p. 39ff). Similarly, Kevin Scharp has argued that we should ‘replace’ our concept of truth with two concepts which, together, are better than our (allegedly inconsistent) ordinary concept for the purposes of ‘serious theorizing’ (Scharp 2013: 134).

These are not isolated cases, mind you. Peppered throughout the literature on the so-called species problem² in philosophy of biology one often finds claims to the effect that a particular species concept is better than another for certain theoretical purposes.³ And one only needs to look back at different applications of the Carnapian method of explication to see that evaluative claims about concepts and vocabularies are a fairly common occurrence in contemporary philosophy.⁴

Often, as in the examples above, such comparative claims are relativized to specific purposes. The presumption then is that a concept cannot be good or bad *simpliciter*. But this is by no means the only game in town. Ingo Brigandt, for example, has argued that many concepts should be understood as having a built-in *epistemic goal*, which “consists of those theoretical, explanatory, or investigative aims that are pursued by this concept’s use” (Brigandt 2010: 36). For example, according to Brigandt, the explanation of patterns of inheritance is an epistemic goal *constitutive* of the classical gene concept—the gene concept associated with Mendel’s work and to the research program led by Thomas Hunt Morgan (Morgan et al. 1915). Note that the claim here is not just that the concept was introduced in order to predict inheritance patterns. Rather, the claim is that a concept that is not put to use for the purposes of predicting inheritance patterns would not *be* the classical gene concept of Mendelian genetics.⁵ If Brigandt is right—and here I do not take a

¹ The terminology here is borrowed from A. Burgess and Plunkett (2013a).

² Kitcher (1984); Mayr (1996); Ereshefsky (2008); *inter alia*.

³ See, for example, Franklin-Hall (2007) and references therein.

⁴ For a more systematic overview of some of the relevant literature, see A. Burgess and Plunkett (2013a,b). For illuminating discussion of Carnap’s ideal of explication and its relation to some of the work on the species problem in philosophy of biology, see Kitcher (2008).

⁵ This at least is suggested by some of his remarks. For example, Brigandt (2010: 36): “each theoretical concept, at least each central biological concept, consists of three components of content: (1) the concept’s

stand on that—then one can make sense of a non-instrumental way of evaluating concepts: we can evaluate how good a concept is relative to its constitutive epistemic goal.⁶

In this chapter, however, I do not aim to advance or take issue with specific claims about which concepts are better than others. Nor do I wish to discuss whether conceptual evaluation must always take place against the backdrop of a specific purpose. My concern, rather, is with an explanatory question that, as far as I can tell, has received little attention in the literature.⁷ In brief, the question is: what makes for good concepts, epistemically? But let me say something about each of the key terms that figure in it, to get a better sense of what an answer to our question will look like.

1.1 ‘Makes’

Take any specific claim about some concept being better than another (whether or not *better than* is relativized to a specific purpose). To a first approximation, the question I’m interested in takes the following form: in virtue of what is the one better than the other? Or: what is it *because of which* one is better than the other?⁸

Now, it might appear as if the answer to this question is straightforward, at least if we are interested in comparisons among concepts relative to a specific purpose. For if what we want is to explain why a concept is better than another relative to a given purpose, the answer should be because the one is a better *means* to achieving said purpose than the other. In other words, if conceptual evaluation is always relative to a goal or purpose—whether or not that goal is constitutive of the relevant concepts—it is merely a matter of instrumental evaluation. And what explains how good something is relative to a given purpose is just that it contributes, in the right sort of way, to that purpose.

True, on this way of thinking there would still be room for disagreement as to what ‘the right sort of way’ is. Should we think that what makes a concept better than another for a given purpose is that using the one is more likely to causally bring about

reference, (2) its inferential role, and (3) the epistemic goal pursued with the concept’s use”. See also the discussion in Brigandt (2003), where he argues that different branches of biology are best understood as having different homology concepts, partly because the homology concepts used in different branches are put to distinct theoretical uses.

⁶ See Thomasson (Chapter 21, this volume) for an alternative take on this strategy—if concepts have constitutive functions, one can evaluate concepts in terms of how well they serve the relevant functions.

⁷ Admittedly, it is not entirely clear whether the explanatory question I will be interested in is entirely neutral on the more ‘first-order’ questions about which concepts are better than others. The analogous question about the relationship between meta-ethics and normative ethics has been the subject of some interesting discussion recently (see, e.g., Enoch 2011; Fantl 2006). I do not know what the outcome of that debate will be. But I expect whatever considerations bear against (or in favor) the claim that our answers to first-order normative questions can be constrained by answers to questions in metaethics can also be brought to bear on the question whether the explanatory question about conceptual evaluation puts any constraints on the first-order debate. That said, I expect that, whether or not these projects constrain one another, many first-order claims about some concept being better than another will be compatible with different answers to the explanatory question I will be interested in.

⁸ Cf. Fine (2012); Rosen (2010). For illuminating discussion of the role of such *because* claims in normative theory and their relationship to the more ‘metaphysical’ discussion of grounding and the *in virtue of* relation, see Berker (2018).

the achievement of that purpose than using the other one? Or should we instead think of the relevant notion of contribution in non-causal terms?

But it would be a mistake to think that in relying on a given purpose in order to compare different concepts we are thereby presupposing some type of teleological story of what makes a concept better than another. After all, one could think of the relevant purposes as determining a particular dimension of evaluation, rather than a goal or something to be promoted. This would leave it open whether the relevant purposes are being used to pick out the particular ranking of concepts or whether they are doing something more—for example, whether they are part of what explains why the concepts are ranked in that particular way.⁹

This becomes especially apparent when we focus, as I will, on questions about conceptual evaluation for *theoretical purposes*. I would be surprised if authors writing about a concept being better than another for theoretical purposes were resistant to being paraphrased as writing about evaluation of concepts from an epistemic perspective. And we can clearly make sense of the possibility of evaluating beliefs, say, from an epistemic perspective without taking on a consequentialist or teleological view on epistemic rationality.

Our question—what makes for good concepts, epistemically—can then be reformulated as: what is it in virtue of which some concepts are epistemically better than others? Or: what explains why some concepts are better than others, epistemically?

1.2 ‘Concepts’

It should not surprise you that philosophers can mean very different things by ‘concept’. And it should surprise you even less that cognitive scientists and philosophers also use the term in importantly different ways—the former use it to refer to so-called mental representations, the latter (often) to mean what gives such representations their representational import.¹⁰ It is hard to expect much agreement, then, on questions of conceptual evaluation unless we specify at the outset what we’re taking concepts to be.¹¹

Now I have said little about what I take concepts to be. This is by design—I intend to remain as neutral as I can on the nature of concepts. The terrain surrounding these issues is slippery: one cannot enter it without getting entangled with questions about modularity, the ‘grain’ of content, and so on.¹² Still, I cannot avoid making some assumptions about what concepts are. My hope is that the assumptions I will make are rather uncontroversial. Fortunately, as we will see, for the purposes of this chapter

⁹ Here, it is helpful to make an analogy with the debate over whether all moral theories admit of a ‘consequentialist’ counterpart. One could grant that any deontological theory can be formulated as a version of consequentialism without thinking that the consequences of a given action play a role in explaining why it has the moral properties that it does. See Portmore (2009) and references therein. Cf. also the discussion in Hammond (1988); Stalnaker (2002).

¹⁰ See the Introduction to Margolis and Laurence (1999) for a guided tour of the issues.

¹¹ Things would be far worse, I suspect, if our interest were in questions of conceptual *engineering*: of creating and changing concepts. For those questions, it is hard to see how to make much progress without substantive assumptions about what concepts are. After all, on some views (e.g. Fodor 1975), changing a concept is entirely out of the question. Here, however, I will have nothing to say about conceptual change—my focus will be exclusively on conceptual evaluation.

¹² For a helpful overview, see Margolis and Laurence (2014).

it will turn out not to matter how we answer many of the questions about the nature of concepts.

First, I will be assuming that concepts are the sorts of things that words mean. This need not imply, of course, that every concept corresponds to the meaning of some word in English, nor that the meaning of every word in English is given by a concept. But I intend talk of concepts to be substitutable throughout with talk of meanings, so that, for example, talking about the epistemic merits of the concept *species*¹³ is tantamount to talking about the epistemic merits of what the word ‘species’ means. Second, and partly as a consequence of our first assumption, I will presuppose that concepts are not mental representations but rather what would give such putative representations their representational import. In talking about concepts, then, I need not presuppose anything like the Language of Thought Hypothesis: as long as it makes sense to talk about the meaning of a word, everything I will say is neutral on questions about cognitive architecture. Third, I will assume that concepts have non-trivial representational content, and that the epistemic merits of a given concept are at least in part determined by their representational content. I am thus ruling out for present purposes views on which, for example, some of our ‘predicative’ concepts do not correspond to a property (except perhaps in a pleonastic sense), at least if the properties in question are supposed to do any explanatory work.¹⁴

1.3 ‘Good, Epistemically’

I’m interested in a particular type of evaluation of concepts, what I’ve been calling *epistemic* evaluation. I want to say a bit more about what I have in mind by that.

I think we can get a handle on what epistemic evaluation of concepts amounts to by taking seriously the idea that better science by and large incorporates better concepts. The hope is that any grip we have on the notion of *epistemically* better science will allow us to get a handle on a way of evaluating concepts that is uncontroversially epistemic in nature. For our purposes, then, a broadly consequentialist strategy will suffice. I will assume that concepts that are conducive to achieving what are clearly epistemic aims and purposes are better, epistemically, than those that are not.

What are the relevant goals and purposes? I want to leave that open. I will only assume that our epistemic goals include forming theories that have the familiar theoretical virtues: accuracy, explanatoriness, fruitfulness, etc. How exactly each of these should be cashed out is a good question, but one I will set aside for present purposes.

At any rate, remember that my question is not: *which* are the good concepts, epistemically? Here I will take for granted that we have a reasonably good sense of

¹³ Throughout, I will use italics to name concepts, relying on context to disambiguate from other uses. I reserve the use of small caps for labeling theses, like SIMPLE below.

¹⁴ I have in mind expressivist or non-cognitivist treatments of many ‘predicative’ concepts (views like those in Gibbard 1990; Blackburn 1998; Yalcin 2011; Pérez Carballo 2016; among others). Presupposing these views away only makes my work here harder: after all, if we hoped for a unified theory of conceptual evaluation, and some concepts that should fall under the scope of such a theory cannot be understood as determining or corresponding to a property, it would be a mistake to build our theory of conceptual evaluation on facts about the corresponding properties.

which concepts best help us achieve our epistemic goals. My question is rather: what is it about those concepts that makes them good concepts, epistemically? What, in other words, is it that explains the fact that they are the most conducive to our achieving our epistemic goals? (What is it *in virtue of which* good concepts are good?)

