Kinds of Kinds: Normativity, Scope and Implementation in Conceptual Engineering

Sarah Sawyer

University of Sussex

s.a.sawyer@sussex.ac.uk

Abstract

In this paper I distinguish three kinds of kinds: traditional philosophical kinds such as truth, knowledge, and causation; natural science kinds such as spin, charge and mass; and social kinds such as class, poverty, and marriage. The three-fold taxonomy I work with represents an idealised abstraction from the wide variety of kinds that there are and the messy phenomena that underlie them. However, the kinds I identify are discrete, and the three-fold taxonomy is useful when it comes to understanding claims about the normativity, scope, and implementation of conceptual engineering. In particular it reveals: first, that conceptual engineering in a broad sense is not distinctively normative; but second, that there is a distinctive normativity present in conceptual engineering in the narrow sense of amelioration; third, that conceptual engineering in the narrow sense of amelioration is only possible for terms that refer to kinds whose grounds we can change; and fourth, that the amelioration of such kind terms can be brought about as a result of changing the grounds of those kinds. The overall aim is to draw attention to the differences between kinds of kinds to make sense of the diverse and often conflicting claims in the conceptual engineering literature, both from its proponents and from its detractors. The hope is that paying attention to the differences between kinds of kinds will provide a better understanding of the landscape.

Keywords

Conceptual engineering; conceptual analysis; meaning; concepts; natural kinds; social kinds; philosophical kinds; implementation challenge; normativity; grounding.

1. Introduction

In this paper I distinguish three different kinds of kinds: traditional philosophical kinds, such as truth, knowledge, causation, necessity, and identity; natural science kinds, such as atomic number, force, spin, charge, and mass; and social kinds, such as class, poverty, economic growth, marriage, and misogyny. My aim is not to provide a comprehensive taxonomy of kinds. The three-fold taxonomy I work with represents an idealised abstraction from the wide variety of kinds that there are and the messy phenomena that underlie them. However, the kinds I identify are nonetheless discrete kinds of kinds, and the three-fold taxonomy is therefore useful when it comes to understanding claims about the normativity, scope, and implementation of conceptual engineering. In particular it reveals: first, that conceptual engineering in a broad sense is not distinctively normative; but second, that there is nonetheless a distinctive normativity present in conceptual engineering in the narrow sense of amelioration; third, that conceptual engineering in the narrow sense of amelioration is only possible for terms that refer to kinds whose grounds we can change; and fourth, that the amelioration of such kind terms can be brought about as a result of changing the grounds of those kinds. The overall aim is to draw attention to the differences between kinds of kinds in order to make sense of the diverse and often conflicting claims that are present in the conceptual engineering literature, both from its proponents and from its detractors. The hope is that paying attention to the differences between kinds of kinds will provide a better understanding of the landscape. It will become obvious as we proceed, but should be noted from the outset, that I presuppose a form of realism about the kinds in question which some may find controversial; but if alternative claims about the normativity, scope and implementation of conceptual engineering depend on an anti-realist account of philosophical, natural, and social kinds, then this itself will be an interesting result that has not yet been widely acknowledged.

The structure of the paper is as follows. In section 2 I characterise the differences between the relevant kinds. In section 3 I consider the implications of the differences for the claim that conceptual engineering is a distinctively normative endeavour. In section 4 I consider the impact of this for different views about the scope of conceptual engineering. In section 5 I offer an understanding of the implementation challenge and how to solve it. I conclude briefly in section 6.

2. Kinds of Kinds

Philosophical questions arise in connection with phenomena of all kinds. But what makes those questions specifically philosophical is plausibly their focus on the necessary features rather than the contingent features of the phenomena in question. Thus, philosophers do not primarily ask what laws of nature obtain in the actual world, but instead ask what it is for something to be a law of nature (in any possible world). Philosophers do not primarily ask what causal relations obtain in the actual world, but instead ask what it is for a causal relation to obtain (in any possible world). And philosophers do not primarily ask whether clairvoyance is a reliable method of belief-formation in the actual world, but instead ask what the implications of such reliability would be for epistemic warrant (in any possible world). For ease of reference, I will call the necessary features of the phenomena of interest to philosophy as traditionally conceived 'philosophical kinds'.[1] In the examples provided, the relevant philosophical kinds are *lawhood*, *causation*, and *epistemic warrant*, respectively.

I do not claim that philosophical kinds can be defined in terms of their necessary features; nor do I claim that every necessary feature of a kind is of interest or relevance to philosophy. Rather, the point is that traditional philosophical kinds retain their features across the possible worlds in which they are instantiated. Laws of nature are not Humean regularities in one world and nomic necessitation relations in another. Causal relations are not a matter of property transfer in one world but counterfactual dependence in another. Epistemic warrant is not a matter of reliability of a certain kind in one world but a matter of internal coherence in another. Philosophical kinds are in this sense cross-world kinds. There are many philosophical kinds besides these, and they all have this feature of being cross-world kinds: truth, knowledge, naming, necessity, justice, moral goodness, identity, and so on. Arguably, this is what makes philosophical kinds open to traditional philosophical investigation rather than the subject of the empirical sciences. It is what makes conceptual analysis rather than empirical science an appropriate methodology for the investigation of philosophical phenomena, where conceptual analysis is understood not as the analysis of concepts, but as the investigation of the necessary features of phenomena by conceptual rather than empirical means. The claim that conceptual analysis is not the analysis of concepts may sound controversial, but reflection on examples indicates that it is the referent of the concept and not the concept itself that is of fundamental interest. What philosophers want to know about are

---

[1] One might instead talk of essences, but this complicates the dialectic, so I stick to talk of necessary features throughout.

laws of nature, causation, and epistemic warrant, not merely our concepts of these things (although concepts are, of course, of philosophical interest in their own right).

Of course, philosophers often draw on real-world examples when constructing philosophical theories about philosophical kinds. And answering philosophical questions typically requires empirical knowledge. Indeed, answering some philosophical questions, such as those raised by quantum physics and psychiatry, for example, clearly requires a good deal of specialised empirical knowledge. Empirical science itself gives rise to philosophical questions, and philosophical theories are designed to accommodate the empirical data. But philosophical theories and questions are often about the necessary features of the relevant phenomena nonetheless.

The claim that philosophical kinds are cross-world kinds is also consistent with the recent program of experimental philosophy, which aims to test intuitions about traditional philosophical phenomena through experimental means.[2] One strand of experimental philosophy—negative experimental philosophy—aims to cast doubt on the use of intuitions in answering traditional philosophical questions; a second strand of experimental philosophy—positive experimental philosophy—aims to make progress on traditional philosophical questions by distinguishing trustworthy intuitions from untrustworthy ones. Neither strand undermines the nature of traditional philosophical kinds as cross-world kinds. Indeed, in so far as experimental philosophy is interested in our actual intuitions rather than in the necessary features of phenomena, it is better understood not as philosophy *per se*, but as an empirical investigation into whether traditional philosophy should rely on intuitions, and hence as an empirical investigation into whether traditional philosophy can yield knowledge of the necessary features of phenomena. Finally, it is enough for my purposes that there are some philosophical kinds as I have defined them, and for the remainder of the paper 'philosophical kinds' will be used specifically to refer to cross-world kinds in the sense explained.

