

1

Introduction

A Guided Tour of Conceptual Engineering and Conceptual Ethics

Herman Cappelen and David Plunkett

Introduction

In *The Will to Power*, Nietzsche writes the following:

Philosophers . . . have trusted in concepts as completely as they have mistrusted the senses: they have not stopped to consider that concepts and words are our inheritance from ages in which thinking was very modest and unclear. . . . What dawns on philosophers last of all: they must no longer accept concepts as a gift, nor merely purify and polish them, but first make and create them, present them and make them convincing. Hitherto one has generally trusted one's concepts as if they were a wonderful dowry from some sort of wonderland: but they are, after all, the inheritance from our most remote, most foolish as well as most intelligent ancestors. . . . What is needed above all is an absolute skepticism toward all inherited concepts.¹

Nietzsche here articulates a radical skepticism about all inherited concepts. Philosophers should question whether the concepts we have are good enough and should engage in conceptual critique. What emerges, thinks Nietzsche, is the following: we should not just improve the concepts we've been given, reforming or "polishing" them in minor ways, but also create new ones—concepts not tainted by the "most foolish of our ancestors".

Even if you think Nietzsche's claim is more than a bit hyperbolic, you might think some more moderate version of his view is justified. For example: maybe *some* of the concepts we have inherited are defective, or at least not as good as they could be for our current purposes. In contrast, maybe you think Nietzsche is wildly off here in his radical stance. Maybe we have good reason to think that our current stock of concepts is just fine for the purposes at hand. Or maybe you think that, even if that stock of concepts could be better, it doesn't make sense to think about changing our concepts—or at least not the core concepts that really

¹ Nietzsche (1901/1968: 220–1, section 409). Thanks to Michael Beaney for pointing us to this passage.

matter in philosophy (e.g., TRUTH, MEANING, CONTENT, or VALUE).² The issues that this quote from Nietzsche brings up thus includes questions such as: What are the dimensions of assessment for concepts? Which philosophical concepts are defective and how can we improve them? How important are facts about the history (or “genealogy”) of our use of concepts to the assessment of our current concepts? These are questions at the heart of the fields that we call “conceptual engineering” and “conceptual ethics”.

If you care about these and related questions, this volume should be of interest to you. It is the first ever volume devoted entirely to conceptual engineering and conceptual ethics. Our hope is that it will help shape and promote what we (the editors) take to be an important, exciting, and underexplored part of philosophy. In this Introduction, we first try to delineate the field and explain why we are using two different expressions (‘conceptual engineering’ and ‘conceptual ethics’) to describe the topics in the book. We then turn to some of the central foundational issues that arise for conceptual engineering and conceptual ethics, and finally we outline various views one might have about their role in philosophy and inquiry more generally.

1. ‘Conceptual Engineering’ and ‘Conceptual Ethics’

The title of this volume uses two expression to describe its topic: ‘conceptual engineering’ and ‘conceptual ethics’. Why? The answer to this is not straightforward. We don’t think these expressions come with fixed meanings. The previous literature has used them in different ways and so do the authors in this volume. These terms are often used without precise definitions by those working in the field. Moreover, when they are given more precise definitions by philosophers, these definitions often contradict those given by others. As editors, we could have played the terminology police for those contributing to this volume. But that would be an exercise in futility. Instead, we have decided to let a thousand (or at least a few) flowers bloom. Contributors use central terms, such as ‘conceptual engineering’, ‘conceptual ethics’, ‘revision’, and ‘amelioration’, in different ways, often explicitly so. That’s how it should be given that this is currently a fast moving literature involving philosophers from many different background and sub-fields. That said, for the purposes of this Introduction, we will offer our own characterizations of conceptual engineering and conceptual ethics, with one of us (Cappelen) offering a characterization of conceptual engineering, and the other of us (Plunkett) offering one of conceptual ethics.³ The basic reason we split up our discussion in this way is that one of us (Cappelen) likes to use the expression “conceptual engineering”, whereas the other one (Plunkett) thinks

² In this chapter, we use small caps (e.g., CAT) to pick out concepts, single quotation marks (e.g., ‘cat’) strictly to mention linguistic items, and double quotation marks (e.g., “cat”) for a variety of tasks, including quoting others’ words, scare quotes, and mixes of use and mention.

³ Our use of the term ‘conceptual engineering’ stems from Cappelen (2018), drawing chiefly on Scharp (2013) and Eklund (2015). Simon Blackburn also briefly uses ‘conceptual engineering’ in the opening pages of Blackburn (2001) in a related way, as does Brandom (2001). Our use of the term ‘conceptual ethics’ draws from Burgess and Plunkett (2013a,b).

that a number of the issues involved are best described as issues in “conceptual ethics” rather than “conceptual engineering”.

Conceptual Engineering

According to Cappelen (2018), *conceptual engineering is concerned with the assessment and improvement of concepts*. However, since it’s unclear and controversial what concepts are (and whether there are any), it’s better to broaden the scope along the following lines:

Conceptual engineering = (i) The assessment of representational devices, (ii) reflections on and proposal for how to improve representational devices, and (iii) efforts to implement the proposed improvements.⁴

Here are some issues that are central for those working on conceptual engineering:

- What are the relevant representational devices? Possible answers include: concepts (as they are construed in some part of psychology or philosophy), lexical items, and the semantic values of lexical items.⁵ A closely connected cluster of questions concerns whether they are in language or thought or both. Different conceptual engineers will give different answers and that will have enormous implications for how the field is understood and practiced.
- Given an answer to the first cluster of questions, we can ask: What kinds of defects can representational devices have? Throughout the history of philosophy, a variety of defects have been proposed: cognitive defects (that undermine our ability to reason properly), moral or political defects (that undermine moral or political values of various sorts), theoretical defects (that undermine progress within some theoretical field), or semantic defects (where the semantic value is incoherent, incomplete, or missing). For illustrations of all of these and a more detailed taxonomy of defects, see Cappelen (2018: chapter 2).
- Once you have detected a defect in a representational device you care about, it’s natural to think about how to improve it. What are the ameliorative strategies? There are four basic options once you’ve identified a defect in C: (i) Do nothing—just live with it (can’t improve it, can’t get rid of it), (ii) Abandonment of C (it’s so defective it can’t be improved), (iii) Improvement of C, (iv) Replacement of C (for certain purposes, in certain contexts).
- Once you have settled on an ameliorative strategy, you might want to do some work to implement it, that is, you might want to engage in a bit of activism on behalf of your ameliorative strategy. If that’s something you want to do, it raises an ‘implementation challenge’: how are ameliorative strategies best implemented?

⁴ Why call it ‘conceptual’ engineering when it’s about representational devices more generally? Purely for aesthetic reason: ‘representational devices engineering’ doesn’t roll off the tongue in the way ‘conceptual engineering’ does.

⁵ Cappelen (2018) suggests we think of the basic case as that of revising the extensions and intensions of expressions.

If you're interested in conceptual engineering, you don't need to focus on all of these issues. Some will focus on discovering defects, some on ameliorative strategies, others on conceptual activism, and yet others want to do the whole shebang.

Conceptual engineering is not usefully construed as a branch of any other part of philosophy. It will draw on insights from philosophy of language, philosophy of mind, epistemology, political philosophy, philosophy of science, ethics, and other fields. That, of course, is also true about these other fields (i.e., they will draw on insights from each other). A case can, however, be made that conceptual engineering is prior to or more fundamental than all other philosophical disciplines. The argument is simple and obvious: reflection and argumentation in any part of philosophy must rely on concepts (epistemology relies heavily on, e.g., KNOWLEDGE and JUSTIFICATION; ethics on, e.g., OUGHT and BAD; and so on for each branch of philosophy.) As Nietzsche correctly observes in the passage at the beginning of this Introduction: it's intellectually irresponsible to throw yourself headlong into an inquiry before questioning the concepts you're using in that inquiry. So conceptual engineering, as construed here, comes first.⁶

Conceptual Ethics

Now that we have a rough characterization of conceptual engineering from one of us (Cappelen), here is a rough one of conceptual ethics, endorsed by the other one of us (Plunkett), drawing on previous co-authored work with Alexis Burgess.⁷ Broadly, conceptual ethics concerns a range of normative and evaluative issues about thought, talk, and representation. Those include issues about which concepts we should use, ways in which concepts can be defective, what we should mean by our words, and when we should refrain from using certain words. (Which issues one thinks belong on this list, as well as how these issues are related to each other, will obviously depend on one's further philosophical commitments.) As the label suggests, some of the core issues in conceptual ethics concern *concepts* (assuming, for now, that there are such things). These include, centrally, normative issues about which concepts one should use (and why) and evaluative issues about which concepts are better than others (and why). Concepts can here be understood in rough terms as *constituent components of thoughts*, leaving it for different theorists to fill out that schematic characterization in different ways. As with conceptual engineering, parallel questions can of course arise for other representational devices beyond concepts (e.g., words).

The use of the term 'ethics' here in 'conceptual ethics' is meant very broadly, to cover "both the study of what one should or ought to do (dually, what can permissibly be done) as well as the study of which actions and outcomes are good or bad, better or worse".⁸ Thus, this use of 'ethics' is *not* meant to privilege moral/political norms in particular (vs., e.g., those that find their central home in epistemology, metaphysics, aesthetics, etc.).⁹

⁶ This raises tricky questions about the nature of the concepts used to think about conceptual engineering—for more on that, see Cappelen (2018: chapter 1).

⁷ Burgess and Plunkett (2013a,b). ⁸ Burgess and Plunkett (2013a: 1094).

⁹ 'Conceptual ethics' is obviously not an ideal label. Many will still hear 'ethics' more narrowly—for example, as tied to distinctively *practical* norms of the sort that have their home in moral and political

Conceptual ethics is a branch of normative and evaluative inquiry, just as (at least certain parts of) epistemology, aesthetics, ethics, and political philosophy can be understood as branches of it. Thus, just as with other branches of normative and evaluative inquiry, people can approach conceptual ethics with very different philosophical commitments, very different views about how to make progress in it, and very different substantive views within it. Moreover, they can also approach it for very different reasons. For example, someone might be interested in conceptual ethics purely as an interesting part of philosophical theorizing. Or one might be interested in it because one is trying to actually change existing thought and talk.