2 The Simple Hypothesis

What I will call the *simple hypothesis* is perhaps the first answer to our question that comes to mind. The hypothesis starts out with a metaphor. The rough idea is that some properties, unlike others, correspond to true ‘joints’ of nature. This can be motivated in more than one way, but here is one that is relatively easy to state.

Suppose you think that for any set of things—including mere *possibilia*, if you like—there is a property corresponding to that set: elements of the set are exactly those things that instantiate the relevant property. (You might think that properties *just are* such sets, but you need not.) Then you will have to agree that any two things—really, *any* two things—have uncountably many properties in common, and thus (or so it seems) are as alike to one another as any two other things. (A few missing premises are left implicit in the preceding—not because they are unimpeachable, but because I do not intend to argue that the solution to this problem I sketch below is the only way out.) And that, we can agree, would be madness: staplers and raccoons, to pick just one pair, have little in common, or at any rate much less in common with one another than two peas in a pod.

As is well known, you can avoid this conclusion if you insist that, among the uncountably many ways of grouping things together, a relatively small class of them has special status—and correspondingly, a small class of properties has special status. It is the sharing of one or more of *those* properties that makes for genuine similarity. Peas in a pod share many more of those special-status properties—and are thus much more like one another—than a stapler and a raccoon do.

The idea of a metaphysical distinction among properties—‘natural/non-natural’, ‘sparse/abundant’, ‘joint-carving/gerrymandered’, and so on—can be fleshed out in multiple ways.¹⁵ But at this point we have enough materials to formulate a schematic version of the simple hypothesis (I will use ‘elite’ as a stand-in for any one of ‘natural’, ‘sparse’, ‘joint-carving’ and so on):

SIMPLE (FIRST PASS): Concepts are better, epistemically, because and to the extent that they correspond to elite properties.

Again, it is worth highlighting the fact that the notion of ‘eliteness’ is supposed to be non-epistemic in nature.¹⁶

¹⁵ Armstrong (1978); Lewis (1983); Schaffer (2004); Sider (2013). See also Dorr and Hawthorne (2013) and references therein.

¹⁶ The claim thus is not merely that what makes concepts better than others, epistemically, is a purely ‘objective’ matter. (Exactly what this means is a good question, but glossing it with the familiar cluster of metaphors will have to do for now. Cf. Rosen 1994.) Rather, the presumption is that the notion of ‘eliteness’ is, in a sense that would need to be made precise, a purely metaphysical one.

This simple hypothesis has at least two considerable virtues (beyond simplicity, that is). First, it seems to get something right about our pretheoretic judgments about the comparative merits of some concepts. Intuitively, there is something epistemically better about our concept *green* over the concept *grue*. And it certainly *seems* as if this is due to the fact that whereas all green things genuinely have something in common, not all grue things do. Second, the simple hypothesis offers a reduction of one vexed question—which concepts are epistemically better?—to a question which, if no less vexed, is at least supposed to be independently pressing. We need, the story goes, the elite/non-elite distinction for other purposes—for giving an account of Laws of Nature and for giving a theory of intentionality, to name a few. It would be beautifully economical if we could also appeal to it in order to give an account of conceptual evaluation.

Of course, the simple hypothesis, as stated, cannot be right. The notion we are after is a comparative notion—that of a concept being epistemically better than another. And as long as we think, as I think we must, that not all good concepts are equally good, we will want to make finer distinctions than those afforded by a binary elite/non-elite distinction.

On the face of it, this is not an insurmountable obstacle. For one might in principle define a suitable relative of ‘eliteness’ that is comparative in nature. This would essentially involve specifying a ‘more elite than’ relation that applies to properties, and replacing SIMPLE (FIRST PASS) with:

SIMPLE: A concept is epistemically better than another if and because the property corresponding to the former is more elite than the one corresponding to the latter.

(Again, on this picture, the ‘more elite than’ notion would need to be specified in purely non-epistemic terms. The goal, recall, is to reduce the question of epistemic evaluation of concepts to questions about the purely metaphysical merits of the corresponding properties.)

But already at this level of abstraction, we run into difficulties.¹⁷ The first and most straightforward one is this. Elite properties, we are sometimes told, make up a minimal supervenience base for all facts.¹⁸ As a result, high-level properties will not be elite properties. Now, this need not mean that all high-level concepts are equally bad. The ‘more elite than’ relation should presumably make distinctions among non-elite properties. But it does mean that the concepts of fundamental physics are epistemically better than the concepts of, say, psychology. And while

¹⁷ Admittedly, many of the problems I will flag in the remainder of this section arise because of the decidedly Lewisian understanding of eliteness I am working with. An alternative gloss on the relevant notions—one that takes seriously the lessons from Fodor (1974) (see also Antony 2003: esp. p. 12f)—would allow for perfectly elite properties not only at the most fundamental level (say, the level of fundamental physics), but also at the level of chemistry, biology, etc. Cf. Schaffer (2004), and what he labels the ‘scientific conception’ of elite properties: “On the scientific conception, the properties invoked by total science are ontologically on par. All carve out joints of nature. Muons, molecules, minds, and mountains are in every sense equally basic” (p. 94). Ultimately, I will grant the proponent of SIMPLE that there is some conception of eliteness that is in reasonably good order, so I will set aside views like Fodor’s for the sake of brevity. As should become clear, my arguments in §3 apply equally to both a Lewisian and a Fodorian view of eliteness.

¹⁸ See, for example, Lewis (1984, 1983). For an argument to the contrary, see Eddon (2013). Cf. also Dorr and Hawthorne (2013: 10–13).

one might be willing to grant that there are some—or even *many*—epistemic dimensions along which the concepts of fundamental physics are epistemically better than those of psychology, it is hard to believe that there is no epistemic dimension along which psychological concepts are not worse than those of fundamental physics.

The second, related difficulty is a bit less straightforward, for it can only be formulated against particular assumptions about the fine structure of the ‘more elite than’ relation. At a very abstract level, the worry is that even if we grant that concepts corresponding to perfectly elite properties are better than the rest, it isn’t clear that a plausible way of defining ‘more elite than’ for non-elite properties will mesh with reasonable assumptions about the relative merits of high-level concepts.

It is easier to see this with a specific proposal for how to understand the ‘more elite than’ relation. Suppose that we think, as some do, that a property is more elite than another iff its simplest definition in terms of fundamental properties is less complex than the simplest definition of the other (for some suitable notion of simplicity).¹⁹ Consider now your favorite example of a concept of a multiply realizable property—say, the concept of pain. Presumably, the corresponding property will admit of a characterization as a disjunction, each of whose disjuncts corresponds to a way in which the property could be realized. It follows now from our toy theory of ‘more elite than’ that a concept for the property obtained by removing one or more disjuncts from the initial property will be epistemically better than the initial concept. So, the concept of pain-in-carbon-based-beings, for example, would turn out to be epistemically better than the concept of pain. And this, surely, should cast doubt on the current proposal.²⁰

Now, all of this might be avoided if instead of starting with a primitive elite/non-elite distinction we start with a comparative notion specified in non-epistemic terms.²¹ We could, for example, start with the assumption that there are objective facts about genuine similarity—for example, the fact that a horse and a rhino are more similar to one another than a raccoon and a stapler are to one another. And we could use these in turn to get both an elite/non-elite distinction *and* a notion of ‘more elite than’ which may well avoid some of the earlier difficulties.

For example, once we’re allowed access to facts about objective similarity, we can define a ‘more elite than’ relation as follows: property *F* is more elite than *G* iff the most dissimilar *F*s are more similar to one another than the most dissimilar *G*s.²² Perfectly elite properties, on this view, would just be the properties that are maximally elite.²³

¹⁹ Lewis made something like this suggestion, in passing, in Lewis (1984: 228). For a discussion of some of the difficulties involved in carrying this out in reasonable detail, see Sider (2013: section 7.11.1).

²⁰ A more subtle example, borrowed from Hall (2011): *being methane* is a much less useful concept, epistemically, than *being a saturated hydrocarbon*, but the ‘canonical’ definition of the latter will be much more complicated than the one of the former. Cf. Hall (2011: 21ff).

²¹ Or if, following some speculative remarks in Sider (2013), we also allow for primitives other than ‘elite’ in our definition of ‘more elite than’, for example, *law-like*.

²² Cf. Rosen (2015: 191).

²³ A notion of similarity could also be appealed to in giving a definition of eliteness directly, by letting elite properties be *convex*, in the sense that whenever *x* and *y* are more like one another than *y* and *z* are, then if *x* and *z* have property, so does *y*. Cf. Gärdenfors (2000: 70f), as well as Oddie (2005: 152ff). For discussion of the issues arising from putative connections between naturalness and similarity, see Dorr and Hawthorne (2013: 21–7).

Still, to the extent that we think that, by and large, a property is more elite than another iff it is more fundamental, we will be hard pressed to vindicate the thought that, at least along some epistemic dimension, concepts from special sciences are better for some purposes.

But let us suppose, for the sake of argument, that we can get an account of eliteness that is not compromised from the outset. The question is whether we can build a plausible theory of conceptual evaluation on top of it.

3 Against the Simple Hypothesis

I think SIMPLE cannot be right. I want to rehearse a few arguments to that effect. Ultimately, there are things a believer in SIMPLE can say in response to each of them. I will mention some of them. Thus I do not presume that any of the arguments here are decisive. The hope rather is that all of them, together, can be taken as motivation to seek an alternative story—one that might too be sensitive to the metaphysical contours of our objects of inquiry, but which will not be determined by it.

3.1 *Initial Skepticism*

To a naive reader, the strategy behind SIMPLE might not seem very promising. For it seems eminently reasonable to insist that not all concepts can be said to correspond to a property or relation.²⁴ Singular concepts, quantificational concepts, logical concepts—to name a few—are not naturally construed as corresponding to some property we could go on to evaluate for eliteness. And yet we would be giving up the game if we were to rule out all such concepts as being candidates for conceptual evaluation, or if instead we gave up on giving a unified theory of conceptual evaluation.

Yes, we cannot rule out the possibility of generalizing a theory of eliteness so that it encompasses more than properties and relations.²⁵ But we should acknowledge—*pace* Sider—the awkwardness of thinking that the epistemic benefits afforded by the logical vocabulary our best theories rely on is best explained by appealing to comparisons between the metaphysical counterparts of (say) conjunction and the Sheffer stroke.²⁶

None of this is to say that the burden is squarely on the side of SIMPLE and its ilk. But it might be enough to give us pause before embracing optimism about the prospects of SIMPLE.

Another reason for being skeptical of the strategy recommended by SIMPLE comes from general considerations about the possibility of explaining normative phenomena in purely non-normative terms.²⁷ The view that concepts corresponding to elite properties are better epistemically than those corresponding to less elite ones gives

²⁴ Not that this cannot be done, in principle. For a sense of how such a story could go, see J. P. Burgess (2005); Dasgupta (2009).

²⁵ See Sider (2013) for what is perhaps the most sophisticated attempt at doing that.

