

Topic Continuity in Conceptual Engineering and Beyond

Tristram McPherson^a and David Plunkett^b

^aDepartment of Philosophy, The Ohio State University, Columbus, OH, USA; ^bDepartment of Philosophy, Dartmouth College, Hanover, NH, USA

ABSTRACT

One important activity in conceptual ethics and conceptual engineering involves proposing to associate a new semantics with an existing word. Many philosophers think that one important way to evaluate such a proposal concerns whether it preserves the “topic” picked out by the existing word, and several have offered competing proposals concerning what is required to preserve topic. Our paper is focused on the conceptual ethics question of how conceptual engineers should use the term ‘topic continuity’. We provide and defend a context-sensitive answer to this question. Our answer is motivated by the idea that there are several distinct considerations that we can reasonably care about (and which many conceptual engineers already do care about) in thinking about “topic continuity”, and, moreover, that how best to weigh them against each other can vary from context to context. On our proposal, ‘topic continuity’ can function as a useful representational device that enables coordination by inquirers with respect to these concerns. We conclude by locating our account in a broader way of thinking about topic continuity across a range of inquiries.

ARTICLE HISTORY Received 8 September 2021; Accepted 9 September 2021

KEYWORDS Conceptual ethics; conceptual engineering; topic; topic continuity; methodology

Introduction

Both individuals and communities can (and often do) change their views about a given topic. For example, over the course of his life, John Rawls changed his views about justice, and in the early decades of the twentieth century, physicists substantially changed their views about the nature of the atom. Such cases contrast with cases where there is a *change in topic*, rather than a change in view about the same topic. This would occur if, for example, Rawls switched from theorizing about justice to discussing his

CONTACT Tristram McPherson  dr.tristram@gmail.com  Department of Philosophy, The Ohio State University, Columbus, OH 43210-1132, USA

© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

favorite movies, or if a scientist switched from theorizing about atoms to studying the migratory patterns of birds.

While we have introduced the contrast between change in view and change in topic using relatively clear cases, there are also cases that are not so clear. For example, early in Plato's *Republic*, Socrates discusses the view that justice consists in treating friends well and enemies badly.¹ It can be natural to wonder whether this is really a (bizarre) view about the topic of *justice*, or a theory of *something else*. These sorts of observations motivate the thought that it would be useful to have a theory of topic change and topic continuity, and, in particular, one that illuminates (i) what topic continuity consists in, and/or (ii) how to assess whether there has been a change in topic in a given case.

In recent years, these kinds of issues about "topic continuity" have loomed large in philosophical discussions about "conceptual ethics" and "conceptual engineering".² Put roughly, conceptual ethics concerns certain normative and evaluative questions about thought and talk, such as questions about which concepts we should use, and why, and what we should mean by our words, and why.³ In turn, conceptual engineering (again, put roughly) incorporates such normative and evaluative inquiry into a project of introducing or reforming concepts (or other representational or inferential devices), and then trying to implement the use of those new or revised concepts.⁴ One important activity in conceptual ethics and conceptual engineering involves proposing to associate a new intension with an existing word. A common thought about such proposals is that, if implemented, they might well "change the topic" that the word is currently typically used to discuss. And it's not hard to see why. After all, if someone proposes that by 'justice' we should mean *shoes*, then there is a good case to be made that, if this change were implemented, people would have changed topics in a deep way: namely, they would now be talking about shoes, and not about justice.

The idea that certain proposals in conceptual ethics and conceptual engineering involve a change in topic has often been offered as the basis

¹*Republic* 332d (in (Cooper 1997)).

²In this paper, we use single quotation marks (e.g. 'bicycle') to mention linguistic items. We use double quotation marks (e.g. "bicycle") for a variety of tasks including quoting others' words, scare quotes, and mixes of use and mention. We use small caps (e.g. BICYCLE) to pick out concepts.

³Our use of 'conceptual ethics' here draws from (Burgess and Plunkett 2013a) and (Burgess and Plunkett 2013b).

⁴This way of thinking about the relation between conceptual engineering and conceptual ethics draws on (Burgess and Plunkett 2020) and (Cappelen and Plunkett 2020). Note that there is a range of important issues (discussed further in those papers) that we are glossing over in these rough sketches of what conceptual engineering and conceptual ethics are, and issues of how they relate to each other.

for objections to those proposals. The idea, in short, is that this sort of shift in topic constitutes an objectionable “break” that we should seek to avoid.

To get a rough feel for the idea here, consider P.F. Strawson’s objection to Rudolf Carnap’s method of “explicating” key philosophical concepts (which, roughly, consists of conceptual engineering to improve the precision and epistemic benefit of a representational system).⁵ Strawson argues that Carnap’s method fails to answer our original philosophical questions (which were articulated using our ordinary, non-explicated concepts). Instead, Carnap’s proposals in effect broach and answer *new* questions, formulated using the new, explicated concepts. Strawson’s worry, although directly about Carnap, can be generalized to all “reforming” proposals that attempt to associate a new character, intension, or extension with an existing word. In his recent book *Fixing Language*, Herman Cappelen calls this generalized challenge “Strawson’s Challenge”. As Cappelen writes, the core worry is that even if such revisionary proposals were successfully implemented, “they do not provide us with a better way to talk about what we were talking about; they simply change the topic”.⁶ He fleshes this worry out in a bit more detail:

The objection is that the answers employing terms with new extensions fail to answer the original questions. These answers concern something new – not what we were originally talking about when we used the [original expressions]. We have the illusion of an answer, but it’s a purely verbal illusion. There’s a lack of continuity of inquiry; the old questions are not being answered. We’re answering new questions.⁷

Within the recent literature on conceptual ethics and conceptual engineering, many philosophers endorse what one might think of as a *partially concessive* reply to Strawson’s Challenge, or associated worries. In short, this kind of reply grants the Strawsonian idea that the preservation of topic matters, and claims that some, but not all, proposals in conceptual ethics and conceptual engineering allow that to happen.⁸ Not surprisingly, different philosophers have quite different views about which proposals allow that to happen, and why. These differences regarding specific cases are often tied to more general disagreements about how to theorize about topic continuity.