This second point brings out an important aspect of the relation between conceptual ethics and those components of conceptual engineering that involve trying to actually change conceptual or linguistic practices. If one has such practical aims, then studying conceptual ethics might help. This is parallel to how studying normative political philosophy might help those interested in changing actual existing political institutions, or how studying normative aesthetics might help those creating art. But just as it would be a mistake to think of political philosophy *solely* in terms of the role it might play for the project of creating better political institutions, so too would it be a mistake to think of conceptual ethics solely in terms of the role it might play in practical projects of changing actual conceptual or linguistic practices. Conceptual ethics is a free-standing area of normative and evaluative inquiry, and some questions in it that might turn out to be of little use to those involved in actually trying to change conceptual or linguistic practices, or to those trying to engineer concepts.

In what follows, we will use the term ‘conceptual engineering’ and ‘conceptual ethics’ in roughly the ways introduced here. How exactly conceptual engineering and conceptual ethics relate to each other is something which there is live debate about in this area, and which the editors of this volume themselves have spent considerable time arguing about. Many of those engaged in these debates use the term ‘conceptual engineering’ and ‘conceptual ethics’ in ways that differ from our own. With that in mind, it should be emphasized that all of the issues we discuss below can be translated into other ways of using the labels ‘conceptual engineering’ and ‘conceptual ethics’ preferred by different philosophers (including some in this volume).

2. A Bottom-Up Characterization of Conceptual Engineering and Conceptual Ethics

We just gave what can be thought of as a top-down characterization of the topic of this volume. Another way to approach this is through examples. We could point to a range of paradigmatic cases as exemplifying conceptual engineering and conceptual ethics. We then hope that audiences will be able to find some relevant cluster of similarities between the examples.¹⁰ Here are four paradigmatic cases:

philosophy. But other possible labels have their own drawbacks. For example: ‘conceptual assessment’, another possible label here, makes it sound as if this area solely concerns *evaluative* claims about concepts (e.g., which concepts are better than others), leaving out *normative* claims (e.g., about which concepts an agent should use).

¹⁰ See Burgess and Plunkett (2013a) for a similar approach.

- Carnapian explication is an effort to improve on defective meanings. For Carnap the central defects have to do with vagueness and indeterminacy. Improvements—what Carnap calls “explications”—produce better meanings. Improvements for Carnap focus primarily on theoretical virtues.¹¹
- Much of Sally Haslanger’s work on race and gender has two components: it criticizes our gender and race concepts, and then suggests ameliorative strategies to improve those concepts. The defects she finds have to do with social and political effects of the meanings, and the ameliorations are also measured along those dimensions (e.g., by how much they can help us in the pursuit of social justice).¹²
- Peter Railton argues that moral philosophy should involve a methodology that is largely continuous with that of the natural and social sciences (this is the core of his methodological naturalism). Based on what he sees as the best practices within scientific inquiry, he argues that, in doing moral philosophy, we shouldn’t just rely on our folk concepts. Instead, we should reform the meanings of our words to zero in on the topics that really matter, and in providing explanations of the phenomena at hand. Railton then offers improved moral language that can be used in this way, including, for example, reforming definitions of key terms such as ‘moral goodness’ and ‘morality’.¹³
- Matti Eklund and Kevin Scharp explore the idea that TRUTH is inconsistent. If it is, that might be a serious defect (at least in some contexts). What might be needed is an improved, consistent, truth concept (or, in Scharp’s case, multiple concepts). We might, they claim, be okay with using an inconsistent concept in certain areas of our life. But, for the purposes of doing advanced theoretical work in such areas as linguistics and logic, it would be better to avoid doing that, if possible.¹⁴

Here’s the bottom-up way to introduce the topics of this volume: It’s about *that kind of activity* or *these kinds of issues*. The kind of thing they are doing, or the kinds of things they are discussing.¹⁵ The assumption then is that those activities form an interesting kind—a kind of activity or method or subject matter (or maybe all three at once). The idea of there being *some* sort of interesting kind here is a working hypothesis. One challenge for those working in the field is to try to substantiate it. Some chapters in this volume support that assumption, and some argue against it.

¹¹ See Carnap (1947/1956).

¹² See Haslanger (2000). It should be noted that Haslanger’s views on conceptual engineering (and associated issues about the nature of concepts) have changed over time. So what we present here is only one strand of her thinking about the topic. See the collected papers in Haslanger (2012) for a fuller view of her thinking on the topic.

¹³ See Railton (1986a,b), as well as many of the other papers collected in Railton (2003).

¹⁴ See Eklund (2002) and Scharp (2013).

¹⁵ To keep this Introduction at manageable length, we have chosen to not go into great detail of specific cases. But to get a real sense of how the bottom-up approach would work, more details would obviously be needed. For more discussion of examples, see Burgess and Plunkett (2013a,b) and Cappelen (2018).

3. Central Challenges

We turn now to what we take to be some of the central challenges in the areas of conceptual engineering and conceptual ethics. In what follows, we will often put the issues involved in terms of “conceptual engineering” rather than “conceptual ethics”. We do so both for ease of exposition, and because of our belief that engaging with issues in conceptual ethics is an important part of conceptual engineering. However, it should be kept in mind that many of the issues we discuss below apply equally to conceptual engineering and conceptual ethics.

As a heuristic, it’s useful to divide these into two categories:

- Category 1: Domain-general issues.
- Category 2: Issues that arise in thinking about the evaluation and engineering of specific concepts or groups of concepts (e.g., race and gender concepts).

As we’ll emphasize below, there’s arguably no sharp distinction here, and (at the very least) there are many important connections between category 1 and category 2. It’s a rough division. Our focus here (and in the volume as a whole) is on category 1, and how it interacts with category 2. Category 2 is huge, and the details are too diverse for us to even begin to cover them in this brief introductory essay.

In what follows, we discuss six clusters of interconnected issues that arise in work on conceptual engineering and conceptual ethics. Many of these overlap and all of them interact in various ways. It goes beyond the scope of this Introduction to explore all those important connections, so we simply list a range of issues that strike us as both interesting and important, and that we think will be central to debates in conceptual engineering and conceptual ethics in the years ahead.¹⁶

Cluster 1

What are the objects being assessed and improved (and do they exist)? If the aim of conceptual engineering is to assess and improve concepts—or other representational devices—then we are ultimately on the hook for an account of what these objects are. There are very many theories of concepts.¹⁷ Moreover, the term ‘concept’ is used in a variety of ways in philosophy, psychology, and ordinary speech. It is clear that not everyone in the debate means the same thing by the word ‘concept’, or is offering a theory of the same thing. If the objects of assessment are concepts, then one needs a theory of these things that makes it possible for them to be assessed and improved. Furthermore, one needs such a theory on which either their identity conditions are compatible with the idea of amelioration of a single concept, or where it makes sense to think of moving from one concept to another as a form of improvement. If the objects are not concepts, but something else, then that “something else” needs to be characterized (and, again, in a way that meets the above constraint that makes sense of the possibility of some kind of amelioration). For example, if the relevant

¹⁶ Our list of central issues here draws heavily on our previous work, especially Burgess and Plunkett (2013a,b) and Cappelen (2018).

¹⁷ See Margolis and Laurence (1999) for a collection of papers advocating different approaches.

objects are words or meanings, then those things also need to be characterized. Moreover, *whatever* the objects are that are being criticized and improved, there is the threat of skepticism that they actually exist.¹⁸ (Think here, e.g., about Quinean skepticism about the analytic/synthetic distinction as the basis for a certain kind of skepticism about the existence of word “meanings”).¹⁹

Concepts vs. conceptions vs. beliefs: Many people involved in conceptual engineering distinguish concepts from beliefs—for example, understanding concepts as “constituent components” of beliefs (or other attitudes) in one way or another.²⁰ Many others also include “conceptions” as part of the picture here. How exactly are these (and related) things distinguished? How *should* they be distinguished? If the answers here turn out certain ways, then perhaps many instances of conceptual engineering turn out not to be about engineering “concepts” at all, but rather something else.

Metasemantic foundations: Some theorists think of “metasemantics” as the study of the metaphysical foundations of meaning: it provides various accounts of what makes it the case (or grounds) that our words have the meanings they have.²¹ So understood, metasemantics is very important for conceptual engineering (and for conceptual activists in particular): in order to change a meaning, you need to act on (or change in some way) that metaphysical foundation. So different views of metasemantics will generate different views of what conceptual engineering consists in. The parallel point applies about the import of work on the metaphysical foundations of other kinds of content, if one is seeking to assess and improve those kinds of content.

Internalism vs. externalism: An important issue in metasemantics is the distinction between internalism and externalism. In rough terms, according to internalists, meaning depends on facts about the individual (e.g. facts about “what’s in her head”). Externalists deny this: they think meaning is determined at least in part by facts having to do with the history of linguistic usage, or complex use patterns over time, or the judgments of experts, or other things that are not individualistic. This divide might be put in terms of grounding (roughly, asymmetric metaphysical dependence of a certain kind) or supervenience (understood as a purely modal notion). There are important theoretical choice points here in the internalism vs. externalism debate, and the answers will profoundly affect how one thinks about and practices conceptual engineering.²² Some of the key issues here extend beyond just whether one is an internalist or externalist (about either mental or linguistic content). They also include intramural debates among each camp. For example, does the correct

¹⁸ An important question here is whether it is easier to meet this challenge when moving to one of these other objects instead of concepts. One of us (Cappelen) thinks it is. He argues that conceptual engineering concerns extensions/intensions of lexical items, rather than concepts. See Cappelen (2018) for discussion.

¹⁹ See, for example, Quine (1951).

²⁰ For example, see the gloss of concepts in Burgess and Plunkett (2013a).

²¹ See Burgess and Sherman (2014).

²² The literature here is vast. An overview and helpful bibliography is Lau and Deutsch (2016).

externalism involve facts about the future usage of speakers,²³ or is it perhaps tied to “correct” theories of the relevant subject matter in some (perhaps indirect) way.²⁴

Normativity of meaning: How sharp is the distinction between doing descriptive vs. normative work when engaged in the study of concepts, or of other representational devices? On some views, they will necessarily be deeply intertwined. On others, they are pretty far apart. And there is obviously room for a lot of middle ground here. Different theories in metasemantics (either about words or concepts) help inform these different views, as do different views about our thought and talk about semantics. For example, consider discussions about the “normativity of meaning”.²⁵ If claims about meaning (or about mental content) are themselves always normative in some sense, how does this affect our understanding of the relationship of descriptive and normative work on concepts (or other representational devices)? For example: does it collapse the distinction between normative theorizing about concepts (“conceptual ethics”) and descriptive theorizing about concepts? Or are the kinds of normative issues here importantly distinct in some sense?

