19

Philosophy as the Study of Defective Concepts

Kevin Scharp

For the past decade or so, I have tried to make sense of the liar paradox, which is a 2,300 year old problem associated with truth. It goes like this. Think of a sentence that says that that very sentence is not true. Is it true or not true? No matter what answer you give, it is easy to derive a contradiction in just a few steps. The view that I developed is that the liar and the other paradoxes to which truth gives rise are symptoms of an underlying defect in the concept itself. It's not that the reasoning in the paradox involves some trivial mistake or faulty assumption. It is that the concept of truth itself is to blame for these paradoxes that we have found ourselves in for a very long time now. We might say that when we reason to the contradiction in the paradox, we are using all our concepts according to rules that are built into those concepts.

The second half of project is to replace the concept of truth for certain purposes. Truth is a defective concept, and there are certain jobs and it's not very good for. We should replace it with a team of concepts that together can do its job without giving us any of the paradoxes or problems that plague the concept of truth. The job that I really focus on is explaining the meanings or contents of natural language sentences by way of natural language semantics. A very popular form attributes truth conditions to sentences of natural languages. The paradoxes that truth generates mean that it can't do that job very well at all. Anytime you try to use truth to give a semantics for a natural language like English, or really any expressively rich language at all, you end up contradicting yourself. You end up saying things that are inconsistent. And it's the paradoxes that force this upon us.

The replacement concepts, which I call ascending truth and descending truth, can do this job perfectly, and the resulting theory agrees with traditional semantics as a special case everywhere the latter provides coherent results. So it is a lot like advances in science, where the successor theory does everything that the earlier theory did, and solves some extra problems as well.[1]

I call the method followed in this project of replacing truth *conceptual engineering*. I take conceptual engineering to be actively changing some aspect of our concepts—eliminating bad ones, deciding which ones we should use, and which word should

[1] Scharp 2013.

Kevin Scharp, *Philosophy as the Study of Defective Concepts* In: *Conceptual Engineering and Conceptual Ethics*. Edited by: Alexis Burgess, Herman Cappelen, and David Plunkett, Oxford University Press (2020). © Kevin Scharp.
DOI: 10.1093/oso/9780198801856.003.0019

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

express them. Although there are plenty of instances of conceptual engineering in the history of philosophy, it hasn't really been a focus of attention; I borrowed the term from a comment by Simon Blackburn in his little introductory book called *Think*.[2] The idea of conceptual engineering is really taking an active role with respect to our conceptual scheme and changing it when one finds defects in those concepts. In earlier work, I argued at length that the concept of truth does have this kind of defect and is ripe for replacement.

I'm not going to develop that project; instead what I'm going to do is introduce something that I've come to believe as a result of engaging in it. I've come to think that conceptual engineering can and should play a much larger role philosophical theorizing. Indeed, I've come to think that most, if not all, commonly discussed philosophical concepts are inconsistent. Some in the way the truth is inconsistent, others in more subtle ways. Some based on the way they interact with one another, and others just by themselves. As such, I have come to think that philosophy is for the most part the study of what have turned out to be inconsistent concepts.

One way to make sense of this idea of an inconsistent concept is to say that concepts have constitutive principles—principles that constitute that concept in the sense that they tell you which concept it is, and often those constitutive principles are inconsistent with one another or with obvious facts about the world. The concepts that I think are inconsistent include truth, knowledge, nature, meaning, virtue, explanation, essence, causation, validity, rationality, freedom, necessity, person, beauty, belief, goodness, space, time, and justice. So when I say I think that philosophy is for the most part the study of inconsistent concepts, that's really what I mean. I think those are all inconsistent concepts; those are all defective concepts.

1. The Radical Therapeutic Program

I want to stay a little bit about the role that I think conceptual engineering should play in philosophical methodology, and I in order to do that I want to paint a picture of what I think philosophy is like—an account of the nature of philosophy. One way to do this is to appeal to some folks from history of philosophy, namely, Socrates, Nietzsche, and Wittgenstein.

From Socrates, and by 'Socrates' here I mean the early Platonic Socrates, the one from the early dialogues of Plato, we get the idea that the unexamined life is not worth living.[3] I take it he means the life without critical thinking. Subjecting ones' beliefs to critical scrutiny is the key aspect of critical thinking. So if your life is without critical thinking, then it's not worth living. Critical thinking is an essential aspect of living the good life—being in the right way.

Nietzsche. The idea from him is that in the absence of any divine or objective standards for human life we ought to craft our own.[4] That is, you want to take an active attitude toward your own life—for the creation of the structure of one's own life.

And from Wittgenstein, I take the idea that philosophical problems are manifestations of being trapped by our language, and philosophy should take the form of

[2] Blackburn (1999). [3] Plato (1961). [4] Nietzsche (1886).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

therapy that ultimately dissolves the philosophical problems.[5] The key claim here is that the aim of philosophy is to show fly the way out of the fly bottle. That's a nice metaphor that Wittgenstein used, and I'm going to take that metaphor pretty seriously.

Conceptual engineering is taking a Socratic, that is, critical, and Nietzschean, that is, active, attitude toward one's own conceptual scheme. Right now, many of us think that we should already take this attitude toward our beliefs. For example, we should subject our beliefs to a battery of objections, see how well we can reply to those objections, and if a belief doesn't fare well in this process, then that's a good indicator that it should be changed. And by doing this—by subjecting one's beliefs to critical scrutiny—one can craft and sculpt and mold a system of beliefs for oneself, You can do that rather than just doing what most other people do, which is borrowing a set of beliefs from whoever raised you and living throughout your whole life without really thinking about them very much. It turns out that the beliefs you borrow from your ancestors might work pretty well today, but there are also going to be some places where they don't work very well.

I think we should take this same attitude toward our concepts. The central idea of conceptual engineering is that we ought to take the same critical attitude toward our concepts. Likewise, if a concept doesn't fare well under critical scrutiny, the active attitude kicks in and one crafts new concepts to do work wants without giving rise to the problems inherent in the old ones. By doing this, one can sculpt and craft a conceptual repertoire of one's own, rather than just living one's life with the concepts borrowed from one's ancestors. Conceptual engineering is doing to concepts what most of us already think we should be doing to our beliefs. A nice quote from Alexis Burgess and David Plunkett: "Our conceptual repertoire determines not only what we only what we can think and say, but also, as a result, what we can do and who we can be." Through conceptual engineering, we can take some control over what we can think, say, do, and be.