⁵(Strawson 1963), criticizing the kind of method of “explication” found in (Carnap [1950] 1962) and (Carnap [1947] 1956).

⁶(Cappelen 2018, 100).

⁷(Cappelen 2018, 101–102).

⁸For some recent examples of this kind of reply, see (Cappelen 2018), (Prinz 2018), and (Thomasson 2020). See also (Haslanger 2000) for an approach that suggests this kind of reply to “Strawson’s Challenge”.

One kind of theorizing about topic continuity holds fixed the current meaning of 'topic continuity' and uses that term to ask what topic continuity is. We focus on a different issue: namely, the question of what we *should* mean by 'topic continuity'.

In response to this guiding question (which we take to be one in conceptual ethics), we defend a kind of context-sensitive answer. In short, we argue that we should use the term 'topic continuity' to pick out the presence (and degree of presence) of a range of different dimensions that we might describe as constituting a "break" in topic. These dimensions are ones that philosophers involved in debates over conceptual ethics and conceptual engineering are (in our view, *correctly*) sensitive to. They cover a range of different issues, including ones about smooth communication and others about preserving "what matters" to participants in a conversation. Using the term 'topic continuity' as a way to track these dimensions is, we argue, helpful for discussions about the kinds of conceptual engineering proposals that are our focus.

We break up our work in this paper into four main sections. In §1, we clarify and motivate the conceptual ethics question concerning theorizing about "topic continuity" that we focus on in this paper. In §2, we introduce both our positive proposal and the dimensions of a "break" in topic that we think matter, and defend the idea that we want a context-sensitive term to reflect the contextual variance of these dimensions' significance. In §3, we illustrate our proposal by applying it to a case study: Peter Railton's so-called "reforming" definition of the term 'good'. In §4, we then compare our contextualist proposal to some salient alternatives, each of which reflects an important strand of existing discussion about "topic continuity" within conceptual ethics and conceptual engineering. We conclude by briefly exploring the prospects for generalizing from our proposal to theorizing about topic continuity in contexts beyond those concerned with the specific kinds of conceptual engineering proposals that we focus on in this paper.

1. Our conceptual ethics question about "topic continuity"

As we explained in the introduction, this paper addresses a conceptual ethics question about "topic continuity". In this section, we more carefully describe this question. We begin by briefly contrasting it with the questions at the heart of two other sorts of inquiry about topic continuity (or our thought and talk about it). We then clarify and motivate the conceptual ethics question that is our focus.

To begin, let's contrast three sorts of inquiry where the term 'topic continuity' plays a central role.⁹ First, we might use this term, with its current meaning, when seeking to determine what topic continuity is, and/or what conditions entail its instantiation. Call this *first-order inquiry* concerning topic continuity. Second, we might seek to understand how thought and talk that uses the term 'topic continuity' – and what, if anything, that thought and talk is distinctively about – fits into reality.¹⁰ Here, we might investigate the semantics of the word, the metaphysics of the relation it picks out, and the epistemology of how (or if) we can come to know facts about that relation. By analogy with metaethics, call this *meta-level inquiry* concerning topic continuity. Third, we might ask what semantics we *ought* to pair with the word 'topic continuity', given its role in contemporary debates in conceptual ethics and conceptual engineering. Call this *conceptual ethics inquiry* concerning 'topic continuity'.

It is not always clear in the existing literature whether a given theory of "topic continuity" contributes to one or more of these three kinds of inquiry glossed above, nor, if it does, to which one it aims *primarily* to contribute. One possible reason for this lack of clarity is that some authors may be assuming that – given that 'topic continuity' is a quasi-technical term – there won't be that much of a gap between these inquiries, at least in terms of what they yield extensionally. Another possibility is that, on some views about discourse, people may sometimes be implicitly advocating for conceptual ethics views about a term simply by using the term (rather than by explicitly mentioning it).¹¹

Each of the three sorts of inquiry that we have distinguished here can be well-motivated. One motivation for "first-order" inquiry is that the sorts of concerns about conceptual engineering proposals that we gestured at in the introduction appear *repeatedly* (and often seemingly independently) across the literature.¹² This might suggest that there is a *pre-theoretic idea* that these discussions are all latching onto. Consider an example: when Strawson raises his challenge to Carnap, it might seem that one can grasp that there is *some* kind of issue here, which has to do with a potentially problematic *break* between those theses (or

⁹For detailed discussion of a closely related three-way distinction within epistemology, see (McPherson and Plunkett 2021), drawing on earlier work in (McPherson and Plunkett 2017) and (McPherson and Plunkett 2020).

¹⁰For a more careful discussion of a closely related gloss on metaethics, see (McPherson and Plunkett 2017).

¹¹This might happen, for example, by speakers engaging in "metalinguistic negotiation" as described in (Plunkett and Sundell 2013).

¹²For several examples, see (Cappelen 2018, 98).

questions, etc.) formulated with explicated, engineered concepts and those theses (or questions, etc.) formulated using our original, pre-theoretic concepts. Talk of “topic (dis-)continuity” can be understood as picking out this sort of pretheoretic idea of a “break”.¹³ First-order inquiry can be motivated by the conjecture that the current meaning of ‘topic continuity’ picks out the relation at the heart of this issue, and that it would be useful for the projects of conceptual ethics and conceptual engineering if we can develop a clear theory of this relation.

One motivation for meta-level inquiry about ‘topic continuity’ (besides whatever intrinsic interest it has) is that it could help us determine which of first-order or conceptual ethics inquiry it makes more sense to pursue in a given context. For example, meta-level inquiry might provide accounts of (e.g.) the nature of the conceptual role of ‘topic continuity’ or of the relation of topic continuity, which could help to *vindicate* interest in first-order inquiry about topic continuity, holding fixed the current meaning of ‘topic continuity’. On the other hand, meta-level inquiry might show that the existing semantics of ‘topic continuity’ is defective, or that it picks out a relation of no interest. This could help to motivate conceptual ethics inquiry about the term.