Within our control vs. outside of our control: Many of those working on conceptual engineering are interested not just in general theories of conceptual engineering, but in actually bringing about conceptual or linguistic change. If that is part of one’s motivation, it’s important to get clear on whether conceptual change (or meaning change) is something that is within our control. And, if so, how *much* of it can we control, and how well can we control it? We need to get clear on the extent to which these kinds of changes are governed by our decisions, intentions, agreements, and preferences. Maybe, instead, they are governed by mechanisms that are difficult to understand and outside of our control.²⁶ One way these issues about control matter is for evaluating strategies for conceptual activism.

Cluster 2

What are the norms, goods, values, etc., that determine the normative/evaluative facts in conceptual ethics? Philosophers appeal to a wide variety of good, norms, values, etc. in making normative/evaluative claims in conceptual ethics. For example, some appeal to facts about what fundamental reality is like, independent of our thinking about it.²⁷ Others appeal to practical considerations of what would aid us in theoretical inquiry on a given topic, for example, by helping us zero in on an important set of issues, or helping us avoid false beliefs, or helping us smoothly communicate with other inquirers.²⁸ Others still appeal to practical considerations about the ethical/political effects of the use of certain concepts—for example, the way that the use of certain concepts might help promote social justice, freedom, or happiness.²⁹ These different

²³ See, for example, the sort of view advocated for by Derek Ball in his chapter, “Revisionary Analysis without Meaning Change (Or, Could Women Be Analytically Oppressed)?” (Chapter 2, this volume).

²⁴ See, for example, the sort of view advocated for in Schroeter and Schroeter (2014).

²⁵ See Kripke (1982) and Gibbard (2012) for defenses of the idea that “meaning is normative”, and for criticisms of the idea, see, for example, Boghossian (1989) and Hattiangadi (2017).

²⁶ For an argument that much of the relevant changes are not in our control, see Cappelen (2018).

²⁷ See, for example, Sider (2012).

²⁸ See, for example, Eklund (2002) and Scharp (2013).

²⁹ See, for example, Haslanger (2000).

values/norms can interact in any number of ways. For example: it might be that using concept C helps promote social justice because it helps us keep track of important features of social reality we should be studying. This diversity of goods/values/norms raises the question of which of them ultimately matter in conceptual ethics. And how much is there is a general answer here, anyway, as opposed to answers for specific kinds of concepts (e.g., for race and gender concepts, or for truth concepts, etc.)?

What are potential defects and virtues of concepts? What counts as a defect of a concept and what are the potential dimensions of improvement? In the history of this topic, there have been a broad range of proposals. We can classify these into four rough categories: moral-political (e.g., hindering or promoting social justice), epistemological (e.g., hindering or promoting the acquisition of knowledge), cognitive (e.g., hindering or promoting good cognitive functioning), metaphysical defects (e.g., corresponding or not to joints in reality). Which of these proposals help us locate (and explain) genuine defects, and how are those defects related to each other?³⁰

How much do aims matter? We use concepts in very different contexts, in cases where we aim to accomplish very different things—for example, making progress in mathematical inquiry vs. winning an election vs. trying to help build a better society. Many claims in conceptual engineering appeal to the aims an agent has in using a given concept (or set of concepts).³¹ But how much do the aims an agent have matter here? What if (e.g.) she has immoral aims, or aims that won't be productive to furthering inquiry? Maybe what matters here are the aims an agent *should* have? Or maybe aims of any sort (either the ones an agent has, or the ones an agent should have) don't play any sort of fundamental explanatory role here. Perhaps reference to aims is just an unhelpful indirect way of talking about other factors that do the real explanatory work, such as facts about where the joints in reality are, or facts about how people should live, or facts about how our social/political institutions should be organized.

Scope of claims in conceptual engineering: Many of the claims involved in conceptual engineering are aimed at particular people, in particular circumstances. For example, someone might make a claim in conceptual ethics that “we should use concept A, instead of concept B”. Who is the “we” here? It might be every rational agent, or every human being. But, much of the time, it will be a more limited group of people: for example, philosophers involved in the study of language, or people involved in a certain sort of social activism. In many cases (as noted above), the claims will often also be tied to *aims* those people have, or the purposes they have in those contexts. Thus, many claims in conceptual ethics take the form of something along the lines of: “a group of agents should use concept A, instead of B, in circumstances C, for purposes P”.

Cluster 3

What are the plausible ameliorative strategies? And which ones are better than others? Once one has found a defect in a concept, there are many types of

³⁰ For further discussion of this topic, see Cappelen (2018).

³¹ See, for example, Haslanger (2000) and Anderson (2001). For discussion, see Burgess and Plunkett (2013b).

ameliorative strategies that, at least *prima facie*, seem available. Which ameliorative strategy you endorse will be important for thinking about and engaging in conceptual activism. Here are some of those strategies:

- (i) Improve/reform the concept and then use that one.
- (ii) Replace uses of the concept with uses of an “explicated” concept which bears important similarities to the original one.
- (iii) Replace uses of the concept with uses of a very different concept.

So far we have talked about improving/reforming concepts or uses of them, without being explicit about how this is reflected in language. It’s difficult to engage in conceptual engineering without that having linguistic connections. How one thinks about this will in large part depend on how one thinks of the connections between thought and talk. Here are some views about the linguistic implementations of (i)–(iii):

- (iv) Keep the lexical item and associate it with an improved/reformed concept.
- (v) Introduce a new lexical item with associated improved concept(s).
- (vi) Complete rejection: don’t use that expression or the associated concept again.

This picture we just sketched presupposes that conceptual engineering operates on concepts and that these are then associated with (or expressed by) lexical items. As mentioned above, there are alternative views, according to which conceptual engineering operates directly on expressions and their intensions/extensions, so bypasses concepts entirely.³² On this alternative view, here are some strategies that at least *prima facie* seem available:

- (vii) Keep the lexical item and revise the intension/extension.
- (viii) Introduce a new lexical item with a new intension/extension and then let this new lexical item replace uses of the old one.
- (ix) Complete rejection: abandon the lexical item and its associated intension/extension.

As this last paragraph makes clear, how one thinks of the range of ameliorative strategies will depend in large part on what one takes the objects of conceptual engineering to be (e.g., concepts vs. conceptions vs. words), and one’s views about their nature.

Obviously some of these strategies might work well in some contexts and not in others. It is far from obvious that there is a context-invariant way to assess ameliorative strategies.

What is the difference between improving concept *c* and improving one’s beliefs about objects in the extension of *c*? Why choose one strategy over the other? An issue that keeps coming up in discussions of conceptual engineering is the difference between improving a concept, *c*, on the one hand, and improving (or revising) people’s beliefs about objects in the extension of *c*, on the other. Why and when is one strategy superior to the other? How do they interact? The answers to these questions will depend heavily on the answer to the three previous questions. It will

³² See Cappelen (2018).

depend on how one sees the connections between beliefs about *c*-objects and possession of the concept *c*—for some there’s a constitutive connection, for others there isn’t a connection at all.

Connections to ontology: Closely connected to the above point—to what extent can conceptual engineering change non-linguistic and non-conceptual aspects of the world? An obvious connection is this: if conceptual engineering succeeds in a particular case, it will change how people think, talk, and act on the (non-conceptual and non-linguistic) world. However, some conceptual engineers go further. For theorists who think conceptual engineering operates directly on extensions and intensions, the way to describe the effects of conceptual engineering (if the activism succeeds) is as a direct improvement of the non-linguistic and non-conceptual world.³³ On this view, we should not describe the effects of conceptual engineering as an amelioration of the concept of *FREEDOM*, but instead it is freedom itself that has been ameliorated. A final connection is this: for those who think conceptual engineering operates on concepts and also think that some concepts (or our use of them) can be constitutive of some element of non-conceptual reality (e.g. parts of social reality), there’s an interesting connection: amelioration of an important social concept can change the nature of the relevant part of social reality (since the concept is partly constitutive of some element of social reality).³⁴

Practical effects of conceptual engineering: How much practical import does conceptual engineering have? In what ways does it shape our thoughts, actions, and selves? And what about other parts of reality—for example, those that might depend on our conceptual practices in some way? In some cases, philosophers think the practical effects of conceptual engineering can be quite extensive and profound. But others are skeptical that engaging in conceptual engineering will have that big an impact at all (regardless of whether we are in control of that impact or not).

Cluster 4

Limits of revision and change of topic: How much revision is too much? When is a revision a complete change of topic? When would it be okay to change the topic, including perhaps completely abandoning the old topic in doing so? If you revise the concept *COW* to the concept *FOX* (or replace ‘cow’ with ‘fox’), there’s no sense in which you have improved on *COW* or ‘cow’. You’ve just started talking about foxes. You’ve changed the topic. So the general issue then is how far we can go in amelioration without a complete change in topic? What degree of change is acceptable?³⁵

Continuity of inquiry: Consider the debate over whether free will is compatible with determinism. This is a debate that has taken place over time involving many participants. Suppose that at some point during this debate that the concept *FREE WILL* is ameliorated. We now have, in some sense, a better “free will” concept. Does

³³ See Cappelen (2018).

³⁴ For example, this is true on the “dynamic nominalist” view that Hacking (2002) argues for concerning the relationship between naming practices and kinds of people (or, relatedly, social identities).