Philosophical kinds differ in this respect from both the kinds of interest to the natural sciences and the kinds of interest to the social sciences. I will call the former 'natural kinds', and the latter 'social kinds'. Each example of a philosophical question given above was

---

[2] For an overview of experimental philosophy see for example Knobe & Nichols (2017).

explicitly contrasted with the kind of question that is paradigmatically of interest to the sciences: questions about what laws of nature obtain in the actual world, what causal relations obtain in the actual word, and what methods of belief-formation are reliable in the actual world. This is not to say that the natural and social sciences are not interested in cross-world kinds. Indeed, the related distinction between essential and accidental properties is clearly of importance to the natural and social sciences, since the ability to explain and predict the observable data, as well as the ability to manipulate outcomes for our own purposes, whether that's to keep time at sea, to generate renewable energy, or to reduce poverty, depends in part on knowledge of the difference.[3] Counterfactual reasoning is clearly an important feature of empirical science, and as noted above, there is philosophy in science. But the natural and social sciences are not interested in cross-world kinds *per se.* Rather, the fact that the sciences have as their goal the explanation and prediction of observable data in the actual world, and that such explanation and prediction often depend on knowledge of hidden structures and causal mechanisms that are revealed through further empirical investigation, means that scientific investigation of the same surface-level observable data might reveal different natural and social kinds in different possible worlds. Cat-like creatures are not necessarily cats; watery liquid is not necessarily water; the poor are not necessarily lacking in moral worth.[4] But the sciences can remain neutral on whether or not this is so.

This difference between science and traditional philosophy is illustrated by Kripke's (1972) examples of a posteriori necessities. Science discovers, for example, that water is actually H_2O, but the claim that water is therefore *necessarily* H_2O is a philosophical claim, being an instance of the cross-world claim that natural kinds have certain of their properties essentially. That is, it is a philosophical claim that empirical investigation gives us knowledge of essences. Of course, one could reject the claim that water is necessarily H_2O, or the claim that there are necessary a posteriori truths, but this would also be to make a cross-world claim to the effect that natural kinds such as water *do not* have those properties essentially.

---

[3] See Godman, Mallozzi & Papineau (2020) for an argument to the effect that essential properties are super-explanatory.

[4] The first example is a reference to Putnam's robot cats (Putnam, 1962); the second a reference to Twin Earth (Putnam, 1970); the third a reference to the alleged view of the poor in ancient Greece.

But there is a feature of some social kinds that marks them out as different not only from philosophical kinds but also from natural kinds. I say 'some' social kinds rather than 'all' social kinds because it is natural to distinguish different kinds of social kinds. To take one example, Khalidi (2015) distinguishes three kinds of social kinds: those whose instances depend on our attitudes towards them (e.g. *permanent resident* and *prime minister*); those whose instances depend only on our attitudes towards some of them (e.g. *money* and *war*); and those whose instances are independent of our attitudes towards any of them (e.g. *racism* and *recession*). Insofar as this yields an appropriate subdivision, I will be focusing primarily on the first two kinds of social kinds—those that depend on some way on our attitudes towards them. Note also that the claim that certain kinds of social kinds have a feature that sets them apart from philosophical and natural kinds does not undermine the fact that both philosophical and natural kinds play an important role in the social sciences.

To see what distinguishes some social kinds from philosophical and natural kinds, I draw on a three-fold distinction due to Epstein (2015, 2016) between social facts, the grounds of those facts, and the anchors of those grounds. There is, of course, a vast literature on the nature of metaphysical grounding[5] and Epstein's framework is not uncontroversial. For example, there is a substantive question about whether what Epstein refers to as anchors really are distinct from the grounds of the social kinds in question, or whether it is better to understand them as parts of those grounds, or as grounds of those grounds.[6] But I set aside issues such as this, as well as issues that arise within the literature on metaphysical grounding more generally, because the distinctions between philosophical kinds, natural kinds, and social kinds that interests me are independent of the local issues in debates about the nature of metaphysical grounding. I draw on Epstein's three-fold distinction because it provides a useful way of illustrating the differences between kinds that I am interested in, but ultimately I can simply defer to whatever account of metaphysical grounding turns out to be correct. Consider, then, the following social fact:

(1) Al Capone was a criminal.

---

[5] For a general overview of the nature of metaphysical grounding see for example Bliss and Trogdon (2021).

[6] For arguments against Epstein's distinction between grounding and anchoring see Hawley (2019), Mikkola (2019), and Schaffer (2019), and for replies see Epstein (2019a, 2019b).

This social fact is grounded in facts such as the following:

(2) Al Capone ordered the murder of a number of people.
(3) Al Capone carried concealed weapons.
(4) Al Capone failed to pay his taxes.

It might be tempting to say instead that (1) is grounded in the fact that Al Capone broke the law, but this simply raises the further question of what grounds the fact that Al Capone broke the law, which brings us back to facts such as (2), (3) and (4). The grounding relation is an ontological relation and not a causal one. Each of (2), (3) and (4) provides an ontological explanation of (1) in the sense that having performed those acts is what makes it the case that Al Capone was a criminal. They are the acts in virtue of which Al Capone was a criminal. What caused Al Capone to be a criminal, in contrast, might have been greed, or peer pressure, or some other combination of factors, but none of (2), (3) or (4) provides a causal explanation for (1).

Specifying the relation between a social fact and the grounds of that fact leaves open, according to Epstein, the further question of what makes it the case that the grounding relation holds. What makes it the case, for example, that the facts specified in (2), (3) and (4) ground the social fact specified in (1)? Why is it that ordering murder, carrying concealed weapons, and failing to pay his taxes, were acts in virtue of which Al Capone was a criminal? To answer this question is to cite, according to Epstein, not the grounds of the social fact in question, but the anchors of those grounds. In this particular case, the grounding relation is anchored by facts such as that the relevant bills were proposed to the houses of Congress, that they were passed by appropriate vote, that they were signed by the President, that they were subsequently upheld by judicial interpretation and by the verdicts of juries, and so on. These facts are the anchors of the grounds of (1). They anchor the fact that (1) is grounded in (2), (3) and (4). (Again, I remain neutral on whether anchors are best understood as distinct from the grounds of the social kind in question, or as parts of those grounds, or as grounds of those grounds.)

The anchoring relation, just like the grounding relation, is an ontological relation and not a causal one. The fact that the bills were proposed, passed and signed provides an ontological explanation of the grounding relation. They are the facts in virtue of which the grounding

relation holds. What caused the grounding relation to hold, in contrast, might have been a desire to take back control of the city, or political corruption, or some other combination of factors, but the fact that the bills were proposed, passed and signed are not themselves causal explanations of the grounding relation.