We think that conceptual ethics inquiry about ‘topic continuity’ can be motivated without such hypotheses. In recent years, philosophers have put forward a range of strikingly different proposals about the conditions for topic continuity (or proposals that we can fruitfully interpret as being about this issue). For example, some proposals tie topic continuity closely together with the preservation of the *intension* of relevant words.¹⁴ Other proposals emphasize “semantic function”.¹⁵ Others tie it closely to “same-saying” data that we can test empirically.¹⁶ We can distinguish the question of which (if any) of these conditions is currently picked out by ‘topic continuity’ from the question of which (if any) of these conditions it is *useful* for conceptual engineers to attend to in their theorizing. It might, of course, turn out that the word ‘topic continuity’ (as currently used) picks out something important to attend to in inquiry, and, in particular, something important for those working in (or about) conceptual engineering to attend to. But in the conceptual ethics spirit, we might simply aim to target whatever it is that is important here in the vicinity

¹³We take this to be one way of reading (Cappelen 2018).

¹⁴This is one way of reading the view put forward in (Schroeter and Schroeter 2020), drawing on their general metasemantic views in work such as (Schroeter and Schroeter 2014).

¹⁵For example, see (Haslanger 2000), as well as connected discussion in (Haslanger 2020). See also (Thomasson 2020).

¹⁶(Cappelen 2018).

more directly, by asking the following question: what semantics would it be useful to pair with the term ‘topic continuity’ to do theoretical work in evaluating proposals in conceptual engineering? This is the overarching question that we aim to make progress on in this paper.

In other words, we seek to engineer ‘topic continuity’ in a way that makes it *useful* for philosophers working in (or about) conceptual engineering. Somewhat more specifically, we focus on a central class of conceptual engineering projects: those that advocate for changing the semantic value (for example, the *character* or the *intension*) that is associated with an existing word, and which seek to actually implement that change.¹⁷ For brevity, we will call such projects *semantic change proposals*.¹⁸

Many of the most familiar examples of conceptual engineering are plausibly best read as semantic change proposals, from Carnapian explication to Sally Haslanger’s “ameliorative” accounts of race and gender terms.¹⁹ However, there are important sorts of conceptual engineering proposals that are not semantic change proposals. These include proposals that advocate for the introduction of new words, and proposals to abandon or eliminate the use of existing words.²⁰ While questions about topic (dis-)continuity can arise in evaluating all such proposals, they arise in an especially clear way in the context of semantic change proposals. For example, it is Carnap’s proposals to associate new intensions with existing terms that lead to the Strawsonian worries mentioned in the introduction. For this reason, we focus here on providing a useful

¹⁷It should be noted that many such proposals target changing the meaning of the word as used by certain people, in certain contexts, given certain aims, rather than (for example) trying to change the meaning of the term as used by everyone in all contexts. For further discussion of this feature of many proposals in conceptual ethics and conceptual engineering, see (Burgess and Plunkett 2013b) and (Cappelen and Plunkett 2020).

¹⁸We should clarify two things about how we discuss “semantic change” proposals in this paper. First, in this paper, we do not take words to be individuated by their meanings. For example, we will allow that a single word ‘bank’ can have one meaning that is about the sides of a river, and another that is about a kind of financial institution. For example: if you think that there are two different English words written as ‘bank’, everything we say could be re-phrased in terms of groups of homophonous but distinct words. Second, for ease of presentation, we will put much of what we say about “semantic change” proposals (and other kinds of conceptual engineering proposals) in terms of single words (e.g., ‘bank’), rather than about strings of words (e.g., ‘banks in Ohio’). However, all of our core points carry over to conceptual engineering proposals that concern strings of words instead.

¹⁹See (Haslanger 2000). In that paper, Haslanger describes the approach we are focused on here as an “analytic” one, but in later work goes on to describe it as an “ameliorative” one (which we take to be the more helpful terminological choice). See (Haslanger 2020).

²⁰For discussion of examples of such “eliminativist” proposals, see (Burgess and Plunkett 2013b). There are other possibilities as well. On many metasemantic views, the intension of a word can be preserved even over substantial changes in the use of that word. On such views, another natural conceptual engineering project seeks to alter prevailing use, or the common or conventional pragmatic effects of such use. There are also interesting cases that are somewhat intermediate between these options, such as those that employ a version of the sort of “sub-scripting” move advocated for by (Chalmers 2011) in thinking through the idea of “merely verbal disputes” and how to best avoid them.

semantics for ‘topic continuity’ as it is used in the context of semantic change proposals. In the conclusion of this paper, we discuss the prospect of extending our account to other contexts, both in conceptual engineering and beyond.

As we see it, broadly Strawsonian worries help us identify a useful *functional role* to associate with the term ‘topic continuity’. The idea is that one desideratum for semantic change proposals is *continuity* with the prior semantics for the word, in ways that are theoretically important in the relevant context.

It is important to emphasize that we are suggesting engineering ‘topic continuity’ to pick out *one* desideratum for semantic change proposals, not the only, or necessarily most important, one. To see this, consider some *other* plausible desiderata that one might have for such proposals. First, consider that conceptual engineering proposals are often motivated by the thought that if the engineering proposal is implemented, something good will result. Perhaps the engineered semantics will help us to *avoid incoherence in thought*, or *carve nature at the joints*, or *ask and answer more precise questions*, or *ameliorate a certain sort of injustice*.²¹ Focus on one example of this: the idea that the engineered semantics will help us to avoid incoherence in thought. A proposal could fully satisfy this desideratum, while failing to maintain continuity with the prior semantics of the term. Next, notice that even if a conceptual engineering proposal preserves topic, it might be bad to *attempt* to implement it. For example, it might be *futile*, if the proposal has no real chance of widespread uptake.²² Or it might be that *attempting* to implement a certain proposal would generate undesirable confusion or backlash.²³

For these reasons, we think it is important to engineer ‘topic continuity’ in a way that aims to pick out *one* desideratum for semantic change proposals, rather than an overall evaluation of such proposals. Indeed, there might well be sensible semantic change proposals that explicitly aim to break topic continuity, either because in the target case preserving

²¹For discussion of these and other goods that are used in motivating (or assessing) proposals in conceptual ethics and conceptual engineering, see (Burgess and Plunkett 2013b) and (Cappelen and Plunkett 2020).