Conceptual engineering can been seen as in the service of an overarching therapeutic program in the spirit of Wittgenstein. However, Wittgenstein's infamous conservativism, that philosophy should leave everything as is, has no part in this project at all. Our beliefs are not fine as they are. Our concepts are not fine as they are. But we can make them better.

The radical therapeutic program does share with Wittgenstein's methodology the goal of showing the fly the way out of the fly bottle. And conceptual engineering can help. Consider the thesis that philosophy is the study of what have turned out to be inconsistent concepts. Put this idea into a Wittgensteinian program, and you get the following picture. Philosophers are arguing about how best to make sense of concepts that are actually inconsistent. We're trying to figure out how to analyze concepts that are internally defective. The arguments one finds in philosophy journals and books rely on privileging certain constitutive principles here and others there, but ultimately the debates rarely make discernible progress because the concepts being analyzed and the concepts used to conduct the debate are defective. That is one reason

[5] Wittgenstein (1953).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

philosophers end up dealing with so many paradoxes and conceptual puzzles. And it can sometimes seem like that is our whole job.

That's the fly bottle. That's where you are and where I am. Trapped by our most important and cherished concepts into accepting absurdities and reasoning our way to contradiction.

How do we escape? For the past 400 years, the domain of philosophy has been shrinking. That is a sociological fact. Physics, geology, chemistry, economics, biology, anthropology, sociology, meteorology, psychology, linguistics, computer science, cognitive science. Each of those subject matters was a part of philosophy a mere 400 years ago. And we are talking about a discipline that has 2,600 year history; 400 years in 2600 years is nothing. As the scientific revolution ground on, more and more sciences were born. This process is essentially philosophy *outsourcing* its subject matter as something new. As sciences. The process is rather complicated, but an important part of it is getting straight on right concepts to use. That subject matter then gets brought under scientific methodology.

Showing the fly the way out of the fly bottle is taking an active role in this outsourcing process—pushing faster and faster to identify the conceptual defects (the Socratic idea), craft new concepts to avoid the old defects (the Nietzschean idea) with an eye toward repairing the philosophical subject matter for outsourcing as a science.[6] The ultimate goal of the process is the potential end of philosophy. Escape for the fly. The end of philosophy is only potential because it is likely that new technologies and ways of life will give us new inconsistent concepts that are philosophically significant, and these will need to get sorted out as well. So it is not obvious that our stock of defective concepts is ever going to effectively decrease. It really depends on how much conceptual engineering occurs. Speeding it up is up to us. The speed with which we get new defective concepts is mostly *not* up to us. People just make them up as they are needed or wanted. Nevertheless, we can envision a world where we have succeeded in making philosophy evaporate. Sometime after that, it might show up again with new philosophically significant defective concepts. And after that, philosophy might break out during especially rapid technological or social growth. Somewhat like acne.

That's the idea I want to call *the radical therapeutic program*. It calls for actions that might ultimately do away with philosophy.

The scientific element in this radical therapeutic program, which I call *metrological naturalism*, is separable from the conceptual engineering element. However, the two go together well. Metrological nationalism is more successful with consistent concepts, and in order to do conceptual engineering, we need to know what kind of replacement concepts to aim for. One might even say that metrological naturalism without conceptual engineering is empty, and conceptual engineering without metrological nationalism is blind.

Contrast this radical therapeutic program that I've just outlined with the most prominent philosophical methodology of our time, philosophical analysis in general,

[6] Contrast this position with Capellan's recent work in which he argues that we have little or no control over the concepts we use.

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

and the Canberra Plan in particular, which owes much to David Lewis. According to Lewis' methodology, one begins by assembling the platitudes for a philosophical term, and then one tries to figure out which real, relatively fundamental thing the platitudes might describe. If the platitudes are inconsistent, one tries to find something that satisfies a weighted majority of them—try to figure out what comes the closest to satisfying them and that is what the philosophical term designates. That's it.

The dominant, Lewisian methodology is static, and so has nothing to do with change and improvement. Lewis writes,

> One comes to philosophy already endowed with a stock of opinions. It is not business of philosophy either to undermine or to justify these pre-existing opinions to any great extent, but only to try to discover ways expanding them into an orderly system.[7]

I think it is hard to be more wrong than that.

2. Conceptual Engineering

There are several conceptual engineering projects already in philosophy, and I want to lay out two of them. One kind of project is labeled *amelioration* by Sally Haslanger. She argues that we need to change various key terms, especially terms associated with gender and race, for social justice reasons.[8] For example, 'woman,' the word, seems to currently express the concept of an adult human female. However, Haslanger argues that 'woman' is used primarily to subordinate people based on their stereotypical female characteristics. She suggests that the word 'woman' should instead be used to express something like the concept of a person subordinated based on stereotypical female characteristics. She wants to change the word 'woman,' and by doing that, the goal is to fight that subordination by making it explicit in the concept expressed by 'woman.' Rather than having it be something implicit, she suggests making the subordination explicit so that is something that is right there in front of one's face. Then, if the word 'woman' is used in this new way, a supporter of social justice can have as a goal the elimination of women—the elimination of people who are subordinated by appeals to their stereotypical female characteristics. Haslanger's amelioration project is obviously a conceptual engineering project, and there are clear similarities between her project of mine.

There are conceptual engineering projects in contemporary metaphysics as well. For example, Theodore Sider's introduction of the concept of structure as a generalization of David Lewis' notion of naturalness.[9] Another conceptual engineering project in this area concerns the idea that some or all metaphysical disputes are pointless. Metaphysics has come under attack lately, and one of the major criticisms is that metaphysical disputes are merely verbal, where the participants are just talking past one another rather than having a substantive disagreement. For example, how many things there are in a universe with two simple objects (i.e., objects with no proper parts). Are there two things in that universe or are there three things (the two simple things plus the one complex thing made up of the two simple things as parts)?

[7] Lewis (1973: 88). [8] Haslanger (2000). [9] Sider (2011).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

One view is that two people engaged in that dispute are just talking past one another—the dispute is merely verbal, and so not worth having. It is a merely verbal dispute because some key term used in it—maybe the word 'thing' or maybe the quantifier 'there is'—means different things for the people involved. Those who advocate this position are called metaphysical deflationists, and they employ some sophisticated tools with which to formulate and defend this criticism. Probably the most well-known tool is quantifier variance, which is the idea that there are multiple equally good interpretations of what people mean by the existential quantifier, 'there is,' involved in formulating these metaphysical questions. Those engaged in the ontological disputes—the metaphysicians—are simply talking past one another according to this criticism.