³⁵ This issue is central to Strawson’s critique of Carnap on explication in Strawson (1963). For some of the more recent discussion on the topic, see Railton (2003); Haslanger (2012); Eklund (2017); and Cappelen (2018).

this constitute a discontinuity of inquiry: are we now engaged in an investigation of a new question and have we left the old question behind? Or is there some sense in which we can still say that we are discussing the same question as before, that is, whether free will is compatible with determinism? If there has been a discontinuity, what does that mean for the ability of speakers to meaningfully disagree with each other from different sides of the divide? What does it mean for our views about intellectual progress within that inquiry, and the ability of inquiry to build toward something that is “objectively” better in some sense?³⁶

Conceptual engineering and verbal disputes: Consider again the debate over whether free will is compatible with determinism. Suppose as before that at some point the concept *FREE WILL* is ameliorated. Here is an issue that’s closely related to the issue of continuity of inquiry: Isn’t there now a significant risk of people engaging in verbal disputes? Those using the pre-ameliorated meaning for ‘free will’ say free will is compatible with determinism”, and then those using the ameliorated concept say free will is not compatible with determinism”. It will look like they are disagreeing, but since they mean different things, they might be engaged in a verbal dispute. Isn’t the entire project of engaging in conceptual engineering at risk of generating an endless amount of verbal disputes?³⁷

Cluster 5

Conceptual fixed points: Are some concepts or terms so basic that they cannot be engineered (or at least not in a way that is rational, or well-supported by reasons)? Are some concepts or terms so fundamental that we are stuck with them, meaning that evolution, revision, and amelioration are impossible? David Chalmers and Matti Eklund have defended the idea that there are conceptual fixed points (or “bedrock concepts”, as Chalmers call them).³⁸ A central challenge for such views is to identify in a principled way the bedrock concepts and explain what makes them more fixed than those that can be engineered. A separate question is whether bedrock concepts should in some sense be normatively privileged, or whether the fact that they are a bedrock is just a descriptive fact about which concepts we happen to be stuck with in some sense.

The self-reflectiveness of conceptual engineering: Many philosophers engaged in conceptual engineering are, unsurprisingly, interested in the concepts used to articulate and describe conceptual engineering itself. For example, *CONCEPT* is itself an excellent candidate for conceptual engineering. So conceptual engineering can become (and perhaps should become) self-reflective. That self-reflection may change the nature of the activity.

Hypocrisy: How much of a problem is it if one uses (and not just mentions) concept *C* to make an argument that concept *c* should be replaced or revised? Is this an

³⁶ For some of our own take on these questions, see Plunkett (2015) and Cappelen (2018). Note that the issues here about continuity of inquiry have long been at the heart of debates in the history and philosophy of science, especially in the wake of Kuhn (1962/2012).

³⁷ For more on verbal disputes, see Chalmers (2011) and Jenkins (2014).

³⁸ See Chalmers (2011) and Eklund (2015).

objectionable form of hypocrisy? Or is it better described, at least in some cases, as a form of internal critique, or addressing one's opponents "on their own terms"? If so, then perhaps it is sometimes a virtue. The issues here will be particularly important for arguments in conceptual ethics involving foundational normative concepts, such as OUGHT and VALUE, that are hard (perhaps impossible) to avoid using in conceptual ethics. Perhaps it is impossible to avoid either some kind of "vindicatory circularity" or "hypocrisy" in such cases.³⁹ Importantly, these issues of hypocrisy and ineffability not only matter for those engaged in theoretical reflection in conceptual ethics. They also matter for those engaged in conceptual activism. For example: perhaps certain ways of advocating for conceptual change inevitably involve using the very concepts one aims to criticize. If so, then there are potentially issues involving not only hypocrisy here, but also issues of misleading or lying.⁴⁰

Cluster 6

How often are we already engaged in conceptual engineering? And do we need to be aware of doing conceptual engineering in order to do it well? As we will discuss in a bit more detail later on in this Introduction, some philosophers think that much of existing philosophical inquiry involves conceptual engineering to some degree. Many hold that this engagement with conceptual engineering is going on implicitly, perhaps even without the philosophers themselves being aware that is what they are up to.⁴¹ But many are skeptical of such claims. This raises a question: how much does one need to be aware of doing conceptual engineering to count as doing it? It also raises the question: how much does one need to be aware of doing it in order to do it *well*? Is explicit engagement with conceptual engineering always better than implicit engagement with it? For example, perhaps the best methodology for pursuing conceptual engineering (at least in certain kinds of inquiry) is not to focus on conceptual engineering as such, but rather just engage in ongoing inquiry into the intuitively relevant subject matter and then just let conceptual evolution happen naturally as part of the process. For example: for physicists to improve on and introduce new concepts along the way while trying to study physical reality, but without self-consciously ever thinking about part of their activity as an exercise in conceptual engineering.

How important is it to have a correct description of our representational devices before we do conceptual engineering? We've described the aim of conceptual engineering as that of assessing and ameliorating concepts and other representational devices. How important is it to have a correct descriptive account of those devices in order to do the engineering project well? Some analogies spring to mind: to think about how to improve a particular bridge, you need to know about that bridge—the ameliorative work can't be done in isolation from the descriptive work. How helpful are such analogies for understanding the connection between the descriptive and the

³⁹ See Alexis Burgess's chapter "Never Say 'Never say 'Never''?" (Chapter 6, this volume) for further discussion, as well as Eklund (2017).

⁴⁰ See Rachel Katharine Sterken's "Linguistic Intervention and Transformative Communicative Disruptions" (Chapter 20, this volume) for further discussion.

⁴¹ For example, see Plunkett and Sundell (2013); Ludlow (2014); Plunkett (2015); and Thomasson (2016).

normative in the conceptual domain? How sharp is the distinction between the descriptive and the normative when doing conceptual engineering, or when thinking about concepts (or other representational devices) in general?

Conceptual engineering, the method of cases, and the role of intuitions in philosophy: Many philosophers think that the so-called “method of cases” is central to philosophical methodology and that intuitions about cases provide the most important kind of evidence for philosophical theories. One (controversial) way to spell out that view goes as follows: We have a concept, *c*, and our possession of that concept guides the intuitions about *C*-related thought experiments. So we can use intuitions about whether someone knows in, say, a Gettier case, as evidence of whether KNOWLEDGE applies in that case (because the concept somehow guides those intuitions). On this view, intuitions about cases reveal or illuminate core philosophical concepts and that is the reason why the method of cases is central to philosophy.⁴² If, however, your goal is no longer to describe the concepts we have but to improve them—to think of how our concepts should be, then it is much less clear that asking questions of the form ‘Is this a *c*?’ about an imagined case can serve our purposes. There is a significant worry that this method, at best, reveals something about the concept *c* we as a matter of fact have, but our goal now is to think of what the concept *should* be.

Role of conceptual history/genealogy: Some people (notably Nietzsche in the passage quoted at the beginning of this Introduction) support claims in conceptual ethics by appeal to facts about the history/genealogy of concepts. Or, put in a way that will be more accurate on some theories of concepts: they appeal to historical facts about our (or other people’s) engagement and use of those concepts. On the one hand, there is an obvious worry here that appeals to conceptual history/genealogy might fall prey to versions of the genetic fallacy. On the other hand, there are cases where such historical facts seem at least *prima* relevant to our assessment of our current conceptual practices—for example, if we had acquired our concepts by being brainwashed by an evil scientist, that fact should presumably play a role in our assessment and improvement of those concepts. What role should conceptual history/genealogy have in conceptual ethics?⁴³

Let’s take stock of where we are. In this section, we have presented six clusters of domain-independent issues in conceptual engineering. The issues interact and overlap in various interesting ways and should, we hope, make clear that theorizing about conceptual engineering is fertile ground for philosophical exploration. Many of these issues have not yet been systematically explored. There are as of yet few efforts to give unified theories of conceptual engineering. It should also be clear from the outline above that conceptual engineering interacts in various intriguing ways with topics in ethics, philosophy of language, philosophy of mind, philosophical methodology, metaphysics, philosophy of mind, and epistemology. It also interacts with issues in

⁴² For a criticism of this way to thinking about the method of cases, see Cappelen (2012) and Deutsch (2015).

⁴³ For discussion of this issue, see Plunkett (2016).

linguistics, cognitive science, psychology, history, and sociology. Conceptual engineering has implications for those fields, but the theory of it will also draw on results from those fields. Towards the end of this Introduction we return to the issue of how one can see the position of conceptual engineering in philosophy overall.

4. Interaction between Specific Cases of Conceptual Engineering and General Theorizing

The theoretical foundations of conceptual engineering can be interesting in its own right, for much the same reasons any topic in philosophy can be. But many of those interested in conceptual engineering (and many conceptual activists in particular) are primarily motivated by an interest *not* in the general theoretical questions of the kind we just raised above for their own sake, but rather by a concern with specific concepts or words. For example: the concepts PERSON, FREEDOM, TRUTH, race and gender concepts, or concepts used for classifying mental illness in psychology. And in most cases they will also be interested in the lexical items used to express these concepts. Much of the contemporary discussion in conceptual engineering has been driven by concern with specific concepts or words.

In what follows we say a bit about some of the core issues we see involved in the *interaction* between the general and the specific here: how general theoretical issues (of the sort we canvassed in the last section) interact with more “applied” parts of conceptual engineering, focused on a specific concept or set of concepts.

Here is an analogy: A question that arises in *many* kinds of normative theorizing—including in ethics, political philosophy, and epistemology—is in what way (and to what extent) progress on specific cases (or more “applied” issues) is tethered to more general theory. For example: if we want to make progress on issues about debates in moral philosophy about climate change or abortion, how much (and when) should we appeal to a general normative ethical theory (e.g., act-utilitarianism)? What’s the best way to proceed with this? Is there a general, informative theory about this methodological question? Or can we only answer it when we have a specific set of ethical questions on the table, in particular social-historical contexts? Parallel questions arise here in the context of thinking about conceptual ethics.

In a connected vein, we can also ask about the extent to which normative theorizing about a domain (whether about systematic/general issues, or more applied issues) is tied to *meta*-level theorizing about the domain. For example: we can ask about the extent to which work in normative ethics (e.g., about whether act-utilitarianism is correct) and applied ethics (e.g., about abortion) should be informed by work in metaethics. Even framing this question (as well as our questions in the previous paragraph) inevitably raises thorny issues about how (if at all) one should distinguish between these different topics, or projects.

Our goal is not to settle these debates here. Rather it is to flag them, and note their importance for work on conceptual engineering. We also want to make a few general remarks about some important points that should be kept in mind when thinking about these issues.