²²See (Cappelen 2018) for the idea that on some influential views about the foundations of semantics (including those that he endorses), we face a sort of “cluelessness” problem: we don’t have practically useful information about what it would take to implement a conceptual engineering proposal. If true, this might exacerbate futility worries about conceptual engineering proposals quite generally.

²³It is also worth noting that in some cases, it may also be important (either as a matter of sociological fact about which proposals will be adopted, or as a normative matter of which ones *should* be adopted) that the proposal comes from someone who has relevant *standing* to introduce such an amendment. For example, perhaps only authorized legislators may introduce certain changes to the intensions of certain legal terms.

topic continuity is outweighed by other considerations, or because it fails to be a desideratum at all. For example, some theories of slurs entail that ethically pernicious content is built into the semantics of some slurring terms.²⁴ On such views, some *reappropriations* of certain slurring terms, which pair the existing word with ethically acceptable content, might be overall desirable to achieve, in part because they fail to preserve topic.

With these clarifications in hand, we now introduce our proposed account.

2. A contextualist proposal

In this section, we sketch our conceptual ethics proposal for how to use ‘topic continuity’ in the context of evaluating semantic change proposals. We then explain why we think this proposal is attractive.

To begin, consider the notions of “topic” and “continuity” generally. We will understand topics very broadly, in terms of *classifications*. Some topics will be natural (e.g. protons) or intuitive (e.g. size). Others will be gerrymandered (e.g. itches felt on Monday mornings or water bottles) or unintuitive (e.g. wave function collapse). Topics can also be more or less *general* (e.g. electrons vs physical particles). Because topics can be at varying levels of generality, and can be gerrymandered, it is plausible that at least *some* topics are preserved in any given shift in discourse tied to the uptake of a semantic change proposal.

Alongside this striking variety of topics, there are also a wide variety of possibly salient ways that one topic could be *continuous* with another. For example, one topic could be a subset of the other (e.g. as mountain biking is a subtopic of cycling), or the two topics could intersect in any number of more complicated ways (e.g. as with Vietnamese cuisine and vegan cuisine).

For the purposes of developing our proposed account of the semantics of ‘topic continuity’, we focus our attention more narrowly, with respect to both *topic* and *continuity*. First, we will understand the relevant *topics* to be anchored to the terms under discussion. Thus, if we were considering a proposal to engineer a new intension for ‘truth’, we would understand the relevant topic as “truth”, and not, e.g., a broader topic such as “communication”. This is important because we take discussion of topic continuity in the relevant contexts to concern the following kind of issue:

²⁴For a paradigmatic example of this kind of view, see (Hom 2008). If one classifies semantics broadly to include (e.g.) expressivist “content” or conventional implicature, many other prominent accounts of slurs take their slurring content to be semantically encoded.

whether a proposed change to the intension of a given word (e.g. ‘truth’) preserves continuity with respect to the topic tied to that word (here: ‘truth’), not whether *some topic or other* is preserved over such a change.

Second, we will understand the notion of “continuity” in terms of three central functions of language: *communication*, *representation*, and *inference*.

To flesh this idea out, we consider central dimensions of continuity in each of these three respects.

Communicative Continuity

1. Preserving successful direct communication using the term.

Ordinarily, entrenched linguistic conventions make it possible to use language to communicate efficiently. Attempting to alter the intension of a term can threaten to be deceptive or to produce miscommunication.²⁵

2. Preserving substantive as opposed to merely verbal disputes.

A related issue is that when one person uses a term ‘X’ with intension A, and another uses it with intension B, there is a danger that they will end up “talking past” each other and get involved in “merely verbal disputes”: roughly, disputes where they think they are expressing a disagreement about something important, but that is just because they (falsely) think they are both using the same word with the same intension (or with similar enough intensions). In contrast, other disputes manage to express substantive disagreements. Attempting to alter the semantics of a term could lead to verbal disputes.²⁶

3. Preserving smooth indirect communication.

Languages include conventions for when one can felicitously report what someone else says. One natural question about a semantic change proposal is whether (and the degree to which) speakers can felicitously use the engineered term to report the beliefs or claims of those who used the term prior to the engineering. For example: suppose that we engineer a new intension for ‘marriage’. The question here is whether we can report the beliefs of people

²⁵For discussion, see (Sterken 2020).

²⁶For a discussion that ties topic continuity closely to avoiding verbal disputes, see (Knoll 2020). For helpful discussion on verbal disputes in general, see (Chalmers 2011) and (Jenkins 2014). Note that our wording about the importance of avoiding verbal disputes doesn’t take a stand on what exactly it takes for something to be a verbal dispute, or exactly when and why variations in what people mean by their words lead to verbal disputes.

who used the term with its old intension as beliefs *about marriage*.²⁷

Representational Continuity

1. *Preserving the representational content of the word.*

In general, semantic change proposals will alter the representational content of a term to at least some extent. Still, many want to preserve *some* amount of continuity in the representational content of the term. On some views of content, perfect continuity in this respect might consist in identical intensions (on other views of content, contents can be individuated *hyperintensionally*, in which case intension identity would be insufficient for perfect representational continuity).

2. *Permitting users to continue to think and speak about what mattered in prior use of the term.*

In many cases, semantic change proposals are motivated by the idea that the existing term to be engineered is an imperfect vehicle for some important task. Given this, it will often be important to preserve what mattered in prior use of the term. What mattered might vary with the term. For example, the term may have had a role in allowing us to represent or to track a feature of reality, or some evaluative role associated with its use.

Inferential Continuity

1. *Preservation of inferential patterns previously associated with competent use of the relevant term.*

Natural language terms are typically associated with a variety of inferential tendencies which enhance the reasoning capacities of competent users of those terms. A semantic change proposal might alter or disrupt the pattern of inferential tendencies associated with a given term.²⁸

All of these dimensions of continuity are potentially *gradable*. For example, a change might result in more or less radical deception or confusion, more or less frequently. Similarly, a change might preserve more or less of the

²⁷See (Cappelen 2018).

²⁸For an example, consider the sorts of changes in inferential patterns that are sometimes wrought by what Thomas Kuhn called “scientific revolutions”, which involve the introduction of new “paradigms” (Kuhn [1962] 2012).

representational content of the term, and render felicitous more or fewer of the inferential patterns characteristic of prior use of the term.