²⁶ Cf. Taylor (1993: 99): “Even granted that the *physical* world might come jointed, the notion that what makes adding more natural than quadding is the prior jointedness of mathematical reality, rather than the way we think about it, is peculiarly unappealing.”

²⁷ See Greco (2015) for a helpful discussion of the extent to which ‘open question’-style arguments carry over to debates in meta-epistemology.

rise to the question: why is it better, epistemically, to use concepts that correspond to elite properties? And while there may well be an answer to that question that does not implicitly rely on other normative or evaluative notions, I think it would not be surprising if such an answer is not forthcoming.²⁸

Building a case against SIMPLE on the basis of these two considerations would require a fair amount of work.²⁹ But I do not intend to pursue that here. Rather, I want to focus on a few additional considerations which, together, should at least call into question the strategy recommended by SIMPLE.

3.2 Variety of Explanatory Purposes

Consider the following claim:

EXPLANATION. All else equal, a concept is epistemically good to the extent that it figures in good explanations.

Plausibly, this should be a consequence of any reasonable theory of conceptual evaluation. Is it a consequence of SIMPLE?

If ‘explanation’ in SIMPLE is read as *metaphysical* explanation, then EXPLANATION may seem to follow from SIMPLE, at least given certain reasonable assumptions.³⁰ Assume, first, that metaphysical explanation is a transitive relation among facts. Next assume that if x explains y , then x is *more fundamental* than y . Finally assume that x is more fundamental than y to the extent that properties ‘involved’ in x are more elite than those involved in y .³¹ Then we might be able to conclude that the more explanations appeal to fact x the more elite the properties involved in x are.

This is still not quite what we want. At best we could conclude that the concepts involved in good explanations are likely to correspond to more elite properties, and thus (according to SIMPLE) that they are likely to be good concepts. We could not yet conclude that that all good concepts figure in some good explanations.³² But at least something along these lines might work.

²⁸ Sider (2013) anticipates some of these concerns and claims without argument that aiming to have ‘joint-carving beliefs’ is “a constitutive aim of the practice of forming beliefs, as constitutive as the more commonly recognized aim of truth” (p. 61). But as Hazlett (2017) points out, few if any of the arguments that have been offered for the claim that truth is a ‘constitutive aim’ of belief carry over to support the claim that having ‘joint-carving beliefs’ is a constitutive aim of (the practice of forming) belief(s).

²⁹ See Dasgupta (2018) for a careful attempt at carrying some of this out.

³⁰ What, on a metaphysical conception of explanation, makes for *good* explanations? That is an excellent question, but one I have nothing to say about here.

³¹ This is at best merely a gesture in the right direction. For one, because the claim that a property is involved in a fact needs to be cashed out in more detail—perhaps by assuming that facts are structured entities. But more importantly, because most facts will presumably involve properties of different degrees of eliteness, and we need an argument for thinking that x is more fundamental than y only if the least elite property involved in x is more elite than the most elite one involved in y . After all, a no less plausible hypothesis is that x is more fundamental than y only if the most elite property involved in x is more elite than the most elite property involved in y . (Note too that things get much more complicated very quickly if we allow for facts involving infinitely many properties, or if we allow for the possibility that the ‘more elite than’ relation is not a well-ordering.) Thanks here to Maya Eddon.

³² If we think of facts as comprising a structure partially ordered by ‘is more fundamental than’, we could in principle have facts that are only ‘one-step’ above perfectly fundamental facts but on top of which we will find no other facts they explain. This would give rise to properties that are almost perfectly elite but

We run into more significant difficulties, though, once we recall that good explanations are not explanations *tout court*, but rather explanations *of* something. For, presumably, not everything is equally in need of explanation.³³ To take a simple example, consider the contrast between the following two events:

- (1) Nancy, Tom, and William—three complete strangers—each win one of the three prizes on the raffle at the town fair.
- (2) Alice, Barbara, and Carol—three sisters that came together to the fair—each win one of the three prizes on the raffle at the town fair.

The second of these two cries out for explanation in a way that the first one does not. And, plausibly, the more a fact cries out for explanation the more valuable an explanation of that fact will be. So we should expect that good concepts will figure in good explanations *of what is in need of explanation*. Or, better perhaps: better concepts will figure in good explanations of explananda in need of explanation.³⁴

This by itself may not be enough of a concern. A proponent of SIMPLE might, after all, have a story about what makes a fact cry out for explanation that does not in turn rely on epistemic considerations. Trouble is, as far as I can tell no such story has ever been told. And it is not clear how such a story could be told. For whether something cries out for explanation or not seems to be highly context-sensitive—whether some event cries out for explanation seems to depend in part on something like an epistemic state or a body of theory.

To my knowledge, not much has been written on the question what makes something cry out for explanation.³⁵ But three families of proposals seem to have emerged from the relatively small discussion.

which do not figure in any metaphysical explanations. Concepts corresponding to that property would thus be good concepts, according to SIMPLE, but not good according to EXPLANATION.

³³ This claim is not entirely uncontroversial. See, for example, Friedman (1974: 13): “All phenomena, from the commonest everyday event to the most abstract processes of modern physics, are equally in need of explanation—although it is impossible, of course, that they all be explained at once.” Friedman does not offer an argument for this claim. Instead, he claims he “cannot see any reason but prejudice for regarding some phenomena as somehow more natural, intelligible, or self-explanatory than others” (p. 13).

³⁴ For examples of the way in which the choice of concepts is sensitive to what is deemed worthy of explanation, see the discussion the ‘category influence hypothesis’ in Franklin-(2015: 933)—see especially the discussion of how classificatory practices in chemistry, based on atomic numbers, are partly the result of chemists’ interest in explaining and accounting for certain material transformations, as opposed to the behavior of material in centrifuges. Cf. also Hendry (2010: 147): “eighteenth-century chemistry and its nomenclature were shaped by interests in the qualitative patterns of particular kinds of chemical behaviour (combustion, calcination, acid-base reactions) and explaining them in terms of a particular conception of elemental composition. These patterns are determined by sameness and difference in nuclear charge, and quite insensitive to sameness and difference in atomic weight.” Another cluster of examples can be found in the discussion of the species problem in philosophy of biology. Kitcher (1984), for example, identifies nine different species concepts, each of which is best suited to a particular explanatory purpose—the presumption then is that different explanatory goals result in different assessments as to the epistemic benefits of a particular concept. Cf. also Stanford (1995).

³⁵ For a helpful overview of the rather small literature on this topic, see Grimm (2008). See also Wong and Yudell (2015). Related issues are sometimes discussed in the context of finetuning arguments—there, the question is whether the fact that the universe is capable of sustaining life cries out for explanation—in debates over whether the laws of nature or the universe’s ‘initial conditions’ themselves stand in need of explanation, and in the debates over physicalism in the philosophy of mind, where the question is whether

First, there's what Grimm (2008) calls the 'surprisingness' account.³⁶ A fact cries out for explanation, on this view, to the extent that it is unexpected, or surprising. On the face of it, this seems like a promising view. But it is hard to cash it out so that it isn't vulnerable to obvious counterexamples.³⁷ We are all too familiar with the fact that the universe appears to be perfectly fine-tuned for the emergence of life. And yet that alone would not convince those who think that this fact cries out for explanation that it does not.³⁸

Next, there's a view we can extract from the writings of Sylvain Bromberger, what we can call the *p*-predicament account.³⁹ On this view, a fact cries out for explanation to the extent that (a) we believe there is a true answer to the question *why* it obtained, but (b) any answer we can (currently) think of is ruled out by what we know.⁴⁰

On each of these two views, however, whether something cries out for explanation is highly dependent on features of our epistemic situation. And they make it hard to see how something could cry out for explanation *tout court*, as opposed to crying out for explanation for a particular agent at a time, or relative to a particular body of background beliefs.⁴¹

Not all views on what it takes for something to be in need of explanation carry their sensitivity to epistemic considerations on their sleeves. For example, according to Stephen Grimm, a fact cries out for explanation just in case it is not necessary, where the relevant notion of necessity is narrower than broadly logical or even metaphysical necessity—Grimm does not quite say what the relevant notion is, other than that it is "tied to our sense of what sort of capacities a thing has, relative to the *kind* of thing it is."⁴² The idea, I take it, is that objects (contingently) have

there are any 'brute necessities'. See, for example, Levine (1983); Callender (2004); White (2007). White (2005: 3) offers a plausible necessary condition on what it takes for something to cry out for explanation. Unfortunately, it does not help our discussion, for it appeals to the very notion of being in need of explanation that we are trying to illuminate.

³⁶ Grimm attributes this view to Hempel and Peirce. See p. 485 and references therein.

³⁷ Cf. Grimm (2008: 458, n6). A familiar objection due to Peter Lipton is that many things that are not surprising nonetheless cry out for explanation—the rattle of my car might cry out for explanation, even though there is nothing surprising about it at this point (see, for example, Lipton 2004: 26). The proposal in Wong and Yudell (2015) could be understood as a sophisticated attempt at reviving the surprising account—on their view, a fact cries out for explanation when we are surprised by it not 'fitting' with a background theory (p. 2884).

³⁸ To be sure, there is a way of reading 'surprising' on which it is still surprising that the universe we live in fine-tuned for life. But that reading smells remarkably like a metaphor for being in need of explanation.

³⁹ Bromberger (1962, 1971, 1988, *inter alia*).

⁴⁰ Bromberger characterizes a *p*-predicament as an all-or-nothing affair—see, for example, the Introduction to Bromberger (1992). But one could imagine generalizing his definition so one can be in a *p*-predicament to a greater or lesser extent. Presumably, the relevant sense in which we cannot think of an answer will need to be specified more explicitly (in this regard, see Bromberger 1971: 117).

⁴¹ The point here is independent of the suggestion, most famously made in van Fraassen (1980: chapter 5), that whether something is a good explanation of a given explanandum is also highly context-dependent.

⁴² Grimm (2008: 484). Cf. also p. 494: "a situation stands in need of explanation for someone in virtue of the person's sense that there are various alternative ways the subject of the situation (a system, say, or a substance that constitutes the 'A' in a fact such as *A* is *F*) might have been."

certain capacities, which themselves determine the sorts of states they might be in.⁴³ And it is because x could have failed to be F that, according to Grimm, x 's being F is in need of explanation.

Whatever its merits, though, this view fails to do justice to the guiding contrast between, say, three strangers winning the three prizes at the town raffle and three siblings winning the three prizes. On Grimm's view, each of these is in need of explanation. Moreover, there isn't anything to be said for the rather plausible claim that the second one of these is more in need of explanation than the first one, even if both are in some sense in need of explanation.

There is yet another view worth mentioning, inspired by a passing remark of Cian Dorr's.⁴⁴ A fact cries out for explanation just in case it is unlikely, given that it obtains, that it does not have a good explanation. On this view, the fact that three siblings won the three prizes cries out for explanation because it is unlikely, given that they did win, that there is no good explanation of it.