First, the basic issues on the table here aren't idiosyncratic to specifically normative domains. We can wonder—for example—about how much our theorizing about a particular topic in linguistics or biology should be informed by our more general theorizing in those domains, as well as our theorizing about our theorizing in those domains.

Second, it should be clear that in order to engage in conceptual engineering you do not need to have worked-out, explicit views on all these issues mentioned above. If that was required, conceptual engineering would never have happened. The questions here are about the interaction between more general theoretical reflection and more applied issues in conceptual engineering, rather than the issue of whether the former is a prerequisite for the latter.

Third, we have many good examples of someone being good at X without being a good theorist about X (or of having a theory about what being-good-at-X consists in, or possessing views about the best methodology for engaging in X, etc.). For example, many scientists make massively important contributions to science while having bad views in the philosophy of science. And many good tennis players don't have good theories of their own activity. Moreover, there might well be costs to theoretical reflection as well. Perhaps the time it takes to engage in that reflection could have been spent better doing something else. Or, more dramatically, perhaps theoretical reflection will make someone worse at what she does; think of a tennis player who can't serve as well after thinking too much about her serving technique. None of this means that theoretical reflection can't aid people in many cases. But it does suggest we should proceed with caution in assuming theoretical reflection here will be crucial to success.

Fourth, even if the thoughts in the last paragraph suggest some amount of modesty and caution here, we don't want to be overly pessimistic about the contributions of general theory (or meta-theory) about a domain to more "applied" or more specific issues. Such theoretical reflection can, and we think often does, help make contributions. This is especially so when our theoretical reflection is in relatively good epistemic standing, compared to the standing of our theorizing about the more "applied" issues.

Fifth, many find the following view attractive here: the interaction (in terms of evidential import, methodology, etc.) between more applied issues and more general issues in conceptual engineering will go in *both* directions. If some version of that idea is on the right track, then the following becomes important: there will be many interesting questions about *how* the general informs the specific and vice versa. This will, we predict, pattern (at least to a certain degree) in sync with other theoretical domains. General theories tend to take the form of models that abstract from the messiness of particular cases and that is in part what makes systematic theorizing possible. On the other hand, such models will then include idealizing assumptions that often (or sometimes) will make it hard to see how to apply it to particular cases.

Finally, it is worth highlighting that some concepts are tied up with the general theory in a particularly direct way: As we mentioned above, the efforts to conceptually engineer concepts such as CONCEPT, TRUTH, OUGHT, or CONCEPTUAL ENGINEERING will have a direct and immediate impact on the general theory. These are points where the general theory of conceptual engineering and the engineering of specific concepts will be deeply connected.

5. Role and Scope of Conceptual Engineering in Philosophy: Descriptive

One sense we sometimes get when talking to people about conceptual engineering and conceptual ethics are that they are *'hot' new topics*—a *trendy new field*. While we hope it is true that conceptual engineering and conceptual ethics are things that many philosophers will work on and think about, it would be misleading in the extreme if we gave the impression that these are topics/activities that haven't been important throughout the history of philosophy. Many philosophers, working in many different theoretical traditions, across many centuries, have thought of their work as involving some kind of conceptual engineering or conceptual ethics, and/or conceived of the work of other philosophers along such lines (even if they didn't use the terminology we use here).

For example, consider the founding work of analytic philosophy in the early twentieth century. A case can be made that much of this work centrally involved conceptual engineering. For example, Frege's *Begriffsschrift* is a paradigm of conceptual engineering: he aimed to improve language for certain purposes. As he puts it: "If the task of philosophy is to break the domination of words over the human mind . . . then my concept notation, being developed for these purposes, can be a useful instrument for philosophers".⁴⁴ Or take Wittgenstein. In *Tractatus Logico-Philosophicus*, he aimed to draw a line between what could be said and what could only be shown. You *shouldn't*, according to Wittgenstein, try to say what can only be shown. The aim of telling philosophers (and others) about the legitimate and illegitimate uses of language is a *normative* aim.

Next, consider Carnap. His work on explication and language choice are paradigms of conceptual engineering. He writes:

The task of making more exact a vague or not quite exact concept used in everyday life or in an earlier stage of scientific or logical development, or rather of replacing it by a newly constructed, more exact concept, belongs among the most important tasks of logical analysis and logical construction. We call this the task of explicating, or of giving an explication for, the earlier concept. (Carnap, 1947/1956: 8–9)

Carnap's interest in conceptual engineering expands beyond the idea of the explication of concepts. For example, his criticism of metaphysics as lacking a semantic foundation (as being nonsensical), and his proposal for an improved language (inspired by his verificationism), are also deeply bound up with issues in conceptual engineering. Importantly, much of the conceptual engineering in early twentieth-century analytic philosophy wasn't just concerned about purely epistemological or scientific goals. For example, consider Carnap's aim of modifying language to allow multiple people, from multiple places, to engage in collective, rational inquiry. Making that possible was in part a *political* aim, tied to a democratic, enlightenment view of politics that ran through Carnap's work.⁴⁵ This political side of things is also pronounced in Susan Stebbing's worries about how certain ways of using

⁴⁴ Frege (1879/1967: 7).

⁴⁵ For connected discussion, see Galison (1990).

key terminology in politics (e.g., ‘democracy’ or ‘freedom’) hindered clear thinking about social and political issues, which in turn made it difficult to effectively critique the rise of fascism. Stebbing saw analytic philosophy as helping provide tools to combat the relevant problematic sorts of thinking; both in diagnosing what was going wrong with it and in helping us make it better.⁴⁶

As this brings out, Frege, Wittgenstein, Carnap, Stebbing, and other founders of analytic philosophy were extensively engaged in conceptual engineering. So rather than describe conceptual engineering as a ‘hot’ new topic in analytic philosophy, we could instead think of it as simply paying more attention to a key aspect of analytic philosophy that has been with us since its origins.

Moreover, the idea that key parts of philosophy involve conceptual engineering is hardly parochial to self-consciously “analytic” philosophy. Consider here the quote from Nietzsche we introduced at the start of this chapter, in which he claims that “what is needed above all is an absolute skepticism toward all inherited concepts” and that philosophers “must no longer accept concepts as a gift, nor merely purify and polish them, but first make and create them, present them and make them convincing”.⁴⁷ Nietzsche is here advocating a radical skeptical stance with respect to inherited concepts, as well chastising other philosophers for failing to (at the very least) seriously engage that position.⁴⁸ In Nietzsche’s view, it seems, philosophy should involve more conceptual engineering than it in fact has. This attitude is, in turn, reflected in his own work. For instance, a good part of *On the Genealogy of Morality* can be read as a critique of the distinctively *moral* concepts that have shaped much of modern life.⁴⁹ To put it in Nietzsche’s own terms (from *Twilight of The Idols*), we have become “stuck in a cage, imprisoned among all sorts of terrible concepts”, and part of his goal is to help us (or at least *some* of us) find a way out of that cage.⁵⁰

In the passage from *The Will to Power*, Nietzsche (in characteristic fashion) positions his views as a radical break from much of the history of philosophy. But while Nietzsche might well be advocating for a more radical view of conceptual engineering than many have, there is a strong case that some amount of conceptual engineering, involving some amount of skepticism toward our inherited concepts, has played an important role in philosophy throughout its history. On this front, consider here what Strawson says in the introduction to *Individuals* about the difference between revisionary and descriptive metaphysics. He writes that “descriptive metaphysics is content to describe the actual structure of our thought about the world, revisionary metaphysics is concerned to produce a better structure”.⁵¹

⁴⁶ See Stebbing (1939/1941) and Stebbing (1941/1948). Thanks to Bryan Pickel for helpful discussion of these parts of Stebbing’s work.

⁴⁷ Nietzsche (1968: 220–1, section 409).

⁴⁸ Note that Nietzsche, later on in the passage we quoted, gives a nod to Plato for (at least possibly) seriously engaging the kind of radical skepticism about our current concepts that Nietzsche advocates. Plato’s defense of our concepts, thinks Nietzsche, *perhaps* results from him taking the skeptical challenge seriously. He writes: “What is needed above all is an absolute skepticism toward all inherited concepts (of the kind that one philosopher *perhaps* possessed—Plato, of course—for he taught the reverse)” (Nietzsche 1968: 221, section 409).

⁴⁹ Nietzsche (1887/1994).

⁵⁰ Nietzsche (1889/1954: 502).

⁵¹ Strawson (1959: 9).

Strawson describes the revisionist as insisting that metaphysics is “essentially an instrument of conceptual change, a means of furthering or registering new directions or styles of thought”.⁵² Importantly, Strawson gestures at a way of reading the history of philosophy as a division between revisionists and descriptivists. He says: “Perhaps no actual metaphysician has ever been, both in intention and effect, wholly the one thing or the other. But we can distinguish broadly: Descartes, Leibniz, Berkeley are revisionary, Aristotle and Kant descriptive. Hume, the ironist of philosophy, is more difficult to place. He appears now under one aspect, now under another”.⁵³ If Strawson is correct, then conceptual engineering has been a key part of the history of philosophy stretching back centuries.

Moreover, the picture we get from Strawson above might well be understating the import of conceptual engineering to philosophy. If certain views in the philosophy of language are right, then philosophers might be *tacitly* engaging in conceptual engineering much more than they realize. For example, consider Peter Ludlow’s recent work on what he calls the “dynamic lexicon”. His idea, in rough outline, is that speakers regularly adjust and create new meanings for words on the fly in conversation, such that those meanings are “dynamic” and in flux. If Ludlow is correct, then much of conversation (including much of conversation *in philosophy*) involves navigating issues in conceptual engineering; including issues about what the best meaning of a word is for the context at hand, or going forward into further contexts.⁵⁴ A similar view of philosophy—in which it *in fact* involves extensive conceptual engineering (perhaps without the awareness of philosophers that it does so)—emerges in the work of other contemporary philosophers, including, for example, Amie Thomasson and Kevin Scharp.⁵⁵

What we have just said is of course not even the beginning of a sketch of the role of conceptual engineering and conceptual ethics in the history of philosophy. There is rich terrain here to explore in many places, including, for example, its role in the Platonic dialogues, Hume, Kant, pragmatism, Heidegger, Foucault, Deleuze, Tarski, and ordinary language philosophy. We predict that there are rich new perspectives available for readings of the history of philosophy that more actively pay attention to the idea of conceptual engineering. We hope that scholars with better historical knowledge than us will write histories of philosophy from this perspective.