The metaphysicians have strong objections to this criticism, but they also have proposed a new kind of strategy for conducting metaphysical disputes just in case the metaphysical deflationists turn out to be right. The strategy is called *Plan B* by Sider who is one of its primary advocates. Plan B is to give up using natural languages like English for doing ontology, for doing metaphysics, and instead stipulate a fundamental meaning for existential quantifiers in a new language often called *Ontologese*. According to Plan B, metaphysical disputes can then be conducted using this new language and its new existential quantifier, which is stipulated to be fundamental. That is obviously a conceptual engineering project. This project is less focused on improving our conceptual scheme in light of discovering one of our concepts is defective, and it is more focused on avoiding criticism by talking in a new way. Still, it is a conceptual engineering project.

3. Metrological Naturalism

There are two aspects to the radical therapeutic program outlined above: conceptual engineering and metrological naturalism. Let us turn to metrological naturalism. *Metrological* naturalism is a kind of *methodological* naturalism, which is a somewhat unpopular philosophical methodology these days. Methodological naturalism is captured by the following quote from Wilfrid Sellars. "In the dimension of describing and explaining the world, science the measure of all things."[10] Moreover, scientific methods are the most reliable route to true beliefs despite scientific results being fallible. Finally, philosophy should be continuous with the sciences in two senses: (i) sciences do not require justification or grounding from philosophy, and (ii) philosophy and the sciences should pursue similar goals and employ similar methods. Philosophers have typically for a very long time now taken it as our job to assess whether scientists are spending their time productively, and we often assumed that science is in need of some kind of justification from first principles; that is, philosophers have historically taken sciences to be ungrounded unless they have been provided a grounding from philosophy. I suppose this isn't surprising given the philosophical sources of all the sciences, but that is no reason to retain it, and methodological naturalism rejects it. Moreover, methodological naturalism

[10] Sellars (1956).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

enjoins philosophers to look to the sciences for methods and goals. Beyond that vague advice, there isn't much agreement about how to be a methodological naturalist.

Here is a problem for methodological naturalism: it seems like many of the things that the naturalist says are not scientific. For example, no science concludes that science is the best road to truth about the physical world. That is not a scientific result. Another problem is that there is no consensus about what demarcates science from non-science. The demarcation problem is notorious as a standard example of a problem in philosophy that has received tons of attention but resists any solution. And there is no consensus on the nature of scientific methodology either. Obviously there is observation, hypothesis, prediction, and experiment, but beyond those kinds of banalities, it is hard to say exactly what the scientific method consists in, and it is not exactly obvious that someone who is working at the Large Hadron Collider, the particle accelerator, is following the same kind of methodology as a sociologist who studies rural food distribution networks or as a biologist who is investigating fetal development in marine mammals. Also, many philosophical topics are abstract and so resist scientific methods, which tend to emphasize causal interactions. And finally, scientific methods aim for descriptive results, but many philosophy philosophical topics are normative.

The version of naturalism that I want to advocate is called *metrological* naturalism because the Greek word for measure is '*μέτρον*' (*metron*). My version is a *measurement-theoretic* methodological naturalism, which we can therefore call *metro*logical naturalism.

The big idea from scientific revolution is that we can use mathematics to describe, explain, predict, and control the world around us. That is the idea that comes out in Galileo's very nice formulation, "mathematics is the language of nature."[11] It is since the early 1600s that we have had this idea, but it was not really until the 1800s, in fact that late 1800s, that theorists turned their attention to providing a *scientific* understanding of how mathematics is applied in this way in the sciences. The result of these investigations is called *measurement theory*: the study of how mathematics applies to the natural world in the sciences. You can think of measurement theory as an all-purpose foundation for scientific theorizing in much the same way that set theory can be thought of as an all-purpose foundation for mathematics. Metrological naturalism has as a methodological principle that philosophers should use measurement theory as a guide or model in philosophical theorizing.

According to this methodology, we should be using resources of the sciences in philosophical theorizing. There are many ways of doing this, but let me just briefly present three of them. First, cast your philosophical theories of X, where X is whatever philosophical concept you want to be thinking about, as measurement systems for X. There is a large literature on how to construct measurement systems, and we know pretty well how to do this for things like length and weight. Trying to figure out how to construct a measurement system for something like truth or justice is a lot more complicated, but this isn't just an analogy.

[11] Galilei (1623) as translated in Popkin (1966: 65).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

For example, Donald Davidson, over the course of a famous career, crafted a detailed measurement system for belief, desire, and meaning.[12] He called it the unified theory. In the case of length, we start with people's judgments about which of two objects *extends beyond* the other when the two are lined up, and thinking of two objects *end to end* as a single thing having a length just like any regular object. We then translate these judgments into a different language—a language with terms like 'longer than' and 'concatenation.' Call it the relational language. If object A extends beyond object B, then A is longer than B. If A is end to end with B then there is an object C that is identical to the concatenation of A and B. One benefit of this translation is that 'longer than' and 'concatenation' are very clearly behaved with explicit laws defining them, whereas the judgments we started with are pretty messy. From here, the next step is to prove a certain kind of result called a representation theorem, which says that we can translate from the language with 'longer than' and 'concatenation' into a mathematical language, which contains 'greater than' and 'plus' and numerals as well. Proving the representation theorem shows that you can use the numbers talked about by the mathematical language to keep track of the objects talked about by the relational language. Each object gets a number, and the number assigned to A is greater than the number assigned to B if and only if A is longer than B. And the number assigned to A plus the number assigned to B equals the number assigned to the concatenation of A and B. The representation theorem says that you can assign numbers to the objects in a way that makes all these principles, and a bunch of other ones, true. Another kind of result, called a uniqueness theorem, says how many different ways there are of assigning numbers like this. Overall, the measurement system begins with basic judgments about one object extending beyond another and putting two objects end to end. And the measurement system delivers something remarkable! A number for every direction of every object so that we can use these numbers to do everything from carpentry to identifying the distances to distant galaxy clusters. And we also get a plethora of length scales from stadia to megaparsecs. All of that comes from making some basic assumptions about 'longer than,' 'concatenation,' and how to translate them into basic English (extends beyond and end to end) and into mathematics (greater than and plus).