Epstein's three-fold distinction between social facts, the grounds of those facts, and the anchors of those grounds, can be reframed to provide an account of social kinds. On this view, the grounds of a social kind can be thought of as a disjunctive set of properties, where the instantiation of any property in the set is sufficient for membership of the kind in question. Thus the grounds of the social kind *criminal* can be thought of as the disjunctive set of properties {ordering murder, carrying concealed weapons, failing to pay taxes, …, …}.

One particularly significant fact about the grounds of social kinds is that they can be changed, as can be seen by the fact that they vary from time to time and from place to place. For example, rough sleeping and begging became grounds for being a criminal in the UK with the introduction of the Vagrancy act of 1824, subsequently ceased to be grounds for being a criminal in Scotland in 1982, and, following recent votes in favour of repealing the act in both the House of Lords and the House of Commons, will soon cease to be grounds for being a criminal anywhere in the UK. Similarly, homosexuality is not currently grounds for being a criminal in the UK, but it was until relatively recently, and it still remains grounds for being a criminal in a large number of countries across the world. The grounds of the social kind *criminal*, then, are context-dependent and subject to change.[7]

The anchors for the grounds of social kinds are also subject to change across time. We can see this by reflecting on the fact that the same social kind can have the same grounds in two communities and yet the explanation for the grounding relation be different in each case. This is because different processes can serve to establish the relevant grounding relations in different communities. For example, the processes required to change the grounds of the social kind *criminal* in the US are currently different from the processes required to change the grounds of the social kind *criminal* in the UK, and both of these are different again from

---

[7] Other examples include legislative changes relating to the use of terms such as 'cheddar', 'chocolate', and 'champagne'. In this sense, *cheddar, chocolate* and *champagne* are also social kinds anchored by convention.

the processes required to change the grounds of the social kind *criminal* in a dictatorship where rule by decree is the norm. The anchors of the grounds of these kinds of social kinds are all ultimately matters of convention, and the changeability of the grounds of social kinds is itself to be explained by the conventional nature of the anchoring of those grounds.

Philosophical kinds can also have grounds. For example, according to the nomic necessitation theory, the fact that it is a law of nature that all *F*s are *G* is grounded in the fact that the property *F* and the property *G* stand in the relation of nomic necessitation, and according to the counterfactual theory of causation, the fact that striking the match caused the explosion is grounded in the fact that the explosion would not have occurred had the match not been struck. The theories themselves are controversial, but the fact that at least some philosophical kinds have grounds is not. Similarly, natural kinds can have grounds. For example, the fact that an atom is a hydrogen atom is grounded in the fact that it has one proton in its nucleus, and the fact that sodium chloride is a salt is grounded in the fact that it consists of a cation (the positive ion of a base) and an anion (the negative ion of an acid). We can think of irreducible philosophical and natural kinds as a limiting case and say that they are grounded in themselves, or grounded in identities. But the grounds of philosophical kinds and natural kinds are fundamentally different from the grounds of social kinds such as *being a criminal*.

One crucial difference is that, in contrast to the grounds of social kinds, the grounds of philosophical kinds and natural kinds are *not* subject to change. The grounds of philosophical kinds are fixed in virtue of the fact that philosophical kinds are cross-world kinds in the specific sense articulated above. Their grounds do not vary from world to world, let alone from time to time. There is nothing we can do to change them. The grounds of natural kinds are also fixed. All hydrogen atoms, no matter where and when located, have a single proton in their nucleus, since that is the property in virtue of which an atom is a hydrogen atom; and all salts, no matter where and when located, consist of a cation and an anion, since that is the property in virtue of which a compound is a salt. Again, there is nothing we can do to change them.

Since the grounds of these philosophical and natural kinds are unchangeable, one might say that they have no anchors. Alternatively, and perhaps more plausibly, we can say that the grounding relations in each case are anchored by identity relations. Thus the fact that it is a law of nature that all *F*s are *G* is grounded in the fact that the property *F* and the property *G*

stand in the relation of nomic necessitation because being a law of nature is *identical* to being a nomic necessitation relation; and the fact that an atom is a hydrogen atom is grounded in the fact that it has one proton in its nucleus because being a hydrogen atom is *identical* to having a single proton in the nucleus.

It is important to note that a change in the grounds (or anchors) of a social kind over time does not thereby undermine its identity over time or its status as a kind. As Millikan (1999, 2000, 2017) and Bach (2016, forthcoming) point out, social kinds are real kinds that support explanations and predictions in the social sciences. The realist outlook I mentioned at the beginning extends, then, not only to philosophical kinds and natural kinds, but also to social kinds with changeable grounds. Moreover, an accurate description of the causal structures of social reality requires same-tracking of social kinds through change. This is not to say, of course, that a social kind can maintain its identity through *any* change in its grounds, no matter how extensive. But this point follows from the fact that social kinds are real kinds. Only those changes that preserve theoretically significant explanations and predictions across members of the 'kind' before the change and members of the 'kind' after the change will be changes to a single kind. The fact that social kinds are real kinds that can undergo changes to their grounds is also consistent with the existence of borderline cases and vagueness.

According to Millikan (1999, 2000), social kinds are 'historical kinds', where likenesses between members of a single historical kind are explained by historically-situated reproductive processes. This understanding fits nicely with the claim that social kinds can maintain their identity through changes to their grounds, because historically-situated reproductive processes will not guarantee identity of grounds across spatially and temporally distant members of the relevant kind. Millikan contrasts historical kinds with 'eternal kinds', where likenesses between members of a single eternal kind are explained by the natural laws that govern their internal properties. This understanding fits nicely with the claim that the kinds of the natural sciences have grounds that are fixed. Given that likenesses between the members of a single philosophical kind are explained by the fixed nature of the grounds of the relevant kind, philosophical kinds are, in this respect, similar to eternal kinds.

The phenomenon of changeable grounds is not to be confused with the phenomenon of multiple realizability, which is distinct. There are examples of multiply realizable philosophical and natural kinds, but multiple realizability is consistent with the grounds of the

relevant kinds being fixed and unchangeable. For example, according to Putnam (1967), *pain* is a higher-order, multiply realizable kind whose realizers are members of the disjunctive set of properties: {being in brain state *b*, being in octopus state *o*, being in computer state *c*, …}. Similarly, the natural kind *noble gas* is a higher-order, multiply realizable kind, whose realizers are members of the disjunctive set of properties: {being helium, being argon, being krypton, …, …}. But sets of realizers for higher-order philosophical and natural kinds are not changeable in the way that the grounds of certain social kinds are. We cannot make it the case that a particular brain state does not realize pain, and we cannot make it the case that helium is not a noble gas. Nor can we make it the case that something that is not currently a realizer of pain comes to be a realizer of pain, or that what is currently not a noble gas comes to be one. We can, perhaps, build an electronic machine that can feel pain, or discover a new noble gas, but these would constitute discoveries that the sets of realizers for the higher-order properties were larger than we thought; they would not constitute a change to the sets of realizers for those properties. In contrast, we can make it the case that carrying a concealed weapon is no longer a crime, and we can make it the case that coercive behaviour is. These are not a matter of discovery but a matter of deliberate and targeted change.