Of course, it might turn out that that philosophy has *in fact* involved extensive conceptual engineering, but that it is a deep mistake for it to continue in this vein. Or perhaps the reverse is true: perhaps philosophy has *not* in fact involved much conceptual engineering and should involve much more. This raises further normative issues. These issues are not about the descriptive question about the role that conceptual engineering (or conceptual ethics) has played in the history of philosophy (or currently plays in contemporary philosophy). Rather, they are about the question of what role conceptual engineering (or conceptual ethics) *should* play. We turn to this set of questions in the next section.

⁵² Strawson (1959: 10).

⁵³ Strawson (1959: 9).

⁵⁴ Ludlow (2014).

⁵⁵ See Thomasson (2016) and Kevin Scharp’s chapter “Philosophy as the Study of Defective Concepts”, (Chapter 19, this volume). See also Plunkett and Sundell (2013) and Plunkett (2015) for sympathetic discussion of this possibility.

6. Role and Scope of Conceptual Engineering in Philosophy: Normative

So what role *should* conceptual engineering and conceptual ethics play in philosophy? For ease of exposition in what follows, let's take this question in terms of conceptual engineering. Here's a way to break up that normative question into two components:

1. How many parts of (or sub-fields of, or issues in) philosophy should conceptual engineering play a role in?
2. For each part it should be involved in, how important should it be?

Focusing on 1 and 2 makes the answer to the question 'What role should conceptual engineering play in philosophy?' a matter of degree. Below we briefly sketch possible motivations for four kinds of answers:

1. All of All
2. All of Some
3. Some of All
4. Some of Some
5. Nothing

1. *All of All of Philosophy*

On this view, conceptual engineering should be seen as relevant to every issue in philosophy and it's the only thing that's relevant. So philosophy should consist entirely of conceptual engineering.

Kevin Scharp endorses All of All of philosophy. He argues that all philosophical concepts are defective, and that philosophy's task should be to discover those defects and then create replacement concepts. Once that is done, we ship the questions off to the sciences, and that's the end of philosophy.⁵⁶

2. *All of Some*

This is the view you would hold if you think, for example, Scharp's view is correct for some parts of philosophy: in those parts, conceptual engineering is all there is to do. Then there are some other parts of philosophy where there's more to do than conceptual engineering.

3. *Some of All of Philosophy*

According to this view, conceptual engineering is relevant to all philosophical questions, but answering those questions requires more than doing conceptual engineering (so it is relevant to all of philosophy, but isn't all of philosophy).

One important central argument for Some of All is that we have no particular good reason to think that the concepts that we have inherited are *ideal* for philosophical theorizing. The default view should be that they could be improved. That thought is natural even for those who think we have good reason to preserve many of the

⁵⁶ See Scharp's "Philosophy as the Study of Defective Concepts" (Chapter 19, this volume).

distinctions found in natural language. For example, in “A Plea for Excuses”, Austin says that

ordinary language . . . embodies . . . the inherited experience and acumen of many generations of men. . . . If a distinction works well for practical purposes in ordinary life (no mean feat, for even ordinary life is full of hard cases), then there is sure to be something in it, it will not mark nothing.⁵⁷

However, in a point that is often neglected by those who quote the above passage, Austin then goes on to note that “ordinary language is not the last word: in principle it can everywhere be supplemented and improved upon and superseded”.⁵⁸ The challenge here is to recognize when ordinary language is good enough and when it can be improved upon. Philosophy might play a helpful role in that.

One view one could hold in support of Some of All is that the default assumption should be that the ordinary distinctions were not made for or developed to be ideal for philosophical theorizing. As a corollary, part of what philosophers should do is reflect critically on the usefulness of ordinary concepts for philosophical theorizing. A stronger view would hold that the concepts we currently have embody ideologies and power structures that can be repressive both on a political and personal level, so we should always critically examine inherited concepts. On those kinds of views, conceptual engineering isn’t All of All of philosophy, but is a part of all of philosophy.

4. *Some of Some of Philosophy*

On this view conceptual engineering should be seen as relevant to some aspects of some philosophical questions. This is simply a more restricted version of Some of All. If you think the arguments we just sketched for Some of All don’t apply to all philosophical concepts, but to some, this would be your favoured normative view of conceptual engineering.

One motivation for excluding conceptual engineering from some domains has been mentioned above: maybe some concepts are so basic that they are irreplaceable and immutable. Maybe TRUTH is like that. Maybe basic normative concepts, like OUGHT, are like that. If so, these are concepts you can use to engage in conceptual engineering, but are not themselves engineerable. If so, conceptual engineering is relevant in some, but not all philosophical domains.

5. *Nothing*

According to Nothing, conceptual engineering plays no part in philosophy *at all*. This is a view the editors of this volume bet against, but some contributors to this volume (e.g., Ball, Chapter 2, and Greenough, Chapter 11) flirt with or endorse it. Here are three possible motivations for Nothing:

⁵⁷ Austin (1956: 11).

⁵⁸ Austin (1956: 11).

- (i) Conceptual engineering requires the existence of concepts, there are no concepts, so we can't do conceptual engineering. Since ought implies can, we shouldn't do conceptual engineering.⁵⁹
- (ii) Conceptual engineering involves fiddling with concepts, but fiddling with concepts is impossible.⁶⁰ Maybe that is because they are abstract entities (or entities of some other kind you can't fiddle with). Since ought implies can, we shouldn't do conceptual engineering.
- (iii) It simply isn't within the proper domain of philosophy to fiddle with concepts. Philosophy is about reality: knowledge, freedom, meaning, belief, etc. Those phenomena are within the proper domain of philosophy, but the concepts are not.⁶¹

Most of the chapters in this volume are opposed to Nothing. This is natural since the aim, in part, is to provide readers with a broad range of frameworks for theorizing about conceptual engineering. In our view, it is only when many of those frameworks are on the table that proponents of Nothing will have a clear enough target.

7. Conceptual Engineering beyond Philosophy

Almost everything we have said about the potential significance of conceptual engineering (and conceptual ethics) apply beyond philosophy to inquiry more generally (as well as to speech and thought more generally).

Many of the arguments for the importance of conceptual engineering in philosophy are also arguments for its importance in other fields, such as biology, mathematics, physics, psychiatry, law, and politics. We use concepts and words in all areas of inquiry, and they can be better or worse relative to the goals and standards of a given part of inquiry. Thus, conceptual engineering matters for all domains of inquiry.

The point we just made applies beyond inquiry: many activities involve concepts and lexical items in various ways. The activity of cooking a dinner depends in part on the chefs' conceptual repertoire. So does hiking, going to war, and making friends. Some activities, like breathing and blinking one's eyes, might be independent of our concepts, but the range and significance of concept-involving activities is very broad. If that line of thought is correct, conceptual engineering is important not just for linguistic and cognitive activities, but also for many other core elements of human life.

These facts about the import of conceptual engineering beyond philosophy help underscore the potential significance of philosophical work done on conceptual engineering (and conceptual ethics). At the same time, it helps open up potentially fruitful avenues for future research. For once we see the import of conceptual engineering to other areas of inquiry—and to a wide range of activities that people engage in—we can look to see how conceptual engineering works in these other areas

⁵⁹ Note that this presupposes that conceptual engineering requires the existence of concepts in the relevant sense, which one of us (Cappelen) rejects. See Cappelen (2018).

⁶⁰ We here use 'fiddling' to mean roughly the following: do something with or to.

⁶¹ There is, for example, a way of reading Williamson (2007) as advocating a version of this view. In fact, we think his view is more subtle than this, but we won't go into the details of it here.

as input to our philosophical theorizing about the topic. Such investigation might help us not only better understand how conceptual engineering in fact works, but also provide us with helpful material to think about in our theorizing about how it should work.

Conclusion

So far we have tried to remain fairly neutral in our presentation. In conclusion, it's time to put some of our cards on the table: we think conceptual engineering and conceptual ethics both *are* and *should be* central to philosophy. The role of conceptual engineering and conceptual ethics in philosophy and other areas of inquiry has been underexplored, often overlooked, and typically underappreciated. Philosophers have engaged in conceptual engineering and conceptual ethics in various sub-fields in various time periods, and there has been some scattered theorizing about them. But, especially in comparison to a range of descriptive questions about concepts and words—including issues about what they are, and how we use them—there has been relatively little sustained engagement with the normative and evaluative questions about concepts and words at the heart of conceptual engineering and conceptual ethics. We don't want to hang too much on the question of just how central conceptual engineering and conceptual ethics are to philosophy, but whatever degree of centrality one assigns to them, we hope this Introduction and the contributions in the volume helps highlight a broad range of interesting, important, and under-explored questions.

Acknowledgments

This chapter has benefited enormously from many discussions with other philosophers. We owe a particular debt to Alexis Burgess, our co-editor of this volume. We would also especially like to thank Derek Ball, Josh Dever, Matti Eklund, Øystein Linnebo, Tristram McPherson, Tim Sundell, Amie Thomasson, David Braddon-Mitchell, Peter Galison, Sally Haslanger, Peter Railton, Rachel Sterken, and Kevin Scharp. Note that the views in this Introduction are not the considered views of either Plunkett or of Cappelen. Their own considered views can be found in Plunkett (2015) and Cappelen (2018). The aim of this Introduction is to provide an easily accessible entry point for those not already invested in the topic.

References

- Anderson, Elizabeth. 2001. Unstrapping the Straitjacket of 'Preference': A Comment on Amartya Sen's Contributions to Philosophy and Economics. *Economics and Philosophy* 17 (1):21–38.
- Austin, John. 1956. A Plea for Excuses. *Proceedings of the Aristotelian Society* 57:1–30.
- Ball, Derek. Chapter 2, this volume. Revisionary Analysis without Meaning Change (Or, Could Women Be Analytically Oppressed)? In Alexis Burgess, Herman Cappelen, and David Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Blackburn, Simon. 2001. *Think: A Compelling Introduction to Philosophy*. Oxford: Oxford University Press.