In Davidson's work, the same kind of measurement system is laid out for belief, desire, and meaning. He makes some assumptions about the how people think about sentences that they think are true, and some assumptions about people's preferences. Together these allow him to show how to assign truth conditions to the sentences of the language in question and beliefs and desires to the person in question. In his demonstration, Davidson utilizes a fictional character called the radical interpreter, who goes through a certain procedure to figure out meanings, beliefs, and desires, from holding-true and preferences. This procedure is showing a representation theorem for the measurement system. The same idea is worked out in considerably more detail in Robert Matthews' book *The Measure of Mind*.

[12] Davidson (1990).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

So the first bit of advice from metrological naturalism is to cast one's philosophical theories as measurement systems in much the same way that Davidson and Matthews do. Second, focus on semantic theories of philosophical locutions rather than trying to analyze philosophical concepts. Philosophers have been, since the early twentieth century, trying to analyze concepts. Even today the majority of talks at most major conferences are dedicated to philosophical analysis projects, where the philosopher is attempting to put forward a philosophical analysis of some concept like obligation or beauty or reason.

Instead of doing philosophical analysis or conceptual analysis we should be focused on semantics for philosophical terms, because arriving at an adequate semantic theory for a philosophical term can often cut through many of the unnecessary and confusing assumptions associated with that term. One example here would be the semantics for reasons project I have jointly undertaken with Brian Weaver.[13] Getting straight on the semantics for reasons locutions helps tremendously in assessing traditional philosophical debates about reasons. For example, there is a debate between factualists, who say that reasons are facts, and mentalists, who say that reasons are mental states. Understanding the proper semantics for reasons locutions exposes that debate as not very substantive.

Third, utilize the tools of measurement theory for answering philosophical questions. For example, invariance and symmetry play major roles in measurement theory, and they can be used to make sense of the philosophically significant distinction between objective and subjective phenomena. Consider the claim, 'In Scotland, it is twice as hot in May as it is in March'; one might feel compelled to assert this sentence because the average temperature in May is around 10°C and the average temperature in March is around 5°C. However, if we transform these values from the Celsius system to the Farenheit system, we get: 50°F in May and 41°F in March. But of course 50 isn't twice 41. The lesson is that one *can* of course multiply temperature measurements, but multiplication with temperature measurements isn't objective. When one says that it is twice as hot in Scotland in May as it is in March, one is talking about one's way of representing the world (using the Celsius scale) rather than talking about the world itself. Here the key to identifying objective features of the world is invariance. Something is an *objective* feature of temperature if it is invariant across different scales, whereas something is a *subjective* feature of how we measure temperature if it fails to be invariant. In contrast to the case of temperature, multiplication is an objective feature of length because it is invariant across length scales; for example, 20 feet is twice 10 feet, and 6.096 meters is twice 3.048 meters.

It is worth highlighting some features of metrological naturalism:

(i) It says nothing about the methods one uses to arrive at or justify philosophical theories. So metrological naturalism is not opposed to *apriori* methods (e.g., intuitions, deductions) in philosophy.

(ii) It is not opposed to *apriori* or theoretical philosophical claims. In fact, it is plausible to think that certain aspects of a measurement system for some concept will be constitutive of that concept.

[13] Weaver and Scharp (forthcoming).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

(iii) It does not offer analyses of concepts or necessary and sufficient conditions for concept application.
(iv) It does not offer reductive explanations in any sense.
(v) It is not a metaphysical thesis about what exists or does not exist.
(vi) It does not need a leading science, a criterion for what demarcates science from non-science, or an account of scientific methodology.
(vii) It is applicable to abstract topics (e.g., one can do measurement theory for mathematics and logic).
(viii) It is applicable to normative topics (e.g., formal axiology).

There is much more to be said for metrological naturalism both as an independent methodology and as a companion to conceptual engineering, but the above will have to do for this occasion.

4. Elements of Conceptual Engineering

Conceptual engineering is actively changing some aspect of our concepts—eliminating bad ones, adding new ones, deciding which ones we should use for which purposes, and choosing which words should express which concepts. There are plenty of instances of conceptual engineering in the history of philosophy, and we have considered a couple above.

Another term that is used often in this area of research is 'conceptual ethics,' but the two are distinct. *Conceptual ethics* is the study of evaluative and normative issues associated with our concepts and the words that express them. *Evaluative* issues are those pertaining to how good something is; for example, oxygen is a better concept than phlogiston, and the luminiferous ether is not a good concept. *Normative* issues are those pertaining to obligations and permissions—to what we ought to do and what we may do; for example, we ought to use the word 'woman' to express the concept of someone oppressed on the basis of stereotypical female characteristics, and we ought not use the concept of mass from Newtonian mechanics when calibrating the atomic clocks on GPS satellites.

Conceptual ethics is clearly involved in conceptual engineering because the latter often relies on evaluative and normative judgments about our concepts. However, not all conceptual ethics is conceptual engineering. For example, judging that our concept of mental illness is just fine for our purposes is doing conceptual ethics, but it isn't doing conceptual engineering. And the converse holds as well. For example, establishing a relative consistency proof for an axiomatic theory of ascending truth and descending truth, which have been suggested as replacements for our defective concept of truth, is a part of conceptual engineering, but it isn't a part of conceptual ethics. Hence, I do not see the two terms as competitors for describing a single area of philosophy.

Among conceptual engineering projects, two kinds deserve to be singled out as significant and distinct: conceptual revision vs. conceptual replacement. *Conceptual revision* is changing a concept so as to improve it in some way, but the concept persists through whatever changes happen to it. It is the same concept before and after the revision to it. I think some people conceive of Carnap's method of

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

explication as a kind of conceptual revision by adding a degree of clarity to an otherwise fuzzy concept. Sally Haslanger is often read in this way, but it is isn't obvious that it is the most accurate interpretation.

On the other hand, *conceptual replacement* doesn't cause any changes to any concepts at all; rather, these projects introduce new concepts to our conceptual scheme and prescribe a particular role for them to play; this role might already be filled by one of our existing concepts, so a replacement project might suggest that one of our existing concepts isn't cut out for one of the jobs we think it can do. Conceptual replacement is the kind of conceptual engineering project I take up with respect to the concept of truth in my book *Replacing Truth*. There I argue that truth is an inconsistent concept, and I offer two replacement concepts that, together, will do some of the work we have been using truth to do. Note that replacement does not entail elimination—we still retain the defective concept of truth, because in the vast majority of situations, we can use it without running into any trouble whatsoever. However, we do eliminate one or more roles for truth to play once we have the replacement concepts.