Now, on a plausible way of spelling this view out, whether something cries out for explanation will depend on some background epistemic state—a credence function, say, that could be used to measure the relevant conditional probabilities. But one could also spell it out so that the relevant conditional probabilities are perfectly objective.⁴⁵ We would still have to worry about which body of evidence is the relevant one to determine how unlikely it is—relative to those conditional probabilities and that body of evidence—that there is no good explanation of some particular fact. But setting that aside, as far as this view is concerned, there could be many concepts that correspond to elite properties which have none of the epistemic good-making features that appeal to eliteness was supposed to explain. Suppose, to illustrate, that there is some concept that correspond to a perfectly elite property. As far as this view is concerned, it could be that this concept never figures in explaining things that cry out for explanation. After all, facts about what cries out for explanation are not determined by facts about what properties are elite. It would thus be hard to see how the claim that a concept corresponds to a perfectly elite property could be what alone explains why it has the epistemic merits that it has.

⁴³ Presumably, Grimm is not thinking of those natures and capacities as being essential to a given thing. Otherwise, the fact that it is metaphysically possible for A to be F would imply that it is possible for A to be F even if we hold fixed A 's nature and capacities, so that the corresponding notion of necessity would be as strong as metaphysical necessity, contrary to what Grimm suggests.

⁴⁴ The view is discussed in a footnote, modified by a 'roughly', and not obviously wide enough in scope to apply to anything beyond whole theories. See Dorr (2010: 154, n29): "Roughly, for a claim to cry out for explanation is for it to be unlikely *conditional on its being true* that its truth is not explained by that of any better theory." There are multiple ways of making sense of this claim. On one, what cries out for explanation is a complete theory: T cries out for explanation just in case it is unlikely, given T , that T is not explained by some other theory T' that is better (in the relevant sense) than T . The downside of this understanding is that the scope is limited to whole theories. Another interpretation is this: a claim C cries out for explanation just in case it is unlikely, given C , that C is not explained by some theory T' that is better (in the relevant sense) than C . The downside of this way of thinking about it is that it requires that we find a suitable way of comparing particular claims about some event (say) with complete theories. Perhaps the view is that T' is better than C in the relevant sense if it is a better explanation of C than C itself. (Though this would make it hard to see how there could be any claim that does not cry out for explanation, save perhaps for 'self-explanatory' claims, whatever those turn out to be.) Or perhaps the idea is that T' is better (in the relevant sense) if T' is better than some other salient, available putative explanation of C . I cannot tell.

⁴⁵ As in, for example, Williamson (2000: chapter 10).

Incidentally, it is worth highlighting a general upshot of this last observation. The objection to SIMPLE is *not* that it would imply, in contrast to all plausible theories of what it takes to be in need of explanation, that it is an objective matter whether something cries out for explanation. The objection, rather, is that SIMPLE would imply that whether something is in need of explanation has nothing to do with *epistemic* considerations. It could be a perfectly objective matter whether something is in need of explanation relative to a given epistemic situation. So unless we have some sense as to how being in need of explanation could have nothing to do with epistemic considerations, we have some reason to think that SIMPLE cannot be right.

3.3 *Trouble with Uniformity*

A theory of conceptual evaluation should, if possible, be sufficiently general in scope. We want an answer to the question, what makes for good concepts, epistemically, that is relatively discipline neutral. To illustrate: an answer that applies uniformly to both mathematical concepts and concepts from the natural sciences is preferable, all else equal, than one that does not.

How does SIMPLE fare with respect to this desideratum?

On the one hand, SIMPLE generalizes quite nicely to different domains of inquiry. As long as it makes sense to talk of elite properties in a given domain, we can use SIMPLE to answer the question what makes for good concepts in that domain. And while not all proponents of the elite/non-elite distinction might be willing to countenance that such a distinction carries over to mathematical properties, for example,⁴⁶ I will simply assume that the proponent of SIMPLE is not among them.

On the other hand, digging a little deeper reveals that things are not as straightforward as they appear.⁴⁷ According to SIMPLE, what explains why the concept *electron* is a good concept is that the property of being an electron is suitably elite. And what explains why the concept *determined game* is a good concept is that the property of being a determined game is also elite.⁴⁸ Ask now: how are these two explanations supposed to go?

⁴⁶ If elite properties constitute a minimal supervenience base for all of reality, then there is no hope for a non-trivial theory of elite mathematical properties. And if elite properties are qualitative properties, there is even less hope for a non-trivial theory of elite mathematical properties. (Cf. Lewis 1991: section 2.6.) None of this need imply that one cannot give a theory of naturalness for mathematical properties: it's just that whatever that theory turns out to be it will have a rather different look than any of the familiar theories of eliteness. For relevant discussion, see Bricker (n.d.).

⁴⁷ I'm indebted to Tappenden (2008a) for raising the question whether there is a plausible theory of the natural/non-natural distinction that encompasses both mathematical and non-mathematical properties. Cf. also the discussion in Tappenden (2008b) of an argument in Sider (1996) for the claim that preference for one over another way of defining ordered pairs in set theory can only be based on merely pragmatic considerations.

⁴⁸ A (two player, perfect information) game which does not allow for draws is *determined* iff one of the players has a winning strategy: a strategy that will allow her to win no matter what the other player does. Alternatively, if we allow for games where players can draw, we say that a game is determined if one of the players has a strategy that guarantees she will not lose. A fundamental theorem attributed to Zermelo (but see Schwalbe and Walker 2001 for a more accurate account of the history behind the role of Zermelo's work in early game theory) is that every finite game—a game that ends after a finite number of moves—is determined. Not all infinite games are determined. Characterizing the class of infinite games that are determined has proven to be a surprisingly fruitful research program. More on this below.

Consider one particular theoretical benefit we might expect to accrue from epistemically good concepts: fruitfulness. A concept is fruitful, as I will understand the term, when it makes unexpected appearances in dealing with a variety of questions.⁴⁹ One particularly important way a concept can be fruitful is in lending itself to successful prediction: one can typically predict features of an unexamined object falling under that concept on the basis of features of other, examined objects falling under that same concept. So the question is: why would the fact that a concept corresponds to an elite property explain that concept's contributing to successful prediction?

If we zoom in on a particular class of arguably elite properties—the so-called natural kinds—we can get some sense of how to answer that question. Kinds, one might think,⁵⁰ correspond to property clusters that are sustained by certain causal mechanisms. It is the fact that certain causal mechanisms sustain the co-instantiation of certain features that explains why concepts corresponding to natural kinds are projectable.⁵¹

The idea, very roughly, is that properties come in clusters: for example, thermal and electrical conductivity, ductility, malleability, and having a shiny appearance. And this clustering is due to the presence of underlying mechanisms (something having to do with the presence of free electrons, in our example) which ensure that those properties tend to be coinstantiated.⁵² As a result, the property *being metal* is a property that supports inductive generalizations: from the fact that the temperature of a given metal is inversely proportional to its conductivity, one is warranted in inferring that the temperature of other metals is too. This at least gives us a sketch of a story as to how the metaphysics of the property corresponding to our concept *metal* can account for the role that this concept plays in our epistemic practices.

Now, much like the concept *being metal*, the concept *determined game* is a fruitful concept. To see what I mean, let me take a moment to introduce some definitions. Any given set of real numbers can be identified with a unique set of sequences of natural numbers. For any set *A* of sequences of natural numbers, we can define the

⁴⁹ Note here that at least some of the problems raised by Nolan (1999) against taking the fruitfulness of a *theory* to be an epistemic virtue do not arise for thinking of fruitfulness as an epistemically good-making feature of concepts. As Nolan points out, if the fruitfulness of a theory is, as the term might suggest, a matter of it giving rise to new problems or opening up new lines of inquiry, it is hard to see how fruitfulness (or 'fertility', to use Nolan's term) would be an epistemic virtue. To think so would be like thinking that "Faces many problems' or 'Could do better' or 'Much room for improvement' are high praise on the report card of a theory" (p. 267). Similar concerns do not carry over to thinking of fruitfulness of *concepts* as an epistemically good thing.

⁵⁰ See, for example, Boyd (1988); Kornblith 1993); Millikan (1999).

⁵¹ It is telling that such homeostatic property clusters can in principle be sustained by the intentional activities of human beings. See, for example, Mallon (2003). Presumably, such properties would not count as elite on any reasonable theory of eliteness. Otherwise, whether a property is elite would depend on contingent facts about human interests at a given point in time. (In this respect, the discussion of the concept of the GDP in Coyle (2014), and the way in which changes in the economy make it a less than useful concept can serve as an illustration of this point.) That suggests that the explanation for why concepts corresponding to these properties turn out to be epistemically valuable will have little to do with the eliteness of the relevant properties. But this is a line of argument I do not intend to pursue here. (Thanks to Hilary Kornblith for discussion on this point.)

⁵² Cf. Boyd (1999: 82f).

two-player game G_A as follows: players I and II alternate picking a natural number; player I wins iff the resulting sequence is an element of A —else, player II wins. We say that a set of real numbers A is determined iff G_A is determined.

It is well-known that all open sets (countable unions of disjoint open intervals) and all Borel sets (sets obtained from open sets by closing under countable unions, countable intersections, and relative complements) are determined. The *Axioms of Definable Determinacy* state that all ‘definable’ sets of reals are determined—I write ‘axioms’ because there are different ways of making the notion of definability precise, each of which corresponds to a particular axiom of definable determinacy.⁵³ Axioms of Definable Determinacy have been a driving force behind much work in set theory over the last thirty years.⁵⁴ The study of determinacy did not just result in a number of general results about how well-behaved ‘reasonably definable’ sets of real numbers are—for instance, it is known that all determined sets of reals are Lebesgue measurable, have the Baire property, and contain a perfect subset if uncountable—but also served to establish deep connections between two seemingly unrelated branches of set theory: the study of large cardinals and descriptive set theory.⁵⁵

Indeed, it would not be implausible to talk of the notion of determinacy as being projectible: it seems reasonable to conclude, from the assumption that all sets of reals we have been able to define are determined, that all definable sets of reals are determined, even though strong versions of definable determinacy are independent of the standard axioms of set theory.⁵⁶ Moreover, axioms of definable determinacy have been crucial in proving results in a range of different areas which later turned out to be provable in much weaker theories.⁵⁷ Thus, the notion of determinacy seems to have as good a claim at being an epistemically good concept as any natural kind term.

Of course, we cannot explain the fruitfulness of the notion of determinacy in the same way we explained the fruitfulness of the species concept. It would be quite a stretch to posit something analogous to a causal mechanism which sustains (in a non-causal way) the clustering of so-called regularity properties (Lebesgue measurability, etc.) around determined sets of reals. So a very different type of explanation will be needed to account for the fact that concepts corresponding to elite mathematical properties turn out to be fruitful. Such a story might well be forthcoming—perhaps mathematical properties have their own, distinctive way of being elite, one which explains the fruitfulness of the corresponding concepts. But as far as I can tell, no such story has ever been told. (The view that for a mathematical property to be

⁵³ Perhaps the best-known of these axioms states that all sets of reals in $L(\mathbb{R})$ —the smallest transitive model of set theory that contains all reals and all the ordinals—are determined.