- Boghossian, P. 1989. The Rule-Following Considerations. *Mind*, 98, 507–49.
- Brandom, Robert B. 2001. Modality, Normativity, and Intentionality. *Philosophy and Phenomenological Research* 63 (3):611–23.
- Burgess, Alexis. Never Say “Never say ‘Never’”? Chapter 6, this volume. In Alexis Burgess, Herman Cappelen, and David Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- 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, Alexis, and Sherman, Brett. 2014. *Metasemantics: New Essays on the Foundations of Meaning*. Oxford: Oxford University Press.
- Cappelen, Herman. 2012. *Philosophy without Intuitions*. Oxford: Oxford University Press.
- Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Carnap, Rudolf. 1947/1956. *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago: University of Chicago Press.
- Chalmers, David J. 2011. Verbal Disputes. *Philosophical Review* 120 (4):515–66.
- Deutsch, Max. 2015. *The Myth of the Intuitive*. Cambridge, MA: MIT Press.
- Eklund, Matti. 2002. Inconsistent Languages. *Philosophy and Phenomenological Research* 64 (2):251–75.
- Eklund, Matti. 2015. Intuitions, Conceptual Engineering, and Conceptual Fixed Points. *The Palgrave Handbook of Philosophical Methods*, ed. C. Daly. London: Palgrave.
- Eklund, Matti. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.
- Frege, Gottlob. 1879/1967. *Begriffsschrift: Eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Trans. S. B.-M. in Jean van Heijenoort (ed.), *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931*. Cambridge, MA: Harvard University Press.
- Galison, Peter. 1990. Aufbau/Bauhaus: Logical Positivism and Architectural Modernism. *Critical Inquiry* 16 (4):709–52.
- Gibbard, Allan. 2012. *Meaning and Normativity*. Oxford: Oxford University Press.
- Hacking, Ian. 2002. Making Up People. *Historical Ontology*. Cambridge, MA: Harvard University Press.
- Haslanger, Sally. 2000. Gender and Race: (What) Are They? (What) Do We Want Them to Be? *Nous* 34 (1):31–55.
- Haslanger, Sally. 2012. *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.
- Hattiangadi, Anandi. 2017. The Normativity of Meaning. In B. Hale, C. Wright and A. Miller (eds.), *A Companion to the Philosophy of Language*. Oxford: Wiley Blackwell.
- Jenkins, C. S. I. 2014. Merely Verbal Disputes. *Erkenntnis* 79 (1):11–30.
- Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.
- Kuhn, Thomas. 1962/2012. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Lau, Joe, and Deutsch, Max. 2016. Externalism about Mental Content. *The Stanford Encyclopedia of Philosophy* (winter edn), ed. Edward N. Zalta. <<https://plato.stanford.edu/archives/win2016/entries/content-externalism/>>.
- Ludlow, Peter. 2014. *Living Words: Meaning Undetermination and the Dynamic Lexicon*. Oxford: Oxford University Press.
- Margolis, Eric, and Laurence, Stephen. 1999. *Concepts: Core Readings*. Cambridge, MA: MIT Press.

- Nietzsche, Friedrich Wilhelm. 1887/1994. *On the Genealogy of Morality*. New York: Cambridge University Press.
- Nietzsche, Friedrich Wilhelm. 1889/1954. Twilight of the Idols. In W. Kaufmann (ed.), *The Portable Nietzsche*. New York: Viking.
- Nietzsche, Friedrich Wilhelm. 1901/1968. *The Will to Power*. Trans. W. Kaufmann. New York: Random House.
- Plunkett, David. 2015. Which Concepts Should We Use? Metalinguistic Negotiations and the Methodology of Philosophy. *Inquiry* 58 (7–8):828–74.
- Plunkett, David. 2016. Conceptual History, Conceptual Ethics, and the Aims of Inquiry: A Framework for Thinking about the Relevance of the History/Genealogy of Concepts to Normative Inquiry. *Ergo* 3 (2):27–64.
- Plunkett, David, and Sundell, Timothy. 2013. Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint* 13 (23):1–37.
- Quine, Willard V.O. 1951. Two Dogmas of Empiricism. *Philosophical Review* 60 (1):20–43.
- Railton, Peter. 1986a. Facts and Values. *Philosophical Topics* 14:5–31.
- Railton, Peter. 1986b. Moral Realism. *The Philosophical Review* 95:163–207.
- Railton, Peter. 2003. *Facts, Values, and Norms: Essays toward a Morality of Consequence*. New York: Cambridge University Press.
- Scharp, Kevin. 2013. *Replacing Truth*. Oxford: Oxford University Press.
- Scharp, Kevin. Chapter 19, this volume. Philosophy as the Study of Defective Concepts. In Alexis Burgess, Herman Cappelen, and David Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Schroeter, Laura, and Schroeter, Francois. 2014. Normative Concepts: A Connectedness Model. *Philosophers' Imprint* 14 (25):1–26.
- Sider, Theodore. 2012. *Writing the Book of the World*. Oxford: Oxford University Press.
- Stebbing, L. Susan. 1939/1941. *Thinking to Some Purpose*. New York: Allen Lane: Penguin Books.
- Stebbing, L. Susan. 1941/1948. *Ideals and Illusions*. London: Watts and Co.
- Sterken, Rachel Katharine. Chapter 20, this volume. Linguistic Intervention and Transformative Communicative Disruptions. In Alexis Burgess, Herman Cappelen, and David Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Strawson, P. F. 1959. *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.
- Strawson, P. F. 1963. Carnap's Views on Conceptual Systems versus Natural Languages in Analytic Philosophy. In P. A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*. London: Open Court.
- Thomasson, Amie L. 2016. Metaphysical Disputes and Metalinguistic Negotiation. *Analytic Philosophy* 57 (4) :1–28.
- Williamson, Timothy. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.

Abstracts of Chapters

1 Introduction

In this Introduction, we aim to introduce the reader to the basic topic of this book. As part of this, we explain why we are using two different expressions ('conceptual engineering' and 'conceptual ethics') to describe the topics in the book. We then turn to some of the central foundational issues that arise for conceptual engineering and conceptual ethics, and finally we outline various views one might have about their role in philosophy and inquiry more generally.

2 Revisionary Analysis without Meaning Change (Or, Could Women Be Analytically Oppressed?)

This chapter develops a conception of philosophical analysis which makes sense of the idea that a correct analysis can be revisionary (in that it departs from ordinary or expert belief and linguistic usage). The view is superior to the alternatives defended by most proponents of 'conceptual ethics' and 'conceptual engineering' (according to which revisionary theorizing involves replacing words or concepts) because it better explains the arguments we advance when we engage with proposed revisionary analyses. A key idea is that analytic claims can emerge in the course of debate without change of meaning, so that our acceptance (perhaps late in the debate) of some analyticity can fix the meaning of a word as we used it all along. The discussion focuses on Haslanger's revisionary analysis of gender.

3 Minimal Substantivity

This chapter defends a notion of "minimal substantivity" for ontological disputes, focusing on persistence and composition as case-studies. Deflationists argue that these disputes are defective: they are merely verbal, or epistemically underdetermined. Anti-Deflationists emphasize the substantivity of these controversies, either in Quinean terms, or in terms of fundamentality. This chapter aims at striking a midway between these approaches, thus outlining a Minimal Anti-Deflationism. First, it is shown that even if these ontological disputes are merely verbal, they can be recast as non-verbal metalinguistic disputes. These exchanges have in turn an ontological substantivity if we understand them as disputes about which ontological commitments we should undertake. This substantivity is however minimal, because it need

not be articulated either naturalistically or in terms of fundamentality. In closing, an illustration is provided of the fruitfulness of the minimal substantivity notion for cases of conceptual engineering.

4 Reactive Concepts

This chapter recommends that we consider a kind of concept which bears a relation like the one traditional accounts of concepts bear to beliefs, but instead bears it to states individuated not only by their causal inputs, but also by their direct causal outputs. They will be called reactive representations, RRs for short. They are partially representational states which are reactive inasmuch as they bypass interaction with distinct desires to directly motivate behaviour. Associated with these representations are abilities that will be called Reactive Concepts. The chapter argues that taxonomizing mental states this way casts light on the nature of a range of phenomena, including hate speech, crypto-evaluative terms, and phenomenal concepts.

5 Strategic Conceptual Engineering for Epistemic and Social Aims

This chapter advocates strategic conceptual engineering, that is, the employment of a (possibly novel) concept for specific epistemic or social aims, concomitant with the openness to use a different concept for other contexts. We illustrate this approach by sketching three distinct concepts of gender and arguing that all of them are needed, as they answer to different social aims. The first concept serves the aim of identifying and explaining gender-based discrimination. It is similar to Haslanger's account, except that rather than offering a definition of 'woman' we focus on 'gender' as one among several axes of discrimination. The second concept of gender is to assign legal rights and social recognitions, and is to be trans-inclusive. We argue that this cannot be achieved by previously suggested concepts (including Jenkins's) that include substantial gender-related psychological features or awareness of social expectations. The third concept of gender serves the aim of personal empowerment through gender identity. This chapter points to contexts where a concept's role in explanation and moral reasoning can be more important than determining the extensions of concepts.

6 Never Say 'Never Say "Never"'

Is there anything wrong with using a concept *C* in the course of arguing against the use of *C*? You might think not, so long as the argumentation is framed as a kind of *reductio*; or climbing a ladder to be kicked away. On the other hand, you might find such "hypocritical" arguments self-undermining: if they're sound, then we shouldn't accept any of their premises that make use of *C*. Understanding the status of hypocrisy is an urgent question for conceptual ethics (and therefore engineering),

insofar as concepts of philosophical interest tend to be central to our conceptual schemes, in the sense that they tend to turn up in the analyses of lots of other concepts. Developing these points (of view), the present chapter exposes a worrisome feature of the anti-hypocrisy argument just sketched: it seems to be indeterminate whether or not the argument is itself hypocritical.

7 Conceptual Engineering

This chapter develops and defends the Master Argument for Conceptual Engineering: (1) If *W* is a word that has a meaning *M*, then there are many similar meanings, *M*₁, *M*₂, ..., *M*_{*n*}, *W* could have. (2) We have no good reason to think that the meaning that *W* ended up with is the best meaning *W* could have: there will typically be indefinitely many alternative meanings that would be better meanings for *W*. (3) When we speak, think, and theorize it's important to make sure our words have as good meanings as possible. (4) As a corollary: when doing philosophy, we should try to find good meanings for core philosophical terms and they will typically not be the meanings those words as a matter of fact have. (5) So no matter what topic a philosopher is concerned with, she should assess and ameliorate the meanings of central terms. I respond to seven objections to this argument.