To illustrate the distinction between conceptual revision and conceptual replacement, we can think about three distinct readings of Haslanger's conceptual engineering project for the concept expressed by 'woman.' The variance is due to the fact that Haslanger both presents her project as actively choosing to do something to the concept expressed by the word 'woman' and appeals to a particular thesis in philosophy of language. She even suggests that 'woman' right now expresses one concept—the concept of being an adult human female—but that we should, for the purposes of social justice, choose to make it express a different concept—the concept of an adult human who is oppressed on the basis of stereotypical female characteristics. She also sometimes suggests that instead of picking a new concept, we are changing the existing concept expressed by the word 'woman.' Moreover, Haslanger also appeals to a controversial, but highly discussed thesis in the philosophy of language and mind called *semantic externalism*, which entails that the concept one uses to think or the concept that is expressed by a certain word is, to some extent, determined by the physical or social environment of the person thinking or the person using that word. As a result, it could be that 'woman' already expresses the concept of an adult human who is oppressed on the basis of stereotypical female characteristics. It might already express that concept because that is the role the word 'woman' plays in our social structures, whether we know it or not. So, Haslanger's project involves an element of conceptual engineering—choosing for our word 'woman' to express a certain concept—and it sure seems like it doesn't express that concept right now. And her project involves a commitment to semantic externalism as well, which to some extent, takes the control over which concepts our words express away from us. The three distinct readings of what Haslanger is up to are based on privileging these different aspects of her project.

1. A *conceptual revision* project. Prior to Haslanger's work, the English word 'woman' expressed the concept of an adult human female, but Haslanger suggests we ought to change this concept so that its content is *an adult human oppressed on the basis of stereotypical female characteristics*. That is,

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

Haslanger enjoins us to change our concept of woman so that it has a slightly different content. Although Haslanger sometimes talks as if these are distinct concepts, what she means is that they are distinct contents or readings of the same concept, before and after a change in that concept.

2. A *conceptual replacement* project. Prior to Haslanger's work, the English word 'woman' expressed the concept of an adult human female, but Haslanger suggests that we ought to change which concept is expressed by this word 'woman' so that it expresses the concept of an adult human who is oppressed on the basis of stereotypical female characteristics. That is, Haslanger enjoins us to change which concept is expressed by the word 'woman' so that it expresses a different concept. Although Haslanger sometimes talks as if these are the same concept, what she means is that they are distinct but similar concepts.
3. A *belief replacement* project. Prior to Haslanger's work, the English word 'woman' expressed the concept of an adult human who is oppressed on the basis of stereotypical female characteristics. However, no one really believed that it expressed this concept. Haslanger suggests we ought to change our beliefs about which concept is expressed by the word 'woman,' so that we stop believing it expresses the concept of an adult human female and start believing that it expresses the concept of an adult human who is oppressed on the basis of stereotypical female characteristics. That is, Haslanger enjoins us to change what we believe about the concept expressed by the word 'woman' so that our beliefs are true instead of false. Although Haslanger sometimes talks as if we should change the concept expressed by the word 'woman,' what she means is that we should have true beliefs about the concept expressed by the word 'woman.'

The third reading is the one that results from emphasizing semantic externalism over conceptual engineering. The first two projects result from emphasizing conceptual engineering over semantic externalism, and they differ on whether Haslanger offers a conceptual revision project in reading one or a conceptual replacement project in reading two. For what it is worth, I prefer the project described in reading number two, but it doesn't leave much room for semantic externalism to play a role in the account. Moreover, the standard term for the study of how rational agents change their belief systems is belief *revision*, not belief *replacement*. I've chosen the latter to be consistent with my distinction between conceptual revision and conceptual replacement. Belief revision, according to my usage, would be somehow changing a single belief so that it was still the same belief but was somehow different in content. This sort of thing is impossible according to most analytic philosophers because we tend to individuate beliefs by their contents.

5. Tools for Conceptual Engineering

There are many approaches one might take in conceptual engineering. Here I canvass three. Each of these has something to be said for it, and they need not be competitors. Nevertheless, they do have their own problems and pitfalls, so it helps to be clear about which of these tools one is using. Each of these tools is essentially a standard of

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

evaluation—a way of assessing concepts that is relevant to potential conceptual engineering projects. As such, they all belong to conceptual ethics as well, and to the evaluative branch of conceptual ethics in particular.

The *metaphysical approach* is to evaluate our concepts for naturalness or fundamentality—how well they carve nature at its joints. One can see this approach in Sider's recent book, *Writing the Book of the World*. There he advocates fundamentality as the primary epistemic virtue associated with concepts. Usually naturalness or fundamentality are taken to come in degrees, so we can speak of one thing being more fundamental than another and of something being relatively fundamental. Presumably, Sider would accept relative fundamentality as a basic evaluation of concepts—other things being equal, more fundamental concepts are to be preferred.

The *pragmatic approach* is to evaluate our concepts for how well they work. How well does a concept do what we use it to do (or what we ought to use it to do)? This seems to be the genus for ameliorative projects like those advocated by Haslanger, and I often cast my truth project in these terms. Overall, talk of jobs or purposes for concepts is no more heavyweight than talk of things we do with concepts. For example, we formulate truth conditional semantic theories for gradable adjectives (like 'tall'), and many of the most powerful of these semantic theories have the word 'true' in them in way that doesn't seem to be eliminable. From this we can conclude that one job of the concept of truth is serving an explanatory role in these semantic theories.

The *constitutivity approach* is to identify the constitutive principles for a concept and then evaluate those principles using our toolkit for evaluating beliefs. Are they true? Are they justified? Are they helpful? On my view, constitutivity is a pragmatic issue. A constitutive principle for a concept is a principle that is used to guide interpretation—if my interlocutor rejects a principle that I take to be constitutive for a concept that figures in our conversation, then that is a *pro tanto* reason to think that we do not mean the same thing by the word in question. Constitutivity is a descendent of analyticity, but there is no reason to think that constitutive principles should be true or vacuous or uninformative or *apriori*. Moreover, one can reject a principle constitutive of a certain concept without thereby losing possession of that concept. By focusing on constitutivity, we can characterize individual concepts as inconsistent if their constitutive principles are inconsistent with established facts (e.g., truth), and we can characterize mutually inconsistent groups of concepts.