⁵⁴ For an overview of some of the main results, see Koellner (2014); Welch (2015); as well as Maddy (2011: chapter II). For some of the mathematical details, see, for example, Kanamori (2009: chapter 6).

⁵⁵ The canonical reference here is Shelah and Woodin (1990). See Larson (2012); Koellner (2014) for a more comprehensive list of references.

⁵⁶ This is not an isolated example. See, for example, the discussion in Tappenden (2008b) of the quadratic reciprocity theorem and the role that ‘projecting’ on the predicate ‘is a quadratic residue’—and relatedly, the introduction of the Legendre symbol—played in the discovery of the theorem. Unlike in the case of axioms of determinacy, which have been accepted on ‘mere’ non-deductive grounds, the quadratic reciprocity theorem admits of a direct proof. See also the discussion of the concept of field, and the way it contrasts with Frege’s concept of a ‘quantitative domain’, in Tappenden (2005: 20f).

⁵⁷ For further discussion, see Martin (1998).

elite is in part for the corresponding concepts to be fruitful would certainly not do.) And without such a story, a proponent of SIMPLE will be unable to give a unified answer to the question, what makes for good concepts, epistemically?

(Note that what I say here is perfectly compatible with the claim that fruitfulness of mathematical concepts is evidence that they correspond to elite mathematical properties.⁵⁸ What I'm calling into question is the existence of a suitable notion of eliteness for mathematical properties that might explain why concepts corresponding to such properties are fruitful.⁵⁹)

3.4 *Concept and Property*

I have been presupposing, for the sake of compliance with the presuppositions of SIMPLE, that to each concept corresponds a unique property. It is only given this assumption that SIMPLE can be formulated: if a concept corresponds to multiple properties, then we are not guaranteed that concepts can be compared in terms of the eliteness of *the* corresponding properties.

Now, for reasons all too familiar, this presupposition is not quite right—at least given plausible assumptions about how to individuate concepts. For instance, consider the concept *local*. Intuitively, there is no one property—the property of *being local* as it were—that corresponds to that concept. Rather, in different contexts, different properties are predicated of an object that is judged to fall under that concept.⁶⁰

A proponent of SIMPLE could now restrict her view so that it only applies to concepts that do meet this presupposition. So let us banish 'local' from among the candidates to good concepts—banish, too, concepts like *leftmost*, *tall*, and *currently fashionable*. In principle, there is nothing objectionable to this strategy. The worry is that banishing context-sensitive concepts—concepts that correspond to different properties in different contexts—will leave us with little to apply SIMPLE to. In particular, the concern is that by banishing context-sensitive concepts we may end up banishing the very concepts we want a theory of conceptual evaluation to apply to.

What motivates this idea is the observation that many of our seemingly non-context-sensitive terms appear upon reflection to be context-sensitive after all.⁶¹ Take a familiar example, from (Chomsky 1976):⁶²

⁵⁸ Cf. Maddy (2011), whose 'thin realism' is based on the idea that the success of the set-theoretic methods is evidence that those methods are tracking the 'underlying contours of mathematical depth' (p. 82).

⁵⁹ For a nice summary of different attempts at cashing out a notion of mathematical depth—as well as their limitations—see Arana (2015). Arana's discussion is focused on a notion of depth as applied to mathematical theorems. But many of the issues he raises could be transposed into a discussion of depth as applied to mathematical properties.

⁶⁰ Even proponents of so-called semantic minimalism would grant that there is no such thing as the property of being local *simpliciter*. See Cappelen and Lepore (2005: 1). (Admittedly, Cappelen and Lepore express some skepticism that we should treat so-called contextuals as genuinely context-sensitive expressions—see fn. 1, p. 1.)

⁶¹ The relevance of this brief foray into the debate over the extent of context-sensitivity in natural languages to our discussion about concepts should become clearer shortly.

⁶² See also Chomsky (1995, 2000), and many other places. For illuminating discussion of these and related observations, see Pietroski (2003); Collins (2009).

- (3) a. John wrote a book.
 b. The book weighs five pounds.

As Chomsky points out (p. 48), the term *book* behaves differently in the two sentences in (3). In (3a), it must pick out an abstract entity: the book that John wrote can survive the destruction of the original file it was stored in. In (3b), however, it must pick out a concrete entity, capable of having mass. Interestingly, what is going on in (3) cannot obviously be explained by claiming that the two occurrences of *book* correspond to different lexical items. After all,

- (4) John wrote a book that weighs five pounds.

is perfectly acceptable, in contrast with (e.g.)

- (5) Jumbo waved his trunk, which was full of clothes.

which is to

- (6) a. Jumbo waved his trunk.
 b. The trunk was full of clothes.

like (4) is to (3).⁶³

And *book* is not unique in this regard. Consider (from Chomsky 2000: 37):

- (7) a. London is unhappy.
 b. London is ugly.
 c. London is polluted.
 d. London is so unhappy, ugly, and polluted that it should be destroyed and rebuilt 100 miles away

Again, it seems as if *London* denotes slightly different things in (7a)–(7c), and yet the acceptability of (7d) suggests that we are not dealing with multiple lexical items here. Also consider (cf. Chomsky 1970: n. 7 and 1976: 49f):

- (8) a. The temperature is rising.
 b. The temperature is 70°.
 c. The temperature, which was 70°, is rising.
- (9) a. Mary wrote the proof in pen and paper.
 b. The proof will have a lasting impact on the history of mathematics.
 c. Mary wrote a proof, in pen and paper, that will have a lasting impact on the history of mathematics.

To be sure the arguments here are not decisive.⁶⁴ But they strongly suggests that terms like *book*, *proof*, and *temperature*, display a certain amount of unexpected

⁶³ Example (5) is example (19ii) in Chomsky (1976).

⁶⁴ See for example the discussion of Partee's so-called temperature paradox in Montague (1973), as well as the more recent discussion in Lasersohn (2018).

context-sensitivity, so that the property corresponding to each such predicate varies from context to context.⁶⁵

All that said, these examples might appear somewhat irrelevant to the more interesting questions surrounding conceptual evaluation from an epistemic perspective. After all, they all concern what we could call ‘folk’, or ‘common-sense’ concepts. And it is tempting to think that part of what scientific progress consists in is the abandonment of such concepts. Once we narrow our attention to the conceptual tools actually used in mature sciences, we will no longer be faced with the kind of context-sensitivity that might otherwise get in the way of something like *SIMPLE*.⁶⁶

Or so one might think. But unfortunately for the proponent of *SIMPLE*, it looks like many of the concepts that play a significant role in (non-fundamental) sciences can’t quite be assigned to a single property. The phenomenon—which, following Wilson (2006) we can call the ‘multi-valuedness’ of the relevant concepts—is best understood by looking at some examples.

Consider the concept *hardness*.⁶⁷ What is the property corresponding to it? Of course, one *could* say that the property corresponding to the concept *hardness* is, well, the property of hardness. But in trying to find anything illuminating to say about what that property is—something that we might undoubtedly need to do in order to figure out its position on the eliteness scale—we quickly find ourselves at a loss.

There are different more or less familiar tests for hardness—scratching tests, indentation tests, rapping tests, among others.⁶⁸ Which test is deemed appropriate depends on a variety of factors, including the type of material being tested for hardness. As Wilson writes (2006: 336):

[I]n everyday contexts we adjudicate the “hardnesses” of various materials, both comparatively and absolutely, through a wide variety of comparatively easy to apply tests—we might squeeze the material or indent it with a hammer; attempt to scratch it or rap upon it; and so on. In most cases, we will be scarcely aware of the exact technique we will have employed for this appraisal: “Did I rap, squeeze or scratch that piece of wood? I can’t really remember.” In fact, our choice of tests is likely to have been suggested by the material in question: we instinctively appraise a

⁶⁵ It is of course a non-trivial question to explain why these terms get to behave the way they seem to in sentences like (9c), and I have nothing to offer by way of explanation here.

⁶⁶ Indeed, Chomsky himself—insisting as he does that developing ‘referential’ semantics is the wrong way to go about studying meaning—seems to leave it open whether terms introduced in the natural sciences are well-behaved enough so that each term or concept corresponds to a unique property (see Chomsky 2000: 42f). On Chomsky’s view, part of what prevents terms like *London* or *house* from being part of a ‘mature’ science (which Chomsky takes to mean that clauses like ‘*London* refers to London’ cannot be part of any serious scientific study of meaning) is that they are “used to refer to concrete objects, but from the standpoint of special human interests and goals and with curious properties”. And the evidence for this is that, for example, “[a] house can be destroyed and rebuilt, like a city; London could be completely destroyed and rebuilt up the Thames in 1,000 years and still be London, under some circumstances” (p. 21). No term belonging to mature ‘naturalistic inquiry’ should thus display the characteristics of terms like *London* or *house*. If the relevant characteristics include the seeming context-sensitivity of such terms, Chomsky must then think that highly context-sensitive items should be banished from good scientific vocabularies. Cf. Stoljar (2015).

⁶⁷ As should become clear, my discussion here is deeply indebted to Wilson (1982 as well as 2006).

⁶⁸ Wilson (2006: chapter 6, ix).

wood by rapping upon it, a rubber by squeezing, a metal by attempting to make a small imprint; a glass or ceramic by rapping lightly or scratching (not by trying to make a small imprint!).

But the type of material alone does not determine which particular test is appropriate. A hunk of steel, for example, may sometimes be tested using an indentation test—measuring the diameter of an indentation left behind by a steel ball that is placed on its surface for a fixed amount of time, as in the familiar Brinell tests—but also using an abrasion test. And the two tests can yield conflicting verdicts as to the comparative hardness of different types of steel.⁶⁹ It is partly because of this that one finds in the literature on hardness tests claims to the effect that *hardness* is not a ‘fundamental’ property, or that there may be no such thing as the property of hardness.⁷⁰

At this point, it might be tempting to conclude that the concept *hardness* is somehow defective. And indeed, there may be some sense in which it is. But it would be unreasonable to conclude from this that there is nothing to be said for it, epistemically, given the way it is embedded in a network involving other macroscopic concepts like *solid*, *rigid body*, and *force* which seem essential to our understanding of the macro world.⁷¹

Can the proponent of SIMPLE explain what makes *hardness* an epistemically good concept? One possibility would be to insist that there really are multiple concepts in place: *Brinell hardness*, *scratch hardness*, and so on. The mistake, according to this line of thinking, is to think that there is a single concept whose epistemic credentials we need to account for. Rather, there are multiple concepts, each one of which is epistemically good in its own right. It is merely an accident of the way English developed that we have homonyms corresponding to different concepts.