We take these diverse sorts of continuity to be relevant for the following reason: we think each characterizes a dimension of potential continuity or discontinuity in language that someone evaluating a semantic change proposal might reasonably care about. We can locate the heart of this concern within a broader reflection on semantic change proposals as a kind of conceptual engineering activity. Why would a conceptual engineer choose to advocate for semantic change, as opposed to simply introducing a novel word? Presumably because she thinks that there are some values (at least for certain people, in certain contexts) tied to retaining the existing word, and associating it with a new semantics. We can call these values the (putative) *benefits* of the semantic change proposal, and we take continuities along the dimensions we have sketched to be associated with such benefits. However, discontinuities along the dimensions we have sketched here will often be associated with *costs* of moving to the new pairing.²⁹ For example, if a semantic change proposal alters felicitous inferential patterns associated with the use of a term, it could lead users to make more unreliable inferences when they have imperfectly adopted the new pairing. It makes sense for conceptual engineers to track such potential costs. But notice two things. First, the extent of such costs associated with a given discontinuity can vary across contexts. Second, inquirers' concerns with each type of cost can vary across contexts. We think it is useful to use the term 'topic (dis-)continuity' to mark *contextually useful* assessment of such costs.

This puts us in a position to sketch our proposed semantics for the term 'topic continuity' for use in contexts where we are discussing the merits of semantic change proposals:

- (i) The underlying dimensions of "continuity" or "break" relevant to fixing the intension of 'topic continuity' (and related terms such as 'topic discontinuity' and 'topic continuous', etc.) are built into the

²⁹We can usefully contrast two sorts of costs associated with implementing a proposal. Some are costs borne *during transition* to the uniform uptake of the new semantics. For example, "merely verbal disputes" (of the sort discussed in (Chalmers 2011) and (Jenkins 2014)) are a salient danger when only some members of the linguistic community are using the term in the new way. There are also worries about various kinds of deceptive or misleading speech that are salient during transitions (see (Sterken 2020)). Other costs will be most striking with uniform uptake. For example, such a cost would be present if the semantic change makes it harder to speak about what mattered in the prior use of the term.

character of these terms.³⁰ These dimensions are the ones we sketched above under the headings “Communicative Continuity”, “Representational Continuity”, and “Inferential Continuity”.

- (ii) Facts about the interests of participants in the context of utterance fix the following two things:
 - (a) how to weigh the dimensions against each other to get an overall “continuity score”, and
 - (b) *how much* continuity is required for ungraded applicability of the term ‘topic continuity’ (and related terms such as ‘topic discontinuity’ and ‘topic continuous’, etc.).

A few brief comments on this proposed semantics. First, on the basis of the preceding discussion, we are confident that the dimensions of “break” we mention in (i) are central to a useful characterization of ‘topic continuity’. Linguistic items are used to do many things, not all of which will be directly covered by their communicative, representational, and inferential roles. However, we are tempted by the thought that in most cases, they can be used to do these other things in virtue of their communicative, representational, and inferential roles. We recognize that on some ways of theorizing language this may not be true. For this reason, our proposal is, in this respect, tentative. We return to this issue in §4.4.

Second, there are significant questions about how precisely to understand the function from the interests of conversational participants to the contextually-specified intension for ‘topic continuity’ mentioned in (ii). For example: do the relevant facts about the participants’ interests ground the intension facts directly, or are such facts directly grounded in social facts about (e.g.) what has been added without objection to a “conversational scoreboard” etc? These are among the important questions concerning how to theorize about the semantics of context-sensitive expressions.³¹ We propose that *whatever* mechanisms ordinarily do this work for context-sensitive expressions be recruited to do this work here.

Third, condition (ii.b) allows us to usefully treat the term ‘topic continuous’ (etc.) as having both graded uses (for example: “A is somewhat topic continuous with B”) and ungraded uses (for example: “A is topic continuous with B”). By analogy with a natural view about many terms that have

³⁰For determinacy, we here assume a character-based semantics for context-sensitive terms. However, we aim to give ‘topic continuity’ a semantic treatment that is consistent with the natural language semantics for context-sensitive terms, however they are best modeled. So we would be happy to reformulate the theory within a range of alternative frameworks for theorizing about context-sensitivity, such as an indexed truth account.

³¹For a useful introduction to the debates, see (Cappelen and Dever 2016).

both such uses (e.g. ‘flat’), it is plausible that it can be useful to let contextually salient standards fix what threshold is required for ungraded uses. And this is precisely what condition (ii.b) does.

With this sketch of our proposal in hand, we can now make our case for this account as a conceptual ethics proposal for the term ‘topic continuity’ (and closely connected terms, such as ‘topic continuous’, ‘topic discontinuity’, etc.) as it applies to semantic change proposals in conceptual engineering. The core case for the proposal has two parts. First, we think that each of the dimensions of continuity that we have introduced above represents a feature that those evaluating semantic change proposals should often attend to, in light of the potential transition costs involved with discontinuities along these dimensions. So we think it would be counterproductive for an account of the semantics of ‘topic continuity’ to simply rule out any of these dimensions as irrelevant. Second, we think that both the transition costs associated with a given dimension, and any concerns about the sorts of transition costs, can vary across contexts. This is because conversational participants evaluate semantic change proposals with diverse sorts of interests and aims. In light of this, we think that it is useful to have a context-sensitive term that can help conversational participants coordinate on what matters to them in these respects in their conversational context. Our semantics for ‘topic continuity’ treats this term as a label around which to coordinate on these matters.

Careful readers will notice that our account falls within its own scope: that is, we can ask whether our account preserves the topic of “topic continuity”. We think that – at least across a wide range of contexts of evaluation – it does a good job in this respect. Here, for brevity, we will focus on the dimension of *permitting users to continue to speak about what mattered in prior use of the term*. One thing that is striking about existing discussions of topic continuity is the variety of concerns that seem to be animating them. These include everything from concerns about preserving what is important in prior use,³² to preserving one or another sort of function of the relevant term,³³ to fitting in with our existing conventional practices for reporting (and other indirect communication).³⁴ We think that each of these is a concern that it can be useful to track via the use of the term ‘topic continuity’, at least in some contexts. Our account vindicates this variety of concerns by allowing ‘topic continuity’

³²For example, see (Strawson 1963).