The constitutive approach also allows us to model concept change using the elaborate and influential tools of formal epistemology on belief change (e.g., AGM (Alchourrón, Gärdenfors, and Makinson) theory; Alchourrón et al. 1985). These are nice tools and I think they can help us identify conceptual clashes and evaluate proposed conceptual engineering projects. Some examples: (i) one can think of getting rid of a pejorative concept as contraction (belief elimination), (ii) we can define the entrenchment of a concept by the average entrenchment of its constitutive principles, (iii) belief replacement is the adding of a new belief to the set, and that is akin to what happens in conceptual replacement, which is the particular kind of conceptual engineering that I have been pursuing, and (iv) consolidation is an operation on an inconsistent belief base, and this can be thought of as identifying potential replacements in a conceptual engineering.

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

6. Constraints on Conceptual Engineering

I have laid out the radical therapeutic vision of what philosophy is all about. I have presented conceptual engineering as one aspect of a proper philosophical methodology. We should take an active role in altering and improving our conceptual scheme. I have also advocated a scientific element in this philosophical methodology that I have called metrological naturalism.

One of the big things that I think needs to be explored is the extent to which there are important constraints on conceptual engineering. I can imagine a debate over the legitimacy of a certain kind of conceptual engineering project. Consider the following argument.

The pro-choice position (abortion is morally permissible in normal circumstances) and the anti-infanticide position (killing an infant is not morally permissible in normal circumstances) are the right ones to have with respect to abortion and infanticide. Furthermore, when I reflect on the nature of time, I find myself committed to the idea that I do not have temporal parts—I am wholly present (metaphysically) at every moment. If so, then I am an endurantist, and presumably I am an endurantist about a zygote/infant, which is the thing that is wholly present throughout the change from having no rights (as a zygote) to having rights (as an infant). So far, so good, but now I start thinking about metaethics and theories of justice, and I arrive at the plausible view that rights are properties had by objects and that our judicial locutions denote these properties; so I'm a realist about rights. Now I think about it a bit more and I find it difficult to believe that having rights is not an intrinsic property of an entity if one is a realist about rights. So now I have a problem making sense of how a single thing could be intrinsically killable at one time (as a zygote) and then not intrinsically killable later on (as an infant). Now I need to think hard about which concepts of time, persistence, and rights I should use, given my commitments to prochoice/anti-infanticide positions.

Is this a good reason to replace my concept of time, concept of rights, concept of persistence or any of the related concepts appealed to in the inset reasoning above? To be clear, I am not attributing this project to anyone in particular. I am just trying to think through the kinds of conditions that one might want on an acceptable conceptual engineering project.

One way to think about it is that I surely have a reason to change my concepts in this way, but it is *the wrong kind of reason*. Reasons for belief of the right kind are those relevant to the truth of the belief, whereas reasons for having a belief, which are irrelevant to the truth of the belief, are reasons of the wrong kind. There is a larger philosophical literature on reasons of the wrong kind. Pascal's wager is a good example; Pascal's reasons for belief in God are entirely focused on the consequences of having that belief, rather than on whether the belief is true. These same considerations seem to suggest that to the reason for changing our concept of time or our concept of rights is *not* a reason to think the new concept is a good one. Instead, the reason given is maintaining certain moral and social commitments (e.g., pro-choice and anti-infanticide). This is a reason to change the concepts I use, but it is not a reason to think that the new ones are valuable or right for the job. As such I am somewhat uneasy about the conceptual engineering project that might be inspired by the inset argument.

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

On the other hand, conceptual engineering is essentially oriented to action, not belief. Making changes to our conceptual scheme or our language is an activity. And it isn't clear whether there are reasons of the wrong kind for *actions*—notice all the examples given involve beliefs instead of actions. For example, there is no difference in kind between volunteering at a homeless shelter because I want to help the homeless and volunteering at a homeless shelter due to receiving some incentive. Hence, it might be that there are no reasons of the wrong kind to promote moral rightness. So promoting moral rightness is on par with any other reason when it comes to adopting a certain concept of time or a certain concept of persistence, or a certain concept of rights. This result would have tremendous consequences—given the complexity and abundance of connections between concepts, this policy for conceptual engineering would effectively moralize and politicize our entire conceptual scheme. Even logic or mathematics could conceivably be affected.

My attitude on this issue hasn't been entirely stable, but I tend to err on the side of letting a thousand flowers bloom rather than figuring out from the armchair which conceptual engineering projects are kosher. Still, I can envision balking at certain proposals if they were to, for example, promote conceptual confusion or inconsistency for political gain. I think conceptual engineering should always make our conceptual scheme better for us and the concepts we use better for what we use them for. I like to think of conceptual engineering as a wide category, but it certainly has limits.

7. Non-Scientific Exports

I'll consider three objections, one here and the others in the next two sections.

It is clearly false to think that as the subject matter of philosophy shrinks, this subject matter is exported only in the form of sciences. So the entire part of the radical therapeutic program dedicated to helping us escape our predicament of a conceptual scheme dominated by inconsistency (the fly bottle) is baseless.

It is surely right that as the subject matter of philosophy has constricted, science has not absorbed all of it. For example, astrology was a huge part of western philosophy and a primary driver of innovation in astronomy from antiquity until last couple of centuries. Which sciences ended up with this subject matter? None. Throughout the history of western philosophy, there have been changes that eliminate some subject matter from philosophy as not fit for philosophical thinking. Astrology is one example, and sophistry is another expulsion, but it occurred very early in the tradition by those under the influence of Socrates and Plato. In our own modern period, it has become unacceptable within western philosophy to make appeals to God or God's works to explain philosophical puzzles, except at the current time in philosophy of religion and parts of metaphysics. That is a huge change that has happened over the past few centuries. Think of how important appeals to God are in the solution to scepticism according to the orthodox reading of Descartes' *Meditations*.[14] Now imagine trying to publish a paper in a top journal today arguing

[14] Descartes (1641).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

that God's works are the best solution to a problem like the liar paradox. Unthinkable! So philosophy often dumps parts of its own subject matter and these dumps need not be exports to the sciences. They need not even be illegitimate—international relations comes to mind as a topic that started as part of philosophy in the early 1800s, but is now its own respectable discipline alongside philosophy in the university, despite the fact that international relations isn't a science.