There are at least two problems with this strategy, however. First, and perhaps least significantly in this context, this strategy would fail to do justice to the way that the

⁶⁹ Cf. Wilson (2006: 338). Of course, in discussing these issues authors sometimes appeal to qualifiers to avoid confusion. Thus, Walley (2012: 1033) disambiguates and talks about “the relation between Brinell hardness and resistance to abrasion” not being as straightforward as one might have expected. But sometimes, authors expect readers to disambiguate according to context, as in, for example, Malzbender et al. (2002: 52): “hardness is not a fundamental property of materials. Hardness is related to material properties, in particular to the yield strength and the elastic modulus, but this relationship depends on the indenter geometry. We should, therefore, be careful when comparing hardness values from different sources.” Additional examples of this way of exploiting the ambiguity of ‘hardness’ are easy to find in the literature.

⁷⁰ For example, Newey and Weaver (2013: 13): “A hard material is difficult to scratch, wear away by abrasion or to indent. Hardness is not a fundamental property of a material: for each method of measuring it, it is some combination of elastic, plastic, and (in some cases) fracture properties. Hardness can be measured only by comparison with a material used as a scratcher or indenter and has objective meaning only in terms of a specific type of test. For example, glass will scratch steel but fractures more readily under indentation; nylon has a high resistance to wear but not to indentation.” And Cahn and Lifshin (1993: 183): “Although there is probably no such property as ‘hardness’, it is a convenient generic test.” See also the quotes from Samuel Williams in Wilson (2006: 350).

⁷¹ Cf. the discussion of solidity in Wilson (2006: 352–3). As Wilson points out, the notion of solidity that is usefully applied to macroscopic objects—solid objects are those that “preserve their volume under all interactions, as long as they remain integral” (p. 354)—is not the one we would rely on when discussing the microscopic interactions between objects, for it is precisely by “not acting like a solid at the molecular level” that a piece of steel does not crumble into dust under mild applied stress.

term ‘hardness’ is actually used. The apparent ambiguity or polysemy of ‘hardness’ does not seem to pattern with lexically ambiguous terms like ‘bank’. It would be at best an unimaginative joke to say: “there are two different kinds of banks: river banks and financial institutions”. But talk of different kinds of hardness is not out place in the literature on materials science,⁷² as is the claim that hardness is the capacity to resist penetration or abrasion, when there is no property that accounts both for resistance to penetration and resistance to abrasion.⁷³

The second problem is that there seems to be no explanation, on this way of thinking, for why these allegedly different concepts are ‘linked’ together in the way they are. It certainly seems as if part of what lets concepts like *hardness* play the role they do in understanding macro phenomena is that they adjust themselves almost unnoticeably so as to pick out different properties in different contexts, or, if we prefer, so as to take each other’s place (the different concepts) in different contexts. And this seems to be something over and above the putative eliteness of the properties corresponding to each of the ways in which *hardness* is used.

To see what I have in mind, consider a different example, one where the ‘switch’ from property to property is induced not by focusing on a different material but by ‘shifting scales’.⁷⁴ When considering a three-dimensional object at the macro-level, we think of its surface as consisting of a two-dimensional ‘skin’ that surrounds it. Now, there are well-known puzzles that arise from applying these topological notions to macroscopic objects.⁷⁵ But there’s no denying the immense fruitfulness of doing so. This is not just in the sense that our folk-theoretic understanding of macroscopic objects in some sense presupposes that surfaces are aptly thought of in topological terms. These topological notions are essential to continuum mechanics, for example, even if doing so requires a certain amount of theoretical acrobatics.⁷⁶

Consider now what it takes for something to be a surface when looked at from a ‘micro’ point of view. There, it makes sense to think of surfaces as lattice structures

⁷² Cf. the quoted passage in fn. 69.

⁷³ Cf. also Wilson (1982: 575f): “We can meaningfully ascribe an extension (with certain implicit parameters) to the everyday use of ‘is 90°F’ which differs from that appropriate to the temperature experts in our society. There is a natural temptation to claim that the experts assign a somewhat different *meaning* (or sense) to ‘is 90°F’ than the rest of us. The fact that ‘is 90°F’ corresponds to two fairly clearcut *ranges of application* represents an interesting aspect of English linguistics, but this common phenomenon should be regarded as *sui generis* and not lumped with ambiguities of ‘meaning’ such as the word ‘bank’ displays. Intuitive talk of the ‘varying senses’ of ‘is 90°F’ is unexceptionable if understood merely as a description of the double range of application, but it should not mislead us into supposing that the experts ‘grasp’ a distinct concept (especially when physicists employ the term in the same way as the rest of us in ordinary conversation with no sense of disharmony with their laboratory practice).”

⁷⁴ What follows owes much to Wilson (2006); Batterman (2013); Bursten (2018).

⁷⁵ For a nice summary of some of the issues, see Arntzenius (2008); Varzi (2015).

⁷⁶ See, for example, the discussion of how spray can form on the the surface of water in Wilson (2006: 210). As Wilson points out, if we start with a smooth water surface and play out the differential equations governing the flow of wind and water, we will never get the change in topology required in order for a bit of the surface to break off into a water droplet. As a result, some acrobatics are necessary in order to model the generation of the spray. The way Wilson describes it, “[w]hen a change in the fluid’s topology looks imminent, practitioners begin investigating two fluid configurations that run in parallel, one containing the still attached drop and the other describing a drop of similar shape detached from its ocean.”

made up of atoms ‘connected’ to one another by electromagnetic forces. It is by looking at the details of these structures that we can explain, say, why a given metal spoon will bend more easily than a seemingly identical one made of the same material. It is because of slight imperfections on the surface of the spoon—so-called *dislocations*, which are essentially irregularities in the crystalline structure of the material—that a steel spoon can be bent, the quantity and distribution of dislocations determining how easy it is for it to bend. And it is because of the presence of ‘gaps’ between the atoms that make up the surface of a metal beam that oxidation occurs⁷⁷—which in turn explains why so-called cold-welding does not happen in everyday situations.⁷⁸

Our concept *surface*, then, seems to pick out different properties depending on something like a background choice of scale.⁷⁹ A proponent of SIMPLE might thus be inclined to think that there are different concepts here, one for each choice of scale, and offer independent explanations for why each of them is epistemically good.

Now, on the face of it, there is something appealing about the view on which we have different *surface* concepts—a ‘continuum’ surface concept and an ‘atomistic’ surface concept. Yet there are many ways in which these concepts are related: for instance, changes in the microstructure of the surface of the object will in turn affect how the surface should be modeled in continuum mechanics.

Suppose, for instance, we are interested in explaining the way in which hammering a piece of metal might affect its macroscopic properties (through so-called work-hardening, one can strengthen a hunk of steel by making it less ductile). We may need to start by modeling its surface at the microlevel (perhaps because a blow resulted on some cleavage cracks along the surface, introducing dislocations) and model the rest of the steel using a continuum model. But as the hammering of the steel goes on, and the surface dislocations start moving down to the bulk of the steel, our model of the surface may need to switch: extending the atomistic model along the path that the dislocation takes will quickly become prohibitive.

⁷⁷ See Cadavid and Cabot (2017) for an accessible discussion of this point.

⁷⁸ Cf. Wilson (2008: 285f): “The outer molecular layers of an iron bar (which are anti-averaged into the usual two-dimensional boundary conditions of elasticity theory) actually comprise an extremely complex region when viewed up close.”

I was alerted to the possibility of cold-welding in the absence of ‘greasy atmospheric crud’ by Wilson (2006: 373). A particularly dramatic illustration of what can happen to metal once oxides are removed is the well-known anomaly in the deployment of the antenna of the Galileo spacecraft: two metal parts were essentially ‘welded’ to one another, preventing the antenna from fully deploying. See, for example, Miyoshi (1999) for some of the details. Further discussion of the more general issues can be found in the works cited in fn. 92 of Wilson (2006, *loc. cit.*).

⁷⁹ Another example of this type of scale-dependence, also from Mark Wilson, is the concept *force*. As Wilson puts it (Wilson 2013: 54): “The term *force* has a notorious tendency to alter its exact significance as characteristic scale lengths are adjusted. At a macroscopic level, the ‘rolling friction’ that slows a ball upon a rigid track is a simple Newton-style force opposing the onward motion. But at a lower scale length, the seemingly ‘rigid’ tracks are not so firm after all: they elongate under the weight of the sphere to a nontrivial degree. So part of the work required to move our ball against friction consists in the fact that it must *travel further* than is apparent. But when we consider the ‘forces’ on our ball at a macrolevel, we instinctively treat the track length as fixed and allocate the effects of its actual elongation to a portion of the ‘force of rolling friction’ budget.”

Or suppose we are trying to understand the way in which enough stress on a hunk of metal can lead to fracture.⁸⁰ If we rely simply on a continuum mechanical description of the phenomena, we will be unable to predict the emergence of cracks along the surface of our metal—cracks which will grow and result in fracture. This is because introducing cracks requires a change in the topology of our surface, which cannot be predicted by the continuous deformation allowed in a continuum mechanical setting. So in order to account for the emergence of cracks, we need to switch to modeling the surface at a microscale. Unfortunately, we cannot just model the phenomenon of interest entirely at the atomic level, for the number of atoms along our surface is computationally intractable. For these purposes, physicists have developed *multiscale* models, where different models of a surface (say) are ‘linked’ together, so that part of the surface where cracks emerge and grow are modeled at the atomic level and the rest of the surface—along and beyond which the applied stress will flow in a somewhat predictable manner—using a continuum model.⁸¹

And these relations are not clearly explicable in terms of law-like relations among the corresponding properties (for one, the property corresponding to the ‘continuum’ surface concept is one that no concrete object has in this world, so there could not be any law-like connection between instantiations of the property corresponding to the ‘atomistic’ *surface* concept and instantiations of the property corresponding to the ‘continuum’ one). Moreover, thinking of the difference between the two *surface* concepts as a simple case of lexical ambiguity would fail to do justice to the fact that both continuum and atomistic models of the effects of cleavage on a surface are models of the same phenomenon. We would be mischaracterizing multiscale modeling of materials if we thought of them as concerned with different phenomena at the same time, rather than as concerned with the same phenomenon at different scales.

Granted, the proponent of SIMPLE could instead insist that there is a single operative concept here, viz. the ‘atomistic’ *surface* concept. Talk of surfaces as continuous regions is, on this view, strictly speaking false. Still, the fact remains that the ‘continuum’ *surface* concept (assuming, for the sake of argument, that it is a different concept) plays a role in explanations which, according to EXPLANATION, should be deemed an epistemically good concept. Exactly what explains that it is an epistemically good concept is a difficult question. But it is hard to see how the answer will be compatible with SIMPLE.