8 Preliminary Scouting Reports from the Outer Limits of Conceptual Engineering

If we distinguish between the conceptual engineering task of determining what concepts can be built and the conceptual ethics task of determining what concepts ought to be used once built, we make clear the possibility of a conceptual maximalist position that avoids conceptual ethics by holding that the norms of theorizing require a big theory including all truths expressible using any concepts. But the conceptual maximalist then assumes the burden of saying what the range of possible concepts (equivalently, possible languages) is. Despite the attractions of conceptual maximalism, we lack even the beginnings of an answer to the question of what concepts there could be.

9 Descriptive vs. Ameliorative Projects: The Role of Normative Considerations

Feminist philosophers have famously argued that contextual factors play a role in the justification and assessment of scientific theories. A similar question might arise regarding inquiries in metaphysics. Contemporary metaphysicians often assume that contextual factors do not belong in metaphysics. The main aim of this chapter is to argue that normative considerations such as moral and political considerations are relevant in metaphysics. In particular, this chapter explores the nature of descriptive

projects seeking to reveal the operative concept associated with a term (as opposed to ameliorative projects seeking to reveal the target concept that we should associate with the term), and argues that moral and political considerations are relevant not only with respect to ameliorative projects but also with respect to descriptive projects.

10 Variance Theses in Ontology and Metaethics

The chapter illustrates conceptual engineering by bringing up a number of issues in metaontology and metaethics. A prominent debate in metaontology relates to whether some existence concept is metaphysically privileged. On the one hand, ontological realists say yes, and, on the other hand, friends of quantifier variance say no. The chapter brings up the corresponding question in metaethics by asking, is some rightness concept normatively privileged? It investigates this question, and compares the metaethics case and the metaontology case. One aim is to arrive at conclusions regarding possible limits to the project of conceptual engineering.

11 Neutralism and Conceptual Engineering

Conceptual engineers often invoke a distinction between happy-face and unhappy-face solutions to alethic paradoxes. Happy-face solutions are thoroughly specific: they isolate a single, basic principle (“the culprit”). Unhappy-face solutions, meanwhile, are thoroughly non-specific: they merely establish the collective guilt of a group of principles which together produce the paradox. According to this taxonomy, conceptual engineering can only take place via unhappy-face solutions. In this chapter, I: (1) give an expanded taxonomy which allows for both Happy-Face and two forms of Unhappy-Face Conceptual Engineering and show that (2) happy-face treatments represent a limit case. (3) Unhappy-face treatments also represent a kind of limit case. (4) Between these limit cases are treatments which are neither maximally specific nor maximally unspecific but nonetheless specific enough to treat a paradox. (5) Such treatments become thoroughly neutralist when they reject some principle at work in a paradox from a theory-neutral perspective. The upshot is Neutralism—the view that philosophical progress can take place when (and sometimes only when) a thoroughly neutral, non-specific theory is adopted.

12 Going On, Not in the Same Way

This chapter considers, within an externalist semantics, several ways we might understand the project of improving our concepts to promote greater justice. The tools that culture provides us with—such as language, concepts, and inferential patterns—provide frames for coordination and shape our interaction. There are multiple ways these tools can fail us, for example by the limited structure of options they make intelligible. However, we can sometimes reconfigure the resources so that our practical orientations are more responsive to what is good and coordinate in

ways that are just. Such reconfiguration often happens in law; it also occurs in social movements, counter-publics, subaltern communities, and in fascist propaganda. Contestation over meaning is not “mere semantics” for—together with political and material change—it can shape our agency and our lives together.

13 The Theory–Theory Approach to Ethics

One way to approach the theory of reference for proper names is by asking what proper names are good for in the sense of the valuable purposes they serve. Suppose we approach ethical terms and concepts in the same spirit, asking questions like: What purposes do they serve? How could we do something similar but do it better? This chapter explores the implications of this way of thinking about ethical terms and concepts, and explains why a theory–theory or moral functionalist account of them is so attractive when we approach matters from this perspective. The discussion is set inside an avowedly cognitivist, naturalist framework, and touches on the implications of this framework for how to adjudicate debates between rival views in ethics, and the relevance of evolutionary considerations.

14 Conceptual Ethics and the Methodology of Normative Inquiry

This chapter explores two central questions in the conceptual ethics of normative inquiry. The first is whether to orient one’s normative inquiry around folk normative concepts (like KNOWLEDGE OR IMMORAL) or around theoretical normative concepts (like ADEQUATE EPISTEMIC JUSTIFICATION OR PRO TANTO PRACTICAL REASON). The second is whether to orient one’s normative inquiry around concepts whose normative authority is especially accessible to us (such as OUGHT ALL THINGS CONSIDERED), or around concepts whose extension is especially accessible to us (such as BETRAYAL). The chapter aims to make vivid and plausible a range of possible answers to these questions, and important forms of argument that can be used to favor certain answers over others.

15 Conceptual Evaluation: Epistemic

On a view implicitly endorsed by many, a concept is epistemically better than another if and because it does a better job at ‘carving at the joints’, or if the property corresponding to it is ‘more natural’ than the one corresponding to another. This chapter offers an argument against this seemingly plausible thought, starting from three key observations about the way we use and evaluate concepts from an epistemic perspective: that we look for concepts that play a role in explanations of things that cry out for explanation; that we evaluate not only ‘empirical’ concepts, but also mathematical and perhaps moral concepts from an epistemic perspective; and that there is much more complexity to the concept/property relation than the natural

thought seems to presuppose. These observations, it is argued, rule out giving a theory of conceptual evaluation that is a corollary of a metaphysical ranking of the relevant properties.

16 Analyzing Concepts and Allocating Referents

The analysis of concepts is an important part of the philosophical exercise but has to be complemented by the allocation of a referent to any world-tracking concept: this will give an account of what constitutes the phenomenon tracked. But analysis generally leaves room for negotiation about the referent to be allocated, for two reasons. First, because it may underdetermine the specification of conditions that anything must satisfy if it is to count as a referent. And, second, because the specification chosen may itself underdetermine the property that constitutes the referent. The license that philosophy consequently enjoys has to be exercised on the basis of analytically independent criteria linked with the theory within which any realistic concept has to take its place. The lesson is illustrated schematically, it comments on concepts like those of causation, freewill, and value and in a case study of the concept of freedom. It connects with the line argued by a range of recent, more revisionary approaches to philosophy.

17 The A-project and the B-project

Ameliorative philosophical analysis elucidates a concept by looking at the purposes it has for its users. It evaluates them, and gives an analysis on the basis of the purposes the analyst thinks ought to control the concept's use. As such it can look wildly revisionary, an attempt not to tell us what we mean but to change the subject. This chapter sketches a view of meaning on which ameliorative analysis can be understood as not at all subject changing: a word's meaning is the evolving collection of presuppositions its users make, and expect to be recognized as made, when the word is used. Focusing on Sally Haslanger's ameliorative accounts of race and gender, this chapter argues that ameliorative analysis is best understood not as a project that is successful only if it results in the concept's reference reflecting the ameliorative analysis. Conceptual amelioration and engineering are attempts to foster a kind of evolution within a population, a change in the presuppositions that (it is common knowledge) accompany the concept's use. This can be achieved and have worthwhile effects even while no shift in reference occurs.

18 Talk and Thought

This chapter provides an externalist account of talk and thought that clearly distinguishes the two. It is argued that linguistic meanings and concepts track different phenomena and have different explanatory roles. The distinction, understood along the lines proposed, brings theoretical gains in a cluster of related areas. It provides an

account of meaning change which accommodates the phenomenon of contested meanings and the possibility of substantive disagreement across theoretical divides, and it explains the nature and value of conceptual engineering in a way that addresses recent prominent concerns.

19 Philosophy as the Study of Defective Concepts

From familiar concepts like tall and table to exotic ones like gravity and genocide, they guide our lives and are the basis for how we represent the world. However, there is good reason to think that many of our most cherished concepts, like truth, freedom, knowledge, and rationality, are defective in the sense that the rules for using them are inconsistent. This defect leads those who possess these concepts into paradoxes and absurdities. Indeed, I argue that many of the central problems of contemporary philosophy should be thought of as having their source in philosophical concepts that are defective in this way. If that is right, then we should take a more active role in crafting and sculpting our conceptual repertoire. We need to explore various ways of replacing these defective concepts with ones that will still do the work we need them to do without leading us into contradictions.

20 Linguistic Intervention and Transformative Communicative Disruptions

What words we use, and what meanings they have, is important. We shouldn't use slurs; we should use 'rape' to include spousal rape (for centuries we didn't); we should have a word which picks out the sexual harassment suffered by people in the workplace and elsewhere (for centuries we didn't). Sometimes we need to change the word-meaning pairs in circulation, either by getting rid of the pair completely (slurs), changing the meaning (as we did with 'rape'), or adding brand new word-meaning pairs (as with 'sexual harassment'). A problem, though, is how to do this. One might worry that any attempt to change language in this way will lead to widespread miscommunication and confusion. I argue that this is indeed so, but that's a feature, not a bug, of attempting to change word-meaning pairs. The miscommunications and confusion such changes cause can lead us, via a process I call transformative communicative disruption, to reflect on our language and its use, and this can further, rather than hinder, our goal of improving language.

21 A Pragmatic Method for Normative Conceptual Work

How ought we to do work in conceptual ethics? Some have thought that conceptual choice should itself be guided by (heavyweight) metaphysics—for we should be sure that our concepts pick out things that exist or should aim to choose concepts that

really ‘carve the world at its joints’. An alternative is to take a pragmatic approach to conceptual ethics. But pragmatic approaches are often criticized as unable to account for intuitions that some conceptual choices are objectively better than others, and intuitions that the world is structured. As a result, the fear is that a pragmatic approach leaves conceptual choices arbitrary and insusceptible to critique. This chapter confronts such worries and develops a pragmatic method for conceptual ethics that clearly avoids these problems. As a result, we need not rely on heavyweight metaphysics and become entangled in its epistemological mysteries to do conceptual ethics.