In this section I have argued that there are real differences between philosophical kinds, natural kinds, and social kinds. The difference that will be of most relevance to the remainder of the paper is specifically the difference between on the one hand those kinds whose grounds are fixed (i.e philosophical kinds and natural kinds), and on the other hand those kinds whose grounds are changeable (i.e. certain kinds of social kinds). In the next section I appeal to this specific difference in order to distinguish two different ways in which conceptual engineering might be thought of as normative.

3. Normativity in Conceptual Engineering

Conceptual engineering is standardly characterised as 'the project of assessing and improving our representational devices' (Cappelen, 2018: 1), where this is understood as being a distinctively normative rather than a merely descriptive project. Echoes of both the characterisation and its alleged normative import reverberate through the literature, where the contrast between the normative and the descriptive is often used to distinguish conceptual engineering from conceptual analysis, whether explicitly or, more often, implicitly. An explicit statement is given by Nado, who says 'While conceptual analysts are interested in the

concepts we *do* have, conceptual engineers are interested in the concepts we *ought* to have. Their project is prescriptive rather than descriptive' (Nado, 2021: 1509, original emphasis). Similarly, according to Eklund, '... while philosophers have often been concerned with our actual concepts … , philosophers should also be asking themselves whether these really are the best tools for understanding the relevant aspects of reality, and in many cases consider what preferable replacements might be. Philosophers should be engaged in *conceptual engineering*' (Eklund, 2014: 293, original emphasis).[8]

But there are two different ways in which we could improve our representational devices, and these are not always clearly distinguished. The first is improvement that involves the elimination of an old representational device or the introduction of a new one; the second is improvement that involves the revision of a representational device that is already in use. I will refer to the former as 'improvement by elimination/introduction', noting that improvements of this kind could involve either the elimination of a representational device or the introduction of a representational device and need not involve both; I will refer to the latter as 'improvement by revision', noting that improvements of this kind involve the revision of a specific representational device that is already in use and that revisions to a theoretical framework in virtue of the elimination of an old representational device or the introduction of a new representational device do not count as improvements by revision in this sense.[9]

---

[8] Machery (2017) provides an exception to the claim that conceptual analysis is purely descriptive, saying: ' … in practice conceptual analysis is often partly descriptive and partly prescriptive, and it is also often difficult to identify which component of an analysis is descriptive, and which is prescriptive: Analysis often describes the content of a concept and regiments it in one stroke' (217). However, he goes on to say: 'Still, it remains useful to distinguish prescriptive and descriptive analyses, viewed as ideal types' (217), and this distinction is plausibly what conceptual engineers have in mind when they distinguish certain forms of conceptual engineering from traditional conceptual analysis. Note that it is consistent with the distinction between descriptive and normative projects that the former serve as precursors to the latter. See for example Bach (2019), Burgess & Plunkett (2020), Glock (forthcoming), and Jackson (forthcoming). There is also, as I will argue, a kind of normativity built into descriptive projects themselves.

[9] Sawyer (2020b: 1001) introduces a distinction between conceptual engineering in the broad sense, which involves the introduction, elimination or revision of a term, and conceptual engineering in the narrow sense, which focuses specifically on revision. The distinction I draw here effectively divides conceptual engineering in the broad sense into two subcategories: conceptual engineering in the sense of elimination or introduction, and conceptual engineering in the narrow sense of revision.

Of course, in practice we may want to do a bit of both—eliminating old representational devices or introducing new ones, while also keeping but revising others. But the difference between these two ways of improving our representational devices has implications for how we think of the normativity involved in conceptual engineering, as well as its scope. As I will argue, improvement by elimination/introduction is a part of all descriptive projects, including traditional philosophical projects and scientific projects, and hence conceptual engineering, insofar as it involves improvement by elimination/introduction alone, cannot be thought of as *distinctively* normative or as *different in kind* from descriptive science and conceptual analysis, although it can be thought of as having a very broad scope. On the other hand, improvement by revision can be distinguished from purely descriptive projects, and hence conceptual engineering, insofar as it involves improvement by revision, *can* be thought of as distinctively normative and as different in kind from descriptive science and conceptual analysis, but, as I will argue, conceptual engineering in this sense is only possible for representational devices that refer to kinds with changeable grounds, and hence conceptual engineering in this sense has a relatively narrow, albeit enormously significant, scope.[10]

In this section I will explain first, how improvement by elimination/introduction is a part of all descriptive projects, noting that such projects involve the assessment of our representational devices and are governed by the normative question of whether our representational devices are adequate for the explanation and prediction of the relevant data; and second, why improvement by revision is only possible for representational devices that refer to kinds with changeable grounds, noting that, as I will clarify later, such projects involve, in addition, the assessment of the *objects* to which our representational devices refer and are governed by the normative question of whether those objects are adequate for certain *further* purposes such as achieving social justice. The difference between improvement by elimination/introduction and improvement by revision, based as it is on the distinction between kinds whose grounds are fixed and kinds whose grounds are changeable, also has implications for the scope of conceptual engineering, which I mention in this section and clarify further in the next.

---

[10] Precisification is a complex matter, and I do not have the space to discuss it in full here, but for now it is worth noting that some cases of precisification involve improvement by elimination/introduction, while others are better understood as cases of improvement by revision.

Consider the natural sciences, which might be thought to be quintessentially descriptive. The natural sciences have as an aim the explanation and prediction of observable data in the natural realm. Progress towards that aim is accomplished by introducing and developing theoretical (i.e linguistic and conceptual) frameworks. Some theoretical frameworks yield better predictions of the observable data than others, and hence some theoretical frameworks provide better explanations of the observable data than others. Thus the geocentric theory was replaced by the heliocentric theory, phlogiston theory was replaced by the theory of oxidation, caloric theory was replaced by thermodynamics, and so on. In each of these cases, the representational devices (the words and concepts) originally in use were increasingly recognised as inadequate for the purposes of explaining and predicting the observable data and were subsequently replaced by representational devices (words and concepts) that were thought to be an improvement in this regard. In each of these cases, the move from the earlier theoretical framework to the later theoretical framework was governed by the norm that our representational devices should be adequate for prediction and explanation, and thereby involved the assessment of the relevant representational devices against this norm of predictive and explanatory adequacy. In this sense, these are examples of descriptive projects with an essentially normative dimension in which our representational devices are assessed and improved. Natural scientists ask themselves, albeit indirectly, whether our actual representational devices 'really are the best tools for understanding the relevant aspects of reality, and in many cases consider what preferable replacements might be', and natural scientists are interested, again albeit indirectly, in 'the concepts we *ought* to have'.[11]