⁸⁰ Perhaps a more pressing explanandum is the fact that the force required to slide a block over a flat surface is proportional to the block’s weight but independent of the apparent area of contact—the so-called Amonton Law of Friction. The orthodox explanation appeals to the fact that, because of surface asperities, the true area of contact is much smaller than the apparent area of contact and it increases by way of deformation of these asperities in proportion to the applied compressive force. But alternative multiscale explanations—that do not appeal to surface roughness—are available, appealing instead, for instance, to the presence of self-healing cracks on the surface of the materials in proportion to the ratio of horizontal to compressive stress. See Gerde and Marder (2001) for discussion. See also Krim (1996) (and references therein) for a survey of the limitations of explanations appealing to surface roughness.

⁸¹ See, for example, Zhang, Johnston, and Chattopadhyay (2014: 120): “The phenomenon of fatigue involves multiple length scales including crack nucleation in the microscale, coalescence of microcracks in the mesoscale and major crack propagation in the macroscale. Therefore, it is critical to develop a lengthscale-dependent and physics-based model to understand material performance and ultimately assess the survivability of complex structural systems.” See also Rudd and Broughton (2000: esp. section 1).

4 Where to Next?

I wish I had an alternative story to offer, here, as to what makes for epistemically good concepts. Unfortunately, I do not. But I think there are a number of lessons to be drawn from our discussion so far, lessons that any such alternative story must take into account. In closing, I want to briefly go over what I take those lessons to be.

First, it may turn out to be impossible to give an informative answer to the question what makes for epistemically good concepts without making more substantive assumptions about what concepts are. The role that concepts—or meanings, if you would rather avoid talk of concepts altogether—play in our ability to make sense of the world around us cannot be fully accounted for in terms of a list of properties that the relevant concepts correspond to. We thus need a better grip on what concepts are—other than whatever mediates our representation of properties—in order to understand what makes some concepts better suited to our epistemic projects than others.

Second, we should be open to the possibility that no purely ‘metaphysical’ answer to our question will be forthcoming. Nothing in what I’ve said here rules out the possibility that some story that is not sensitive to features of our epistemic situation—one that goes beyond the relative eliteness of the properties corresponding to concepts—might be the best answer to the question what makes for epistemically good concepts. (Perhaps there are some relations among properties such that concepts are epistemically good to the extent that they are associated with a family of suitably related, sufficiently elite properties.) But unless it is roughly as simple as SIMPLE, we have no reason to favor an answer that is independent of our epistemic circumstances to one that is not—the one thing SIMPLE had going for it, which might have led to our thinking that something like it had better be right, was its simplicity. Any story about what makes for epistemically good concepts will, I suspect, turn out to be rather complex. Whether that complexity will come from the metaphysics of properties or from facts about the way in which creatures with our cognitive capacities interact with our rather complex environment is an open question.

Acknowledgements

For helpful comments, discussion, and advice, thanks to Phil Bricker, Maya Eddon, Nina Emery, Sophie Horowitz, Joe Levine, Chris Meacham, Katia Vavova, and participants at the 2016 workshop on Conceptual Engineering at the CSMN, University of Oslo. Special thanks to Alexi Burgess and Hilary Kornblith for extremely helpful comments on an earlier draft of this chapter.

References

- Antony, Louise. 2003. Who’s Afraid of Disjunctive Properties? *Philosophical Issues* 13:1–21.
- Arana, Andrew. 2015. On the Depth of Szemerédi’s Theorem. *Philosophia Mathematica* 23 (2):163–76.
- Armstrong, David M. 1978. *Universals and Scientific Realism*, vol. 1. Cambridge: Cambridge University Press.

- Arntzenius, Frank. 2008. Gunk, Topology and Measure. In Dean Zimmerman (ed.), *Oxford Studies in Metaphysics*, vol. 4 (pp. 225–47). Oxford: Oxford University Press.
- Batterman, Robert W. 2013. The Tyranny of Scales. In Robert W. Batterman (ed.), *Oxford Handbook of the Philosophy of Physics* (pp. 255–86). Oxford: Oxford University Press.
- Berker, Selim. 2018. The Unity of Grounding. *Mind* 127 (507):729–77.
- Blackburn, Simon. 1998. *Ruling Passions*. Oxford: Oxford University Press.
- Boyd, Richard. 1988. How to Be a Moral Realist. In Geoffrey Sayre-McCord (ed.), *Essays on Moral Realism* (pp. 181–228). Ithaca: Cornell University Press.
- Boyd, Richard. 1999. Kinds, Complexity and Multiple Realization. *Philosophical Studies* 95 (1):67–98.
- Bricker, Philip. 1992. Realism without Parochialism. Unpublished manuscript, University of Massachusetts, Amherst. Symposium paper at the Pacific APA meeting, 1992.
- Brigandt, Ingo. 2003. Homology in Comparative, Molecular, and Evolutionary Developmental Biology: The Radiation of a Concept. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution* 299B (1):9–17.
- Brigandt, Ingo. 2010. The Epistemic Goal of a Concept: Accounting for the Rationality of Semantic Change and Variation. *Synthese* 117 (1):19–40.
- Bromberger, Sylvain. 1962. An Approach to Explanation. In R. J. Butler (ed.), *Analytical Philosophy*, vol. 2 (pp. 72–105). Oxford: Oxford University Press. Reprinted in Bromberger (1992: 18–51). Page numbers refer to the reprinted version.
- Bromberger, Sylvain. 1966. Why-Questions. In Robert G. Colodny (ed.), *Mind and Cosmos: Essays in Contemporary Science and Philosophy*, vol. 3 (pp. 86–111). University of Pittsburgh Series in the Philosophy of Science. Pittsburgh, PA: University of Pittsburgh Press. Reprinted in Bromberger 1992: pp. 101–11. Page numbers refer to the reprinted version.
- Bromberger, Sylvain. 1971. Science and the Forms of Ignorance. In Maurice Mandelbaum (ed.), *Observation and Theory in Science*. Baltimore: The Johns Hopkins Press. Reprinted in Bromberger (1992: 112–27). Page numbers refer to the reprinted version.
- Bromberger, Sylvain. 1988. Rational Ignorance. *Synthese* 74 (1):47–64. Reprinted in Bromberger 1992: pp. 128–144. Page numbers refer to the reprinted version.
- Bromberger, Sylvain. 1992. *On What We Know We Don't Know*. Chicago and Stanford: The University of Chicago Press and CSLI.
- Burgess, Alexis, and Plunkett, David. 2013a. Conceptual Ethics I. *Philosophy Compass* 8 (12):1091–101.
- Burgess, Alexis, and Plunkett, David. 2013b. Conceptual Ethics II. *Philosophy Compass* 8 (12):1102–10.
- Burgess, John P. 2005. Being Explained Away. *Harvard Review of Philosophy* 13 (2):41–56.
- Bursten, Julia R. 2018. Smaller Than a Breadbox: Scale and Natural Kinds. *British Journal for the Philosophy of Science* 69 (1):1–23.
- Cadavid, Doris, and Cabot, Andreu. 2017. Oxidation at the Atomic Scale. *Science* 356 (6335):245.
- Cahn, Robert W., and Lifshin, Eric. (eds.) 1993. *Concise Encyclopedia of Materials Characterization*. Oxford: Pergamon Press.
- Callender, Craig. 2004. Measures, Explanations and the Past: Should 'Special' Initial Conditions be Explained? *British Journal for the Philosophy of Science* 55 (2):195–217.
- Cappelen, Herman, and Lepore, Ernie. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- Chomsky, Noam. 1970. Remarks on Nominalization. In Roderick A. Jacobs and Peter S. Rosenbaum (eds.), *Readings in English Transformational Grammar* (pp. 184–221). Waltham: Ginn & Co.

- Chomsky, Noam. 1976. On the Nature of Language. *Annals of the New York Academy of Sciences* 280 (1):46–57.
- Chomsky, Noam. 1995. Language and Nature. *Mind* 104 (413):1–61.
- Chomsky, Noam. 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Collins, John. 2009. Methodology, Not Metaphysics: Against Semantic Externalism. *Proceedings of the Aristotelian Society Supplementary Volume* 83 (1):53–69.
- Coyle, Diane. 2014. *GDP: A Brief but Affectionate History*. Princeton: Princeton University Press.
- Dasgupta, Shamik. 2009. Individuals: An Essay in Revisionary Metaphysics. *Philosophical Studies* 145 (1):35–67.
- Dasgupta, Shamik. 2018. Realism and the Absence of Value. *Philosophical Review* 127 (3):279–322.
- Dorr, Cian. 2010. Of Numbers and Electrons. *Proceedings of the Aristotelian Society* 90 (2):133–81.
- Dorr, Cian, and Hawthorne, John. 2013. Naturalness. In Karen Bennett and Dean W. Zimmerman (eds.), *Oxford Studies in Metaphysics*, vol. 8 (pp. 3–77) Oxford: Oxford University Press.
- Eddon, Maya. 2013. Fundamental Properties of Fundamental Properties. In Karen Bennett and Dean W. Zimmerman (eds.), *Oxford Studies in Metaphysics*, vol. 8 (pp. 78–104). Oxford: Oxford University Press.
- Enoch, David. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Ereshefsky, Marc. 2008. Species. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Fall. Stanford: Metaphysics Research Lab, Stanford University.
- Fantl, Jeremy. 2006. Is Metaethics Morally Neutral? *Pacific Philosophical Quarterly* 87 (1):24–44.
- Fine, Kit. 2012. Guide to Ground. In Fabrice Correia and Benjamin Schnieder (eds.), *Metaphysical Grounding: Understanding the Structure of Reality* (37–80). Cambridge: Cambridge University Press.
- Fodor, Jerry A. 1974. Special Sciences (Or: The Disunity of Science as a Working Hypothesis). *Synthese* 28 (2):97–15.
- Fodor, Jerry A. 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press.
- van Fraassen, Bas C. 1980. *The Scientific Image*. Oxford: Oxford University Press.
- Franklin-Hall, Laura R. 2007. Bacteria, Sex, and Systematics. *Philosophy of Science* 74 (1):69–95. Originally published under the name “Laura R. Franklin”.
- Franklin-Hall, Laura R. 2015. Natural Kinds as Categorical Bottlenecks. *Philosophical Studies* 172 (4):925–48.
- Friedman, Michael. 1974. Explanation and Scientific Understanding. *Journal of Philosophy* 71 (1):5–19.
- Gärdenfors, Peter. 2000. *Conceptual Spaces: The Geometry of Thought*. Cambridge, MA: MIT Press.
- Gerde, Eric, and Marder, M. 2001. Friction and Fracture. *Nature* 413 (6853):285–8.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Greco, Daniel. 2015. Epistemological Open Questions. *Australasian Journal of Philosophy* 93 (3):509–23.
- Grimm, Stephen R. 2008. Explanatory Inquiry and the Need for Explanation. *British Journal for the Philosophy of Science* 59 (3):481–97.
- Hall, Ned. 2011. The Large-Scale Joints of the World. *Humana.Mente* 19:11–39.