Nevertheless, my point is unaffected by this complexity. Philosophy changes in all sorts of ways, and one of the most significant and impactful changes it has undergone is the colossal outsourcing of its material to the sciences. Here is a nice quote from Alexis Burgess and Brett Sherman about what has happened just in philosophy of language and just in the last few decades.

> It's not easy being a philosopher of language these days. Work is hard to come by, and we don't just mean jobs. The subject matter itself seems to be getting smaller and smaller. What were once proprietary issues in the field (like the semantics of names, descriptions, quantifiers, etc.) are now quite rightly seen as scientifically tractable research programs in linguistics and psychology. . . . [T]he marked progress of linguistic semantics obviously owes volumes to the foundamtional work of philosophical luminaries like Frege, Tarski, Davidson, Montague, and Lewis, who helped erect a basic framework for articulating and evaluating claims about verbal meaning. As these foundations have solidified, however, questions once assumed amenable to *apriori* reflection have been exposed as properly empirical quarries. Handmaiden to the science of meaning might be a perfectly respectable job title. But some of us who self-identify as philosophers of language will naturally want to seek out new work.[15]

The process by which philosophy is giving way to the sciences on dozens of fronts is absolutely massive, and it is one of the most significant things that has ever happened in western civilization.

The radical therapeutic program has two parts—it identifies our problem, which is that most or all of our core concepts are inconsistent. And it offers a solution: use conceptual engineering to change our conceptual scheme so that we have concepts that work for us rather than concepts that tangle our thinking, confuse our beliefs, and interfere with our plans. To accomplish this, conceptual engineering aims for certain things in the new concepts, and that is where it relies on metrological naturalism. The solution—the way out of the fly bottle—is to promote the already massive exodus from philosophy that is science. The establishment of science doesn't have to be the only exodus from philosophy for it to be the most significant and for it to be our role model.

8. Defective Concepts in Conceptual Engineering

Another objection: How can I be sure that we won't find awful defects in the concepts employed by conceptual engineering projects themselves?[16]

My reply is that I can't be sure that these concepts aren't defective as well. In fact, I think they probably are defective. High on the list of probably defective concepts is

[15] Burgess and Sherman (2014: 1).

[16] Neil Tennant offered this question at the lecture.

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

the concept of a concept itself. It has well-known problems and many theorists engaged in conceptual engineering projects even go so far as to be concept eliminativists—they think that talk of concepts has no place in a proper conceptual engineering project. Instead, these theorists contend, we can make do with less controversial tools like extensions, which are just sets of individuals denoted by a predicate, and intensions, which are assignments of extensions to various possible situations. Herman Cappelen is a major proponent of the "no concepts" version of conceptual engineering, and there are others as well.[17]

However, there are two things to say to those in the "no concept" wing of the conceptual engineering movement. First, concept eliminativism is itself a conceptual engineering project, and I have yet to see anyone carry that project out in a careful and detailed way. So far, we have some proposals for how to do conceptual engineering without appealing to concepts, but we have very little in the way of reasons to think that this is a good idea. Moreover, there are bound to be more defective concepts utilized by the conceptual engineer than just this one, so if concept eliminativism is appropriate, then presumably other kinds of eliminativism with respect to the tools of conceptual engineering are appropriate as well, and it isn't clear that there will be enough left for the conceptual engineer to use for her projects.

However, the fact is that concept eliminativism is unjustified even if the concept concept turns out to be defective and in need of replacement. The reason is that defective concepts can still be useful, and even those who know they are defective can still employ them without thereby being irrational. For example, I think the concept of truth is seriously inconsistent, but I am not a truth eliminativist. The analogy I like to use is the concept of mass in classical mechanics. Mass is inconsistent concept but is still extraordinarily useful use it for all kinds of things from building houses to landing robots on comets. Think about how insane it would be to use general relativity to, say, design a sturdy bridge. It would be extremely unwieldy and it would take one far longer, and one would end up with the same bridge that would have been designed using Newtonian mechanics. Therefore, although it is likely that the concepts involved in my own methodology are themselves defective, that does not mean they are not useful for this purpose. When one aims to replace some concept, one tries to figure out whether the defect in that concept actually inhibits its utility—whether its defect actually gets in the way of certain applications. If the defect in a concept does undermine its utility for some purpose, then that is a decisive consideration in support of replacing that concept for that purpose. That is exactly the case with our concept of truth; its defects prevent it from effectively performing the role we ask of it in certain applications of natural language semantics. That is, when one tries to provide a truth-conditional semantics for a fragment of natural language that contains liar sentences, then one ends up with an inconsistent semantic theory. If I were to be shown conclusively that one of the concepts involved in my own methodology had a defect *that was impeding its utility in my methodology*, then that would be a problem for me, and I would focus attention on how to effectively replace that concept for my purposes. Therefore, there is a considerable gap between

[17] Cappelen (2018).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

the suggestion that some of the concepts I rely on are defective, and a substantive objection to conceptual engineering as I understand it and practice it. I admit the former, but the latter I have yet to see.

9. Is Philosophy about Concepts?

Here is another objection to what I have said so far: Philosophy isn't the study of concepts at all, so it cannot be the study of what have turned out to be inconsistent concepts.[18] Philosophers do on occasion study concepts, but only as one item among many in other things in the world. For example, there is a difference between the concept of truth and truth itself. Truth is, presumably, a property that things like sentences or theories or propositions can have, whereas the concept of truth is something like a mental representation or a constituent of thought or some other kind of thing that people grasp or possess or understand. Philosophy isn't the study of *the concept of* truth or *the concept of* knowledge or any of the other concepts. Instead, philosophy is the study of certain *phenomena*, like truth, knowledge, freedom, justice, and the rest.

That is all well and good as a start, but as philosophers we must to do better—we need to think a bit deeper about the issue. If our philosophical concepts are as defective as I have suggested, then there is no reason to expect there to be a property of truth or a property of knowledge or any of the rest. At least, not if one thinks of the property of truth as anything like what our concept of truth leads us to think it would be like, and if the property of knowledge is anything like what our concept of knowledge leads us to think it would be like. If the principles for these concepts are inconsistent, then no property can satisfy them. If they are seriously inconsistent, then no property can even come close to satisfying them.