³³For example, see (Haslanger 2000) and (Thomasson 2020).

³⁴For example, see (Cappelen 2018).

to track those of the mentioned concerns that are most relevant in a given context.

With this brief case for our account in hand, the remainder of the paper is dedicated to two tasks. In the next section we illustrate the account with a case study. Then, in §4, we continue the evaluation of our account by comparing it to some salient foils.

3. An illustration: Railton's "reforming definition" of 'good'

In this section we illustrate our proposed semantics for 'topic continuity' by reference to a concrete example: Peter Railton's so-called "reforming definition" of 'good'.³⁵ A natural question about his proposal is precisely whether it preserves the topic of prior discussions conducted using the word 'good'.

As we understand it, what Railton calls a "reforming definition" is a type of what we have called a "semantic change" proposal. On a slightly simplified version of Railton's proposal, the intension associated with 'good' (as it is used in expressions of the form '... X is good for individual A ...') is to be reformed so that it is fixed by facts about what a (nearly) omniscient and perfectly instrumentally rational version of A would want A (the actual, non-idealized individual) to want.³⁶

There are many open theoretical questions that are relevant to determining whether Railton's proposal preserves the topic tied to our current use of 'good'. For example, one important question is whether the proposal actually amounts to any change at all to the intension of the word 'good'. On some views, the substance of Railton's account – or at least something quite close to it – might accurately capture the *current* intension of the word.³⁷

Supposing that his proposal *does* change the word's intension, one question relevant to the application of our account of 'topic continuity' is *what matters* in our current (unreformed) use of the term. Some of

³⁵Railton's "reforming definition" approach draws on work by Richard Brandt, such as (Brandt [1979] 1998).

³⁶(Railton 1986, 173–174).

³⁷This might follow on certain metasemantic views – such as the views advocated for by (Schroeter and Schroeter 2014), (Ball 2020), and (Dworkin 2011) – that (put roughly) incorporate what some would see as normative facts relevant to "reforming" proposals about ethical terms (or perhaps terms more generally) as in fact relevant to what actually determines the meanings of the current terms. Note also that there are "ideal attitude" or "ideal advisor" views in ethics, structurally quite similar to Railton's, such as (Smith 1994), which are not put forward as reforms of the relevant normative or evaluative terminology, but rather as analyses of the meaning of our current terms (or concepts). In other cases of "ideal attitude" or "ideal advisor" views in ethics, such as the view defended in (Lewis 1989), it is not entirely clear whether we should understand the account as a semantic change proposal or not.

Railton's own discussion suggests a proposal as to what matters in our current use of 'good'. Roughly, Railton argues that his proposal for how we should use 'good' should "capture" the "normative force" of the current use of the term, while also improving on that use by permitting robust epistemic access to the relation picked out by the term, and having that term pick out something that is naturalistically respectable.³⁸

Another question concerns how much the reform has *inferential* continuity with prior use. Will people who think thoughts using Railton's reformed definition of 'good' have thoughts with similar inferential patterns as the thoughts of those who think using the unreformed term, or will there be substantial differences?

And finally, we can ask to what degree Railton's proposal achieves *communicative* continuity. Will attempts to use the term with the new intension tend to produce confusion? Will it allow smooth indirect communication involving reports of both unreformed and reformed uses? Consider, for example, the following statements: "G.E. Moore had some surprising views about goodness, and Railton does too" or "Railton has a different view about goodness than Kant".

These are, we take it, important theoretical questions about Railton's proposal. For simplicity, however, we will illustrate our proposal by focusing on two possibilities. The first is that Railton's proposal achieves communicative continuity, but fails to preserve the "normative role" that "really matters" in the prior use of the term.³⁹ By the 'normative role' of a term, we mean, roughly, the role characteristically played in deliberation and evaluation by thoughts which token the concept expressed by the term.⁴⁰ The second possibility is that Railton's proposal *does* preserve that normative role, but nonetheless would lead to regular and substantial communicative breakdowns.

We think that in *some* contexts, conversational participants will reasonably be centrally concerned with preservation of the "normative role" of 'good'. They might reason as follows: *we* are interested in 'good' as a *normative* term, which we use to *guide us* in our deliberations. A term which fails to preserve important elements of that normative role will simply not

³⁸See (Railton 1986, 204–205) and (Railton 1986, 171–172). See also (Railton 2003) for further discussion of the kind of metaethical naturalistic realism Railton aims to defend, which helps motivate the epistemological and metaphysical ambitions of the reforming account offered in (Railton 1986).

³⁹For an argument that reforming proposals along the lines of Railton's fail to preserve important aspects of the normative role of 'good', see (Velleman 1988). In that paper, Velleman focuses on Brandt's proposal, but argues that his core critique raises issues for related proposals, including Railton's.

⁴⁰This use of "normative role" draws from (Eklund 2017), drawing on (Railton 1986)'s discussion of "evaluative role". For critical discussion of Eklund's idea of "normative roles" see (Plunkett 2020).

do what we want ‘good’ to do, and hence will change the topic. On our proposal, in conversational contexts in which participants have these interests, the statement “the second possibility preserves topic, but the first does not” will be true.

We think that in other contexts, participants may reasonably be most concerned with the important *communicative* functions of evaluative terms, in allowing, e.g. constructive evaluative discussion and thereby solutions to coordination problems. On our proposal, in conversational contexts in which participants have these interests, the statement “the first possibility preserves topic, but the second does not” will be true.

This example illustrates the significance (and, we think, the virtue) of our contextualist account. In our two imagined contexts, participants have converged on quite different priorities about what sorts of “break” with prior use amount to a loss of what was important in prior discussion of the topic. We think both sets of participants should be able to use the term ‘topic continuity’ to structure their investigation of possible upshots of implementing a Railton-style semantic change proposal for ‘good’, in ways that are sensitive to their priorities.