- Hammond, Peter J. 1988. Consequentialist Foundations for Expected Utility Theory. *Theory and Decision* 25 (1):25–78.
- Haslanger, Sally. 2000. Gender and Race: (What) Are They? (What) Do We Want Them to Be? *Noûs* 34 (1):31–55.
- Hazlett, Allan. 2017. Understanding and Structure. In Stephen R. Grimm (ed.), *Making Sense of the World: New Essays on the Philosophy of Understanding*. Oxford: Oxford University Press.
- Hendry, Robin Findlay. 2010. The Elements and Conceptual Change. In Helen Beebe and Nigel Sabbarton-Leary (eds.), *The Semantics and Metaphysics of Natural Kinds* (pp. 137–58). Oxford: Routledge.
- Kanamori, Akihiro. 2009. *The Higher Infinite: Large Cardinals in Set Theory from Their Beginnings*, 2nd edn. Berlin: Springer.
- Kitcher, Philip. 1984. Species. *Philosophy of Science* 51:308–33.
- Kitcher, Philip. 2008. Carnap and the Caterpillar. *Philosophical Topics* 36 (1):111–27.
- Koellner, Peter. 2014. Large Cardinals and Determinacy. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Spring. Stanford: Metaphysics Research Lab, Stanford University.
- Kornblith, Hilary. 1993. *Inductive Inference and Its Natural Ground*. Cambridge, MA: MIT Press.
- Krim, Jacqueline. 1996. Friction at the Atomic Scale. *Scientific American* 275 (4):74–80.
- Larson, Paul B. 2012. A Brief History of Determinacy. In Akihiro Kanamori, Dov Gabbay, and John Woods (eds.), *Handbook of the History of Logic*, Volume 6: *Sets and Extensions in the Twentieth Century* (pp. 457–508). Amsterdam: North Holland.
- Lasersohn, Paul. 2018. Common Nouns as Variables: Evidence from Conservativity and the Temperature Paradox. In Robert Truswell, Chris Cummins, Caroline Heycock, Brian Rabern, and Hannah Rohde (eds.), *Proceedings of Sinn und Bedeutung 21*, vol. 2 (pp. 731–46).
- Levine, Joseph. 1983. Materialism and Qualia: The Explanatory Gap. *Pacific Philosophical Quarterly* 64:354–61.
- Lewis, David. 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy* 61 (4):343–77. Reprinted in Lewis (1999: 8–55).
- Lewis, David. 1984. Putnam's Paradox. *Australasian Journal of Philosophy* 62 (3):221–36. Reprinted in Lewis (1999: 56–77).
- Lewis, David. 1991. *Parts of Classes*. Oxford: Basil Blackwell.
- Lewis, David. 1999. *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Lipton, Peter. 2004. *Inference to the Best Explanation*, 2nd edn. London: Routledge.
- Maddy, Penelopea. 2011. *Defending the Axioms: On the Philosophical Foundations of Set Theory*. Oxford: Oxford University Press.
- Mallon, Ron. 2003. Social Construction, Social Roles, and Stability. *Socializing Metaphysics: The Nature of Social Reality* (pp. 327–353). Lanham: Rowman & Littlefield.
- Malzbender, J., den Toonder, J.M.J., Balkenende, A.R., de With, and G. 2002. Measuring Mechanical Properties of Coatings: A Methodology Applied to Nanoparticle-Filled Sol-Gel Coatings on Glass. *Materials Science and Engineering: R: Reports* 36 (2):47–103.
- Margolis, Eric, and Laurence, Stephen. (eds.) 1999. *Concepts: Core Readings*. Cambridge, MA: The MIT Press.
- Margolis, Eric, and Laurence, Stephen. 2014. Concepts. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Spring. Stanford: Metaphysics Research Lab, Stanford University.

- Martin, Donald A. 1998. Mathematical Evidence. In H. Garth Dales and Gianluigi Oliveri (eds.), *Truth in Mathematics* (pp. 215–231). Oxford: Clarendon Press.
- Mayr, Ernst. 1996. What Is a Species, and What Is Not? *Philosophy of Science* 63 (2):262–77.
- Millikan, Ruth Garrett. 1999. Historical Kinds and the “Special Sciences”. *Philosophical Studies* 95 (1):45–65.
- Miyoshi, Kazuhisa. 1999. Aerospace Mechanisms and Tribology Technology: Case Study. *Tribology International* 32 (11):673–85.
- Montague, Richard. 1973. The Proper Treatment of Quantification in Ordinary English. In Jaakko Hintikka, Julius Moravcsik, and Patrick Suppes (eds.), *Approaches to Natural Language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics* (pp. 221–242). Dordrecht: Reidel.
- Morgan, Thomas H., Sturtevant, Alfred H., Muller, Herman J., Bridges, Calvin B. 1915. *The Mechanism of Mendelian Heredity*. New York: Henry Holt.
- Newey, Charles, and Weaver, Graham. (eds.) 2013. *Materials Principles and Practice: Electronic Materials Manufacturing with Materials Structural Materials*. Amsterdam: Elsevier.
- Nolan, Daniel. 1999. Is Fertility Virtuous in its Own Right? *British Journal for the Philosophy of Science* 50 (2):265–82.
- Oddie, Graham. 2005. *Value, Reality, and Desire*. Oxford: Oxford University Press.
- Pérez Carballo, Alejandro. 2016. Structuring Logical Space. *Philosophy and Phenomenological Research*. 92 (2):460–91. Early View, 2014.
- Pietroski, Paul. 2003. The Character of Natural Language Semantics. In Alex Barber (ed.), *Epistemology of Language* (pp. 217–56). Oxford: Oxford University Press.
- Portmore, Douglas W. 2009. Consequentializing. *Philosophy Compass* 4 (2):329–47.
- Rosen, Gideon. 1994. Objectivity and Modern Idealism: What is the Question? In Michaelis Michael and John O’Leary-Hawthorne (eds.), *Philosophy in Mind*. Dordrecht: Kluwer.
- Rosen, Gideon. 2010. Metaphysical Dependence: Grounding and Reduction. In Bob Hale and Aviv Hoffmann (eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford: Oxford University Press.
- Rosen, Gideon. 2015. Real Definition. *Analytic Philosophy* 56 (3):189–209.
- Rudd, R.E., and Broughton, J.Q. 2000. Concurrent Coupling of Length Scales in Solid State Systems. *physica status solidi (b)* 217 (1):251–91.
- Schaffer, Jonathan. 2004. Two Conceptions of Sparse Properties. *Pacific Philosophical Quarterly* 85 (1):92–102.
- Schärp, Kevin. 2013. *Replacing Truth*. Oxford: Oxford University Press.
- Schwalbe, Ulrich, and Walker, Paul. 2001. Zermelo and the Early History of Game Theory. *Games and Economic Behavior* 34 (1):123–37.
- Shelah, Saharon, and Woodin, Hugh. 1990. Large Cardinals Imply that Every Reasonably Definable Set of Reals is Lebesgue Measurable. *Israel Journal of Mathematics* 70 (3):381–94.
- Sider, Theodore. 1996. Naturalness and Arbitrariness. *Philosophical Studies* 81 (2):283–301.
- Sider, Theodore. 2013. *Writing the Book of the World*. Oxford: Oxford University Press.
- Stalnaker, Robert C. 2002. Epistemic Consequentialism. *Proceedings of the Aristotelian Society Supplementary Volume* 76 (1):153–68.
- Stanford, P. Kyle. 1995. For Pluralism and against Realism about Species. *Philosophy of Science* 62 (1):70–91.
- Stoljar, Daniel. 2015. Chomsky, London and Lewis. *Analysis* 75 (1):16–22.
- Tappenden, Jamie. 2005. Proof Style and Understanding in Mathematics I: Visualization, Unification and Axiom Choice. In Paolo Mancosu, Klaus Jørgensen, and Stig Pedersen (eds.), *Visualization, Explanation and Reasoning Styles in Mathematics* (pp. 147–214). Dordrecht: Springer.

- Tappenden, Jamie. 2008a. Mathematical Concepts and Definitions. In Paolo Mancosu (ed.), *The Philosophy of Mathematical Practice* (pp. 256–75). Oxford: Oxford University Press.
- Tappenden, Jamie. 2008b. Mathematical Concepts: Fruitfulness and Naturalness. In Paolo Mancosu (ed.), *The Philosophy of Mathematical Practice* (pp. 276–301). Oxford: Oxford University Press.
- Taylor, Barry. 1993. On Natural Properties in Metaphysics. *Mind* 102 (405):81–100.
- Thomasson, Amie L. Chapter 21, this volume. A Pragmatic Method for Conceptual Ethics. In Alexis Burgess, Herman Cappelen, and David Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Varzi, Achille. 2015. Boundary. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Winter. Stanford: Metaphysics Research Lab, Stanford University.
- Walley, Stephen M. 2012. Historical Origins of Indentation Hardness Testing. *Materials Science and Technology* 28 (9–10):1028–44.
- Welch, P. D. 2015. Large Cardinals, Inner Models, and Determinacy: An Introductory Overview. *Notre Dame Journal of Formal Logic* 56 (1):213–42.
- White, Roger. 2005. Explanation as a Guide to Induction. *Philosopher's Imprint* 5 (2):1–29.
- White, Roger. 2007. Does Origins of Life Research Rest on a Mistake? *Noûs* 41 (3):453–77.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wilson, Mark. 1982. Predicate Meets Property. *Philosophical Review* 91 (4):549–89.
- Wilson, Mark. 2006. *Wandering Significance: An Essay on Conceptual Behavior*. Oxford: Oxford University Press.
- Wilson, Mark. 2008. Beware of the Blob: Cautions for Would-Be Metaphysicians. In Dean Zimmerman (ed.), *Oxford Studies in Metaphysics*, vol. 4 (pp. 275–320). Oxford: Oxford University Press.
- Wilson, Mark. 2013. What is Classical Mechanics Anyway? In Robert W. Batterman (ed.), *Oxford Handbook of the Philosophy of Physics* (pp. 43–106). Oxford: Oxford University Press.
- Wong, Wai-hung, and Yudell, Zanja. 2015. A Normative Account of the Need for Explanation. *Synthese* 192 (9):2863–85.
- Yalcin, Seth. 2011. Nonfactualism about Epistemic Modality. In Andy Egan and Brian Weatherson (eds.), *Epistemic Modality* (pp. 295–334). Oxford: Oxford University Press.
- Zhang, J., Johnston, J., Chattopadhyay, A. 2014. Physics-based Multiscale Damage Criterion for Fatigue Crack Prediction in Aluminium Alloy. *Fatigue and Fracture of Engineering Materials and Structures* 37 (2):119–31.
- Zimmerman, Dean. (ed.) 2008. *Oxford Studies in Metaphysics*, vol. 4. Oxford: Oxford University Press.