For example, a philosopher might decide to study whether truth is a substantive property that can explain things or a deflationary property that doesn't explain anything. This is a huge area of contemporary philosophy covering the last half-century and involving hundreds of theorists and thousands of publications. One might think that such an inquiry has absolutely nothing to do with our concepts—it's about truth, not the concept of truth. But what, exactly, is the property of truth taken to be? Which property is it? It sure isn't the property had by a sentence 'grass is green' just in case it turns out that grass is green, and the property had by the sentence 'snow is white' just in case it turns out that snow is white, and in general the property had by the sentence <p> just in case it turns out that p. Why isn't it this property? *Because there is no such property.* To suppose there is such a property is inconsistent, as shown in the reasoning of the liar paradox. And there are dozens of other paradoxes associated with truth as well. In fact, truth is such a defective concept that no property satisfies even small subsets of the principles we think of as constitutive of truth. So there is no property of being true, not if that property is anything like what the concept of truth leads us to think it would be like. Philosophy cannot be about the

[18] Williamson (2007).

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

property of being true because there is no property of being true for philosophers to investigate.[19]

From the point of view of the radical therapeutic program, there might not be anything like what our philosophical concepts lead us to expect in the world. There might not be properties in the world that correspond to our philosophical concepts. In fact, there probably aren't. Perhaps there is *some* philosophical concept that is consistent enough for there to be something in the world that comes close to satisfying its constitutive principles, but that isn't the case for the vast majority of philosophical concepts. Hence, it makes the most sense to think of philosophy as the study of certain concepts—there isn't much else for it to be about. Even philosophers who think of themselves as studying genuine phenomena in the world are usually just exploring one aspect or another of an inconsistent concept. For example, internalists about knowledge and externalists about knowledge aren't investigating some property—the property of knowing something. Instead, each side takes for granted some of the constitutive principles for the concept of knowledge and uses them to argue against those on the other side in the debate, who take for granted *other* constitutive principles for the concept of knowledge. The debate seems interminable and deadlocked because each side is right—each side has latched onto some aspect of our concept, but each side is wrong as well, in that they reject some other aspect of our concept. The fact that internalists and externalists about knowledge—or pretty much any of the sides in any philosophical debate—can refute each other only shows that all the constitutive principles for knowledge, taken together, are inconsistent. That is, it only shows how defective the concept of knowledge is. Another way of putting the point, in terms of subjects or properties instead of in terms of concepts, would be that the subjects or properties that philosophers might think of themselves as investigating are delineated according to inconsistent principles. So there are no such things. The very idea that there is something like truth or knowledge or freedom or justice or virtue for us to investigate at all is inconsistent. Of course, we have the concept of truth and the concept of knowledge and all the rest, and philosophy is primarily the study of these concepts.

So if there is no such thing as truth or knowledge or freedom or virtue, then what is there? We don't know. And we won't know until we have done far more conceptual engineering.

10. Conclusion

Philosophy, or western philosophy at least, has been focused throughout its history on certain topics or certain concepts—truth, knowledge, value, virtue, freedom, justice, etc. The radical therapeutic program presented here is based on the idea most or all of these concepts are inconsistent. Or, alternatively, most or all of these subject matters are delineated in an inconsistent way. Our philosophical concepts, which are the heart and soul of our conceptual scheme, are organized and distinguished by principles that are themselves inconsistent with one another.

[19] See Scharp (forthcoming) for more details on this example.

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

The result is that just about any time we think or talk about philosophical topics and we try to follow these principles, we end up contradicting ourselves. That is our predicament. The solution sketched here relies on conceptual engineering—charting out the defects in and among our concepts and proposing new concepts that will do the work we demand without causing the problems we currently encounter.

Imagine, for a moment, what it would be like to have a consistent conceptual scheme. No paradoxes. No puzzles. Just clarity. We can do it. You can help.

Acknowledgements

Inaugural address delivered 25 February 2015 at The Ohio State University in Columbus, Ohio. The text has been changed only slightly, but has a few additions. Thanks to everyone present and to those in audiences at the University of Oslo and the University of St Andrews as well.

References

Alchourrón, Carlos, Peter Gärdenfors, and David Makinson. 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *The Journal of Symbolic Logic* 50 (2):510–30.

Blackburn, Simon. 1999. *Think*. Oxford: Oxford University Press.

Burgess, Alesis, and Sherman, Brett. (eds.) 2014. A Plea for Metasemantics. *Metasemantics*. Oxford: Oxford University Press.

Burgess, Alexis, and Plunkett, David. 2013. Conceptual Ethics I. *Philosophy Compass* 8:1091–101.

Cappelen, Herman. 2018. *Fixing Language: Conceptual Engineering and the Limits of Revision*. Oxford: Oxford University Press.

Davidson, Donald. 1990. The Structure and Content of Truth. *The Journal of Philosophy* 87:279–328.

Descartes, Rene. 1641. *Meditations on First Philosophy*. Trans. John Cottingham. Cambridge: Cambridge University Press, 1996.

Galilei, Galileo. 1623. *The Assayer*. Rome: Giacomo Mascardi.

Haslanger, Sally. 2000. Gender and Race: What are They? What do we want them to be? *Nous* 34:31–55.

Lewis, David. 1973. *Counterfactuals*. Oxford: Blackwell.

Nietzsche, Friedrich. 1886. *Beyond Good and Evil*. Trans. R. J. Hollingdale. Harmondsworth: Penguin Books, 1973.

Popkin, Richard. 1966. *The Philosophy of the Sixteenth and Seventeenth Centuries*. New York: Free Press.

Plato. 1961. Apology. In *Plato: Collected Dialogues*. Trans. F. Cornford. Princeton: Princeton University Press.

Scharp, Kevin. 2013. *Replacing Truth*. Oxford: Oxford University Press.

Scharp, Kevin. forthcoming. Conceptual Engineering for Truth: Aletheic Properties and New Aletheic Concepts. *Synthese*.

Sellars, Wilfrid. 1956. Empiricism and the Philosophy of Mind. In Herbert Feigl and Michael Scriven (eds.), *Minnesota Studies in the Philosophy of Science*, Volume I: *The Foundations of*

OUP CORRECTED PROOF – FINAL, 17/12/2019, SPi

Science and the Concepts of Psychology and Psychoanalysis. Minnesota: University of Minnesota Press.
Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
Weaver, Bryan, and Scharp, Kevin. forthcoming. *Semantics for Reasons*. Oxford: Oxford University Press.
Williamson, Timothy. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.
Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. Oxford: Blackwell.