4. Situating our proposal

In order to better understand and evaluate our proposal, we will now situate it with respect to a variety of alternative conceptual ethics proposals about how to use the term ‘topic continuity’ in the context of discussing semantic change proposals in conceptual engineering. Our aim here is twofold: to identify some important choice points regarding the question in conceptual ethics about ‘topic continuity’ that we are addressing in this paper, and to advertise what we take to be the initial attractiveness of our view relative to each of these choice points.

4.1. Reform vs. abandonment

We motivated our proposal by noting two things: first, that there seem to be multiple dimensions of “break” that we might be concerned with, and, second, that interest in these various dimensions will differ across contexts in which we are developing and evaluating semantic change proposals. One might take these very facts to motivate a different conceptual ethics proposal: to simply *abandon* talk of “topic continuity” and replace it with discussion of each of the various dimensions of break. This might seem to better promote clear communication than having a

single context-sensitive term to pick out a contextually varying function of desiderata which are only loosely related.

We have a three-part reply to this alternative idea. First, the case for abandonment cannot simply be that the character we propose for this term is multidimensional and context-sensitive. After all, people often seem to use multidimensional context-sensitive terms “in the wild” in both everyday life and in theoretical inquiry in ways that are productive. Second, we think that the sorts of desiderata that we have built into the semantics of ‘topic continuity’ – desiderata concerning communication, representation, and inference – are aspects of a unified sort of concern. Put roughly, that concern is with continuity regarding what are arguably some of the most (if not *the* most) fundamental functions of language use. Third, we have also suggested a positive reason why it might be useful to have such a term: because it can be used as a success term for *coordination* regarding how to prioritize among an often important and reasonably natural cluster of desiderata on semantic change proposals.

This being said, we are of course open to it being useful in some contexts to focus directly on certain specific dimensions of break, rather than to coordinate on an intension for ‘topic continuity’. A conceptual engineering proposal can earn its keep by providing a discursive tool whose usefulness justifies the effort that getting it into the language requires. This is wholly compatible with the relevant term only being useful to deploy in some contexts.

4.2. Contextualism vs. alternatives

We now consider the possibility of granting that it is useful to retain the word ‘topic continuity’ but questioning whether, as we suggest, the best way of engineering that term is a context-sensitive way. Some of the other salient options here are forms of invariantism, speaker- or evaluator-relativism, and expressivism. We take each in turn.

The invariantist thinks it would be useful to associate ‘topic continuity’ with a *constant* function to intensions, rather than a contextually variant one. One way to be an invariantist would be to think that there is a single elegant feature that should fix the intension of ‘topic continuity’. For example, it could be argued that some notion of *semantic function* is the single feature that we should be tracking with this term. We have two replies here. First, if we are right that it is useful to have a term around which to coordinate on desiderata, as we suggested above, then the mere existence of some illuminating notion of semantic function

does not undercut our positive case for our contextualist proposal. Second, suppose that it would be useful to have both a term that enables us to talk about continuity of semantic function, and the sort of context-sensitive term we have proposed. Then we have a further question: which of these terms is better to pair with the existing word 'topic continuity'? We think that our proposal's ability to capture the diversity of concerns that have been advertised under the banner of "topic continuity" counts in its favor.

Another possibility is to advocate for a *messy* invariantism. For example, one might argue that we are right that all of the dimensions that we mention should matter for "topic continuity", and that *how much* they matter should be variable across contexts. But one might claim that these normative facts are best captured by a complex invariantist semantics, as opposed to a contextualist semantics.⁴¹ To give a feel for the difference: consider that, presumably, on the invariantist account, how the dimensions vary across contexts will be built into the semantics, as opposed to being a function of, e.g. the interests of speakers in the contexts. Whatever the virtues of messy invariantism as a theory of the current meaning of 'topic continuity', we doubt that it is an attractive *conceptual engineering* proposal for two reasons. First, it is unclear what the benefit of building this sort of complexity into an intension would be. Such complexity appears to make 'topic continuity' pick out a gerrymandered mess. Second, it could be much harder to learn such a semantically messy term, which is an undesirable feature of a semantic change proposal.

A third competitor would be to make the content of 'topic continuity' a function not of conversational context but of individual psychology (either of the speaker, or of the interpreter).⁴² We think that this alternative is less attractive as a conceptual engineering proposal than our proposal. This is because, by tying the semantic value of 'topic continuity' to an individual psychology, it is not as well suited to allow for the semantic value of the term to function as the locus of successful coordination between conversational participants (whose psychologies might vary in key ways).

Another alternative to our account would be to engineer an *expressivist* semantics for 'topic continuity'. Put roughly, on an expressivist proposal, the semantics of the term 'topic continuity' is explained (at the most basic explanatory level) by certain "non-cognitive" attitudes (e.g. desires or pro-

⁴¹It may be that this is the best way of reading some of the remarks in (Cappelen 2020) in responding to (Sundell 2020).

⁴²For a relevant view that suggests a role for evaluator-relativism, see (Cappelen 2020, §2).

attitudes) that are expressed by the use of the term. While a lot could be said about this alternative, we will be brief.⁴³ As we have emphasized above, there are a number of dimensions of continuity – in communication, representation, and inference – that evaluators of semantic change proposals should be concerned with. And we noted above that this presents a challenge for proposals to engineer (or retain) a semantics for ‘topic continuity’: why keep a unified term for a seemingly disunified cluster of phenomena? One possible answer is that there is a unified thing that we want to be able to *do* with ‘topic continuity’ judgments. If we could identify such a unified “downstream role”, then we could offer an expressivist account according to which uttering “*x* is topic continuous with *y*” expresses approval of giving that sort of downstream role to *x* and *y*.⁴⁴

The expressivist can take on board our thought that the value of engineering the term ‘topic continuity’ involves promoting a certain sort of coordination – after all, the focus on promoting such coordination is a key part of several expressivist proposals.⁴⁵ The expressivist twist in the case at hand is that the relevant coordination is best promoted by a semantics on which relevant speech acts function semantically to *advocate* for one or another coordination point. For this reason, we think the expressivist alternative deserves to be taken seriously.

One significant worry about the expressivist proposal, however, is that it is not at all clear what the alleged unified downstream role of ‘topic continuity’ talk would be. For example, as we have emphasized in §1, we do not take it to be attractive to treat “topic continuity” as settling which semantic change proposals to go for. Rather, we suggested that it is best understood as one among many defeasible desiderata on such proposals.