Since the social sciences also have as an aim the explanation and prediction of observable data, this time relating to social phenomena, similar considerations apply. Progress towards that aim is once again accomplished by introducing and developing theoretical (i.e. linguistic and conceptual) frameworks, where some theoretical frameworks provide better explanations and yield better predictions than others. Key developments in the social sciences have seen the elimination of theoretical terms such as 'maternal imprinting' and 'racial hierarchy', and the introduction of theoretical terms such as 'recession', 'sustainability', 'indirect

---

[11] It might be thought that these kinds of representational changes in science are not cases of conceptual engineering because they are the by-product of general paradigm shifts rather than intentional, targeted revision of concepts. But the claim that conceptual engineering necessarily involves intentional, targeted revision goes against a widespread view, articulated by Thomasson, following Plunkett and Sundell (2013), that conceptual engineering is something 'we have been doing all along without being aware of it' (Thomasson, 2021: 2).

discrimination', 'oppression', 'the working poor', 'the patriarchy', 'misogyny', and 'post-traumatic stress disorder'. These theoretical changes were also governed by the norm that our representational devices should be adequate for prediction and explanation, and thereby also involved the assessment of the relevant representational devices against this norm of predictive and explanatory adequacy. In this sense, these kinds of examples of theoretical progress in social science are also examples of descriptive projects with an essentially normative dimension in which our representational devices are assessed and improved. Social scientists also ask themselves, then, again albeit indirectly, whether our actual representational devices 'really are the best tools for understanding the relevant aspects of reality, and in many cases consider what preferable replacements might be', and natural scientists are interested, again albeit indirectly, in 'the concepts we *ought* to have'.

Similar considerations also apply in philosophy, where explanation rather than prediction is a primary reason for the development of theoretical frameworks. Key developments in philosophy have seen the introduction of theoretical terms such as 'rigid designator', 'narrow content', 'external reason', 'grounding', 'supervenience', 'epistemic warrant', 'undermining defeater', and so on. These changes were governed by the norm that our representational devices should be adequate for explanation, and thereby involved the assessment of the relevant representational devices against this norm of explanatory adequacy. In this sense, here too, these kinds of examples of theoretical progress in philosophy are examples of descriptive projects with an essentially normative dimension in which our representational devices are assessed and improved.

Improvement by elimination/introduction, then, is an essentially normative dimension of descriptive projects and hence conceptual engineering in this sense cannot be thought of as *distinctively* normative or as *different in kind* from descriptive science or conceptual analysis, although it can be thought of as having a very broad scope. Interestingly, the quotations from Nado and Eklund at the beginning of this section are most straightforwardly read as advocating improvement by elimination/introduction rather than improvement by revision.

What of improvement by revision? This fits more closely with Cappelen's characterisation of amelioration according to which it 'always involves *the extension and intension of a predicate changing over time*' (Cappelen 2018: 62, original emphasis). This characterisation focuses on a change of extension and intension to a term that is already in use within a

theoretical framework and it gives rise to the specific normative question: what ought the intensions and extensions of our terms be? This is the question that drives much of the literature, where it is widely assumed to be a general question applicable to any term. But it is a much more specific question than the general question of what our representational devices ought to be, and has a much narrower target than is often supposed. Assuming that *ought* implies *can*, the question is applicable only to those predicates whose extensions and intensions we can change. After all, if improvement by revision is to be genuinely distinct from improvement by elimination/introduction, then a change to the extension and intension of a term should not be equivalent to the elimination of the term and the introduction of a new term that sounds the same. Something has to provide continuity through the change in extension and intension of the relevant term that makes it, in an appropriate sense, *the same term*.

So what is it to change the extension and intension of a predicate in order to improve it without replacing it? Consider Cappelen's own elaboration. He says:

> 'One way to think of this is: the conditions that need to be satisfied in order to fall into the extension of 'person' (or 'woman' or 'marriage' or 'belief') have been changed by conceptual engineering, and as a result of that the intension and extension of this term has changed. ... Things that were not in the extension pre-amelioration go in, and some of what was in goes out ...' (Cappelen 2018: 62).

The key claim here is that the intension and extension of a term is changed *as a result of changing the conditions that need to be satisfied in order to fall into the extension of the term.*[12] But what is it to change the conditions that need to be satisfied in order to fall into the extension of a term? The account of kinds with changeable grounds discussed in the previous section provides a model answer. Consider the following example. When homelessness was decriminalised in Scotland in 1982, what changed were the grounds of the social kind *criminal*. Being homeless was no longer included in the grounds of that social kind. Changing

---

[12] This is an expression of Cappelen's 'worldliness' account of conceptual engineering as articulated in his (2018: chapter 12). Cappelen's account is insightful and entirely appropriate when applied to social kinds, but it is not applicable to philosophical kinds or natural kinds. The current paper explains why, but see also Sawyer (2020a, 2020b).

the grounds of the social kind *criminal* changed the conditions that needed to be satisfied in order to fall into the extension of the term 'criminal' because changing the grounds of the social kind *criminal* is *just what it is* to change the conditions that need to be satisfied in order to fall into the extension of the term 'criminal'. Thus the intension and extension of the term 'criminal' changed as a result of changing the grounds of the relevant kind.

Improvement by revision, then, can be understood as involving a change to the grounds of a kind, where this brings about a change to the extension and intension of the relevant kind term. However, improvement by revision understood in this way entails that it is only terms that refer to kinds with changeable grounds that can be improved without being replaced. And since neither philosophical kinds nor natural kinds have changeable grounds, we cannot bring about a change in their grounds, and hence cannot, on this understanding, change the extension and intension of either philosophical kind terms or of natural kinds. But without an alternative model of how to change the extension and intension of a term without that amounting to the elimination of the term in question and the introduction of a homonym, there is no reason to think improvement by revision is possible for anything but social kind terms.

This result may sound surprising. However, alleged examples of changes to the extension and intension of philosophical or natural kind terms are not obviously best understood in this way. On reflection, some alleged examples are plausibly cases of improvement by elimination/introduction, and some arguably turn out not to be philosophical or natural kind terms after all. For example, the decision to redefine measurement terms such as 'kilogram' and 'meter' with reference to mathematical constants rather than physical objects does not constitute a decision to redefine natural kind terms because measurement terms are conventional terms that allow us to map objective relations between objective properties of objects. Other alleged examples look to involve epistemic change rather than semantic change. For example, Sawyer (2020b, 2021) argues that many of the alleged examples of conceptual engineering do not involve concept revision but instead involve revisions to our conceptions, thus providing an updated understanding of what were already the referents of our terms.[13] A particularly controversial example here is the example of the widely-cited IAU 2006's resolution on the definition of 'planet'. This is typically understood as an example of a

---

[13] See also Sawyer (2018).

change to the extension and intension of a natural kind term, but can also, instead, be understood as a discovery, despite the fact that the definition was proposed and adopted by vote.[14] Note that what changed in 2006 was our epistemic state, and it was this that led to the new proposed definition. Those who voted were experts in the field and had their eye on what classification system would work best theoretically. And they recognise that their definition may have to change when their astronomical knowledge increases. The case deserves attention in its own right, but for now I merely note that it is far from obvious that the term 'planet' provides an example of a improvement by revision to a natural kind term.

In this section I have argued that there are two different ways in which we might improve our representational devices. The first is improvement by elimination/introduction and is a part of descriptive projects generally, where descriptive projects involve the assessment of our representational devices and are governed by the normative question of whether our representational devices are adequate for the explanation and prediction of the relevant data. The second is improvement by revision, which, I have claimed, is only possible for representational devices that refer to kinds with changeable grounds. This kind of improvement involves, as I will clarify later, the assessment of the *objects* to which our representational devices refer and the normative question of whether those objects are adequate for certain *further* purposes such as achieving social justice. The difference between improvement by elimination/introduction and improvement by revision, based as it is on the distinction between kinds whose grounds are fixed and kinds whose grounds are changeable, also has implications for the scope of conceptual engineering, which is the topic of the next section.

4. The Scope of Conceptual Engineering

Conceptual engineering is often portrayed by its advocates as having a maximally broad scope. This can be seen by the wide variety of examples in the literature, which include terms such as 'belief', 'knowledge', 'truth', 'causation', 'moral goodness', 'freedom', 'gene', 'foetus', 'person', 'woman', 'race', 'family', 'money' and 'marriage'. Included in these examples are terms that refer to philosophical kinds, terms that refer to natural kinds, and

---

[14] For an account of how the decision was reached, see:

https://www.iau.org/static/publications/ga_newspapers/20060812.pdf

terms that refer to social kinds. The broad scope view is stated explicitly by Cappelen, who says '…conceptual engineering is important for all parts of philosophy (and, more generally, all inquiry)' (Cappelen, 2018: 1). It is echoed by others, such as Thomasson, who says conceptual engineering 'has applications all over philosophy … [and] also … outside of philosophy … in medicine, psychology, biology, physics, government and the law' (Thomasson 2021: 1-2). But this broad scope view depends on understanding conceptual engineering as the project of assessing and improving our representational devices in the general sense that groups together the two distinct ways in which we might improve our representational devices: improvement by elimination/introduction, and improvement by revision. As such, it fails to mark out conceptual engineering as distinct from other forms of theoretical work in which we aim to describe various aspects of reality, whether philosophical, natural, or social. The broad scope view is endemic in the literature. Nado (2021) provides another particularly clear example.

I do not wish to legislate on terminology. My point is, rather, that claims that are true only of conceptual engineering in the narrow sense of amelioration should not be made of conceptual engineering in the broad sense, and that a clearer understanding of the difference between kinds of kinds would help in this respect, both amongst its advocates and amongst its detractors. For example, conceptual engineering is criticised by Deutsch (2020) for precisely this reason. According to Deutsch, alleged examples of conceptual engineering involve either the (stipulative) introduction of technical terms, such as 'supervenience', 'grounding' and 'epistemic warrant', or what he calls the 'revelatory' analysis of terms that are already in use, such as 'knowledge', 'justice' and 'free will'. But neither activity is new, according to Deutsch, and both are already undertaken in the context of conceptual analysis. To some extent I agree. As noted in section 3 above, the introduction of technical terms such as 'supervenience', 'grounding' and 'epistemic warrant' is a normal part of the development of descriptive theoretical frameworks adequate to explaining the data, and hence is not definitive of conceptual engineering. It takes place in the context of conceptual analysis, and hence in philosophy, but it also takes place both in natural science and in social science. And the analysis of terms such as 'knowledge', 'justice' and 'truth' is revelatory rather than revisionary because, drawing on the terminology introduced in sections 2 and 3 above, analysis of philosophical kind terms reveals the fixed grounds of the cross-world philosophical phenomena to which those terms already refer, and therefore does not involve a change to the intensions or extensions of the terms.

What Deutsch's criticism ignores, however, is the fact that only philosophical kinds can be the target of conceptual (revelatory) analysis, since it is only philosophical kinds that are cross-world kinds. Neither natural kinds nor social kinds have grounds that can be revealed by conceptual analysis. Instead, empirical inquiry is required. And given the unique nature of social kinds and their changeable grounds, we need, in addition to conceptual analysis and empirical inquiry, to acknowledge conceptual engineering in the narrow sense of improvement by revision. Kinds with changeable grounds that are anchored by convention give rise to the more specific normative question of whether those kinds are the way they ought to be. This is because in those cases, and in those cases alone, we have the capacity to change the grounds of the relevant kinds, thereby changing the intensions and extensions of the relevant terms. This is the sense in which improvement by revision involves the assessment of the *objects* to which our representational devices refer and raises the normative question of whether those objects are adequate for certain *further* purposes such as achieving social justice

This, then, leads us to a natural distinction between conceptual analysis, scientific theorising, and conceptual engineering in the narrow sense of amelioration. Conceptual analysis is the appropriate methodology for revealing the nature of cross-world kinds and is therefore the appropriate methodology for identifying and characterising philosophical phenomena. Scientific theorizing is the appropriate methodology for revealing the nature of natural and social kinds, where what is of interest is the development of a theoretical framework adequate for the explanation and prediction of the observable data. Conceptual engineering in the narrow sense of amelioration, in contrast, is the appropriate methodology for questioning the nature of the social kinds revealed by empirical inquiry and is set apart from both conceptual analysis and scientific theorizing by its unique connection to the more specific normative question of what the grounds of those kinds ought to be, and hence by its unique connection to the related question of what the intensions and extensions of our terms for those kinds ought to be. This is perhaps why the most convincing examples of successful projects in conceptual engineering come from the social realm; these examples, and these examples alone, cannot be reconstrued in purely epistemic terms, as revelatory discoveries about phenomena we have already identified. We do not *discover* that coercive behaviour is a crime, but make it so.

Of course, there are relations between the methodologies. For example, the normative, ameliorative project depends essentially on the scientific, descriptive project, as Bach (2019) has rightly emphasised. (See also footnote 9.) But the distinction between improvement by elimination/introduction and improvement by revision allows us also to distinguish conceptual engineering in the broad sense that covers all three distinct methodologies, and conceptual engineering in the narrow sense of amelioration, which is an important and significantly distinctive subcategory of conceptual engineering in the broad sense.

With the scope of conceptual engineering in the broad sense distinguished from the scope of conceptual engineering in the narrow sense, we are now in a position to look at the implications for the Implementation Challenge, which is the topic of the next section.

5. The Implementation Challenge

Conceptual engineering is generally thought of as involving a change to the intensions and extensions of our terms. This has led to widespread acceptance of the claim that the *direct*, or *immediate* goal of conceptual engineering, even if not its *ultimate* goal, is to bring about meaning change.[15] The distinction between the direct goal and the ultimate goal is important here, since, as Riggs says, '[c]onceptual engineers usually motivate their projects by claiming that changing meaning can change other things we care about, not by claiming that changing meaning is valuable for its own sake' (Riggs, 2019: 7).[16] But the claim that the direct goal of conceptual engineering is to change the meanings of our terms is widespread, and has given rise to the Implementation Challenge where this is understood specifically as the question: how exactly are we to bring about changes to the meanings of our terms? And understood as this question, the Implementation Challenge has come to seem particularly intractable. Cappelen, for one, expresses scepticism about our ability to change the meanings of our terms, claiming that the mechanisms for meaning change are inscrutable—that the mechanisms of meaning change are 'not known to us and might in effect be unknowable'

---

[15] Although Nimtz (2021) suggests that concepts can be engineered by engineering social norms, the social norms he cites are linguistic norms, and so he too is committed to the widespread view that the Implementation Challenge is primarily a semantic challenge.

[16] Isaac et al. (2022) characterise this distinction as a distinction between the *goal* of conceptual engineering (which I have called the *direct*, or *immediate* goal) and the *purpose* of conceptual engineering (which I have called the *ultimate* goal).

(Cappelen, 2018: 74)—and that meaning change is out of our control—that 'no individual or group has a significant degree of control over how meaning change happens' (Cappelen, 2018: 72). Echoing this scepticism, Deutsch says 'a speaker's stipulation that an existing term shall now have a different semantics doesn't suffice to *give* the term a new semantics. … we are, all of us, including would-be conceptual engineers, ignorant of what more must be done, or what else must occur, in order to bring about semantic changes' (Deutsch 2020: 6, original emphasis). The problem might seem particularly intractable on an externalist theory of meaning, where, as Burgess and Plunkett put it, 'our social and natural environments serve as heavy anchors … for the interpretation of our individual …talk' (Burgess & Plunkett 2013: 1096)[17], but since even the internalist accepts a notion of semantic meaning over and above individual speaker meaning, the problem is in fact a general one, as has been widely recognised.[18]

Responses to the Implementation Challenge have been concessionary, acknowledging that meaning change is to a large extent inscrutable and out of our control, and suggesting ways in which some more limited form of meaning change might be accomplished nonetheless, perhaps in limited contexts, or perhaps indirectly. Thus Pinder (2021) suggests that conceptual engineering should be understood as aiming to change speaker meaning rather than semantic meaning, Jorem (2021) agrees that conceptual engineering will have to focus on speaker meaning rather than semantic meaning and goes on to suggest ways in which a speaker might encourage others to adopt her way of speaking, and Koch (2021) claims that although we don't have any immediate mechanisms for bringing about meaning change we can nonetheless have 'collective long-range control' over semantic meaning. Thomasson (2021), meanwhile, suggests looking into the mechanisms of semantic meaning change as it occurs naturally so as to gain some insight into how we might try to bring it about artificially.

But once we see that the more specific normative question of what the intensions and extensions of our terms ought to be is relevant only to conceptual engineering in the narrow

---

[17] See also Riggs (2019). The implications of the claim that externalist theories of concepts preclude conceptual change are explored in the current paper and also in Sawyer (2021).

[18] The point is explicitly acknowledged, for example, by Cappelen (2018: 82) and Jorem (2021: 194). See also Pollock (2019), who defends a version of internalism but accepts the widespread view that the Implementation Challenge is primarily a semantic challenge and that internalism does not provide a ready answer to it.

sense of amelioration, and that it is inherently connected with the question of what the grounds of our kinds ought to be, we can see that the Implementation Challenge is a semantic challenge only in a derivative sense. It is at root the practical challenge of how to change the grounds of our social kinds. If we succeed in changing the grounds of our social kinds, the change to the intensions and extensions of our social kind terms will come about for free. This is because, as noted earlier, to change the grounds of a social kind is to change the conditions that need to be satisfied in order to fall into the extension of the relevant social kind term. For social kind terms, then, unlike philosophical kind terms and natural kind terms, we can effect semantic change, and doing so requires neither an understanding of the mechanisms of meaning change nor the ability to control people's linguistic behaviour. Semantic change is the result of social change, not its necessary precursor.[19]

Reconceiving the Implementation Challenge in this way thus offers some release from the supposed intractability that arises from seeing it primarily as a semantic challenge. There are three points to note. First, reconceiving the Implementation Challenge in this way is consistent with the claim articulated and defended by Sterken (2020) that linguistic interventions and communicative disruptions can serve as important catalysts for social change. More generally, it is consistent with the claim that a change in speaker meaning can serve as the starting point for the implementation of conceptual engineering. But neither of these is a necessary condition on the implementation of conceptual engineering, and implementation will not have been achieved until the grounds of the relevant social kinds have changed and the intensions and extensions of the relevant terms have changed as a result. Second, changing the grounds of our social kinds is still a challenge. But it is not a new challenge, and it is not one that is raised specifically by conceptual engineering. It is the perennial challenge of how to bring about social change by changing community practices, traditions and laws. Finally, the power to change the grounds of social kinds is not equally

---

[19] I am assuming for present purposes, as is common throughout the conceptual engineering literature, that a change to the extension of a term will also thereby involve a change to the intension of that term, and hence involve a change to the meaning of that term. However, given the stability of social kinds and our reference to them there is reason to take seriously the idea that the meaning of social kind terms might also be stable through changes to the grounds of the kinds to which they refer. I leave elaboration of this point for another occasion. But it is an irony of the literature that concepts are widely rejected as being obscure and unexplained, while meanings are routinely assumed to be clear and in need of no explanation.

available to all.[20] Social policies on education, healthcare, taxation, welfare, voting eligibility, public debate, free speech, and much more, affect an individual or group's power to bring about a change in the grounds of social kinds. What this means is that changing the anchors of the grounds of social kinds is an important additional means by which to promote social justice, one which is fundamentally itself an exercise in changing the grounds of further social kinds.

Conceptual engineering at the implementation level is the engineering of social phenomena, but it is specifically the engineering of social phenomena driven by conceptual reflection on how those phenomena ought to be. This brings about a change to the extension of the relevant term while maintaining continuity of reference to the same social kind over time.[21]

6. Conclusion

In this paper, I have argued that understanding the distinctions between kinds of kinds—specifically the distinction between kinds whose grounds are fixed and kinds whose grounds are changeable—provides a clearer understanding of the normativity and scope of conceptual engineering and of what is involved in its implementation. The normative question that is relevant to conceptual engineering in the narrow sense of amelioration is the question of what the grounds of our kinds ought to be. This arises for social kinds alone, since social kinds alone have changeable grounds. The implementation of conceptual engineering in the narrow sense in turn requires the grounds of our social kinds to be changed, and this is what brings about a change to the extensions of our social kind terms. But the social kinds, being real kinds, can persist throughout the changes to their grounds, as can our reference to them.

References

---

[20] For ideological constraints on conceptual engineering, see Podosky (2018), Queloz & Bieber (2021) and Shields (2021). See also Podosky (forthcoming) and Kitsik (forthcoming).

[21] If the proper scope of conceptual engineering in the narrow sense is the social realm and social kinds can persist through changes in their grounds, then we have a solution to the problem of topic continuity for conceptual engineering: sameness of reference through a change to the intension and extension of a social term is a necessary pre-condition of conceptual engineering.

Bach, T. (2016). Social categories are natural kinds, not objective types (and why it matters politically). *Journal of Social Ontology,* 2 (2): 177-201.

Bach, T, (2019). Real kinds in real time: on responsible social modelling. *The Monist,* 102: 236-258.

Bach, T. (forthcoming). Same-tracking real kinds in the social sciences. *Synthese.*

Bliss, R. and Trogdon, K. (2021). Metaphysical grounding, *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2021/entries/grounding/>

Burgess, A. & Plunkett, D. (2013) Conceptual ethics I. *Philosophy Compass,* 8 (12): 1091-1101.

Burgess, A. & Plunkett, D. (2020) On the relation between conceptual engineering and conceptual ethics. *Ratio,* 33 (4): 281-294.

Cappelen, H. (2018). *Fixing Language*: *An Essay on Conceptual Engineering*. Oxford: Oxford University Press.

Deutsch, M. (2020). Speaker's reference, stipulation, and a dilemma for conceptual engineers. *Philosophical Studies,* 177 (12): 3935-3657.

Eklund, M. (2014). Replacing truth? In A. Burgess & B. Sherman (eds) *Metasemantics: New Essays on the Foundations of Meaning*: 293-310. Oxford: Oxford University Press.

Epstein, B. (2015). *The Ant Trap: Rebuilding the Foundations of the Social Sciences.* New York: Oxford University Press.

Epstein, B. (2016). A framework for social ontology. *Philosophy of the Social Sciences,* 46 (2): 147-167.

Epstein, B. (2019a). Replies to Hawley, Mikkola, and Hindriks. *Inquiry*, 62 (2): 230-246.

Epstein, B. (2019b). Anchoring versus grounding: reply to Schaffer. *Philosophy and Phenomenological Research,* 99 (3):768-781.

Glock, H.-J. (Forthcoming). Conceptual engineering and conceptual analysis: Historical and conceptual connections. In M. G. Isaac & K. Scharp (eds) *New Perspectives on Conceptual Engineering (Vol. 1).* Springer.

Hawley, K. (2019). Comments on Brian Epstein's *The Ant Trap*. *Inquiry*, 62 (2): 217-229.

Isaac, M. G., Koch, S. & Nefdt, R. (2022). Conceptual engineering: a road map to practice. *Philosophy Compass,* 17 (10).

Jackson, F. (Forthcoming). Conceptual engineering in context. In M. G. Isaac & K. Scharp (eds) *New Perspectives on Conceptual Engineering (Vol. 2).* Springer.

Jorem, S. (2021). Conceptual engineering and the implementation problem. *Inquiry*, 64 (1-2): 186-211.

Khalidi, M. A. (2015) Three kinds of social kinds. *Philosophy and Phenomenological Research,* 90 (1): 96-112.

Kitsik, E. (forthcoming). Epistemic paternalism via conceptual engineering. *Journal of the American Philosophical Association.*

Knobe, J. and Nichols, S. (2017). Experimental philosophy. *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2017/entries/experimental-philosophy/>.

Koch, S. (2021). The externalist challenge to conceptual engineering. *Synthese,* 198: 327-348.

Kripke, S. (1972). *Naming and Necessity*. Oxford: Blackwell.

Godman, M., Mallozzi, A. & Papineau, D. (2020). Essential properties are super-explanatory: taming metaphysical modality. *Journal of the American Philosophical Association,* 6 (3): 316-334.

Machery, É. (2017). *Philosophy within its Proper Bounds.* Oxford: Oxford University Press.

Mikkola, M. (2019). Grounding and anchoring: on the structure of Epstein's social ontology. *Inquiry* 62 (2): 198-216.

Millikan, R. (1999). Historical kinds and the 'special sciences'. *Philosophical Studies,* 95 (1-2): 45-65.

Millikan, R. (2000). *On Clear and Confused Ideas: An Essay about Substance Concepts.* Cambridge: Cambridge University Press.

Millikan, R. (2017). *Beyond Concepts: Unicepts, Language and Natural Information.* Oxford: Oxford University Press.

Nado, J. (2021). Conceptual engineering, truth, and efficacy. *Synthese,* 198: 1507-1527.

Nimtz, C. (2021). Engineering concepts by engineering social norms: solving the implementation challenge. *Inquiry*. DOI: 10.1080/0020174X.2021.1956368

Pinder, M. (2021). Conceptual engineering, metasemantic externalism and speaker-meaning. *Mind* 130 (517): 141–163.

Plunkett, D, and Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosopher's Imprint*, 13 (23): 1–37.

Podosky, P-M. C. (2018). Ideology and normativity: constraints on conceptual engineering. *Inquiry*. DOI:10.1080/0020174X.2018.1562374

Podosky, P-M. C. (forthcoming). Can conceptual engineering actually promote social justice? *Synthese*.

Pollock, J. (2019). Conceptual engineering and semantic deference. *Studia Philosophica Estonia,* 12: 81-98.

Putnam, H. (1962). It ain't necessarily so. *Journal of Philosophy,* 59: 658-671.

Putnam, H. (1967). Psychological predicates. In W.H. Capitan and D.D. Merrill (eds) *Art, Mind, and Religion*: 37-48. Pittsburgh: University of Pittsburgh Press.

Putnam, H. (1970). Meaning and reference. *Journal of Philosophy,* 70: 699-711.

Queloz, M., & Bieber, M. (2021). Conceptual engineering and the politics of implementation. *Pacific Philosophical Quarterly*, online first. DOI: 10.1111/papq.12394

Riggs, J. (2019). Conceptual engineers shouldn't worry about semantic externalism. *Inquiry,* online first. DOI: 10.1080/0020174X.2019.1675534

Sawyer, S. (2018). The importance of concepts. *Proceedings of the Aristotelian Society,* 118 (2): 127-147.

Sawyer, S. (2020a). The role of concepts in fixing language. *Canadian Journal of Philosophy*, 50 (5): 555-565.

Sawyer, S. (2020b). Truth and objectivity in conceptual engineering. *Inquiry,* 63 (9-10):
1001-1022. DOI: 10.1080/0020174X.2020.1805708

Sawyer, S. (2021). Concept pluralism in conceptual engineering. *Inquiry.*
DOI: 10.1080/0020174X.2021.1986424

Schaffer, J. (2019). Anchoring as grounding: on Epstein's *The Ant Trap. Philosophy and
Phenomenological Research,* 99 (3):749-767.

Shields, M. (2021). Conceptual domination. *Synthese,* online first. DOI: 10.1007/s11229-021-
03454-4

Sterken, R. (2020). Linguistic intervention and transformative communicative disruptions. In
A. Burgess, H. Cappelen & D. Plunkett (eds) *Conceptual Engineering and
Conceptual Ethics*: 417-434. Oxford: Oxford University Press.

Thomasson, A. (2021). Conceptual engineering: when do we need it? How can we do it?
Inquiry. DOI: 10.1080/0020174X.2021.2000118