4.3. Which contextually variable factors should determine reference?

Recall that, according to our proposal, what ‘topic continuity’ refers to is context-dependent, and varies across contexts based on facts about the interests of conversational participants. In this respect, our account is

⁴³For simplicity here, we simply assume that expressivism is a competitor to our proposal. However, on some “metasemantic” interpretations of expressivism (such as those advanced by (Chrisman 2016) and (Ridge 2014)) it is not clear that expressivism is a direct competitor to the proposals we have canvassed so far.

⁴⁴More sophisticated models of the expressed content are of course possible (see (Gibbard 2003), (Schroeder 2008), and (Ridge 2014), for example). Here we are offering the simplest picture for the sake of illustration.

⁴⁵See, for example, (Gibbard 1990), (Blackburn 1998), and (Björnsson and McPherson 2014).

conventionally contextualist. One might worry that this sort of account gives objectionable power to conversational *bad actors*: participants who value something worthless or bad in the existing semantics or usage of a given term. To be more concrete, suppose that the current use of the term being considered for semantic change non-accidentally tends to produce racist inferences. And suppose that a bad actor insists that in this case, topic continuity requires a high degree of inferential continuity, such that any semantic change proposal that eradicated the racist inferential tendency would fail to preserve topic. It may seem like a vice that our semantics allows the bad actor's interests to affect the intension of 'topic continuity' in conversations that she participates in.

One might seek to avoid this problem by altering which contextually variable features fix the intension of 'topic continuity'. For example, one could propose that the contextually-relevant features are the *epistemically justified* and/or *ethically acceptable* interests of the participants. We have two things to say about this proposal. First, we think it comes with significant costs. For example, in contexts where the relevant ethical or epistemic facts are opaque or highly contested, it will be much harder to tell, on this account, what the semantic value of 'topic continuity' is. Second, we think it is important to keep in mind that, as we have emphasized above, topic continuity is just *one* desideratum on semantic change proposals. The fact that a proposal will tend to produce racist inferences is a very weighty objection to it, even if that proposal preserves topic. Thus, the fact that bad actors have the ability to affect the intension of 'topic continuity' in a particular context does not mean that they can somehow render objectionable conceptual engineering proposals acceptable. In light of this, we think the cost of bad actors being able to affect which proposals count as "topic continuous" in certain contexts is not a substantial detriment to our proposal.

4.4. Which factors should be built into the character of 'topic continuity'?

On our account, the underlying dimensions of variation relevant to determining the intension of 'topic continuity' are the ones we summarize under the headings of "Communicative Continuity", "Representational Continuity", and "Inferential Continuity". One could embrace the general structure of our contextualist proposal but reject this account of the relevant underlying dimensions.

Here we want to be concessive. We especially want to allow that there may be further dimensions of continuity that we have not captured in

our list, which could be useful, in a non-trivial number of contexts, to track in our use of ‘topic continuity’. This is especially true given the existence of reasonable controversy regarding how to theorize about linguistic representation; some views on this topic might emphasize further functions that we do not.⁴⁶ Alternatively, one might think that we could get by with fewer desiderata on the list. For example, one might think that it would be desirable for the semantics of ‘topic continuity’ to track only the preservation of whatever mattered to people in previous conversations using the term.⁴⁷ One central worry about this sort of proposal is that, in some cases, previous users of the term might be confused or ignorant about what mattered in their use of the term. More generally, we think it is reasonable for people evaluating a semantic change proposal to be able to intelligibly argue about whether – say – the fact that the proposal will foreseeably produce a certain amount of merely verbal disputes during the transition makes for an important “break” in this context. We see no reason why previous users should be left to legislate this matter.

4.5. General lessons from the discussion of alternatives

Our discussion in this section of the merits of our approach relative to a range of important alternatives has been far too brief to be dispositive. Nonetheless, we think that it helps to further motivate our proposal, by explaining why there is something attractive to say about our account at each of the main choice points that structure our preferred approach. We think that the discussion in this section is valuable in another way as well. We think that the conceptual ethics question about ‘topic continuity’ is an important one, and that a variety of important alternatives deserve to be carefully developed and evaluated. We hope to have contributed to this project in this section.

5. Conclusion

In this paper, we have advanced a view about how we should use the term ‘topic continuity’ when doing philosophical work in or about conceptual engineering. More specifically, we’ve focused on what we have

⁴⁶One possible example: on some views, language centrally functions to structure the *commitments* that we take on in making speech-acts (e.g. (Brandom 1994)). On such views, continuity of commitment-implications of a term might be paramount to theorizing about topic continuity.

⁴⁷Thanks to Jennifer Nado for helpful discussion of this sort of suggestion.

called "semantic change" proposals in conceptual engineering, by which we mean proposals for changing the meaning of an existing term, such that the term (at least in certain contexts, or when used by certain people) becomes associated with a new intension or character. We've advocated for a contextualist account on which the underlying dimensions of "continuity" or "break" relevant to fixing the intension of 'topic continuity' (and related terms such as 'topic discontinuity' and 'topic continuous', etc.) are built into the character of these terms. The dimensions we argued are relevant are those that we discussed under the headings "Communicative Continuity", "Representational Continuity", and "Inferential Continuity". In turn, we argued that facts about the interests of participants in the context of utterance fix both (a) how to weigh the dimensions against each other to get an overall "continuity score" and (b) *how much* continuity is required for ungraded applicability of the term 'topic continuity' (and related terms such as 'topic discontinuity' and 'topic continuous', etc.). We hope that this paper helps to put the conceptual ethics question about 'topic continuity' where we think it belongs: at the heart of contemporary theorizing about the projects of conceptual ethics and conceptual engineering.

In order to keep our project in this paper manageable, we have focused on engineering 'topic continuity' for use in a specific type of context: the evaluation of semantic change proposals. We now want to briefly discuss what can be learned from this paper for discussions of topic continuity in other contexts.

To begin, reflect on the range of contexts in which concerns about something like "topic continuity" might arise. These include conceptual ethics discussions that consider *abandoning* the use of a term (at least in certain contexts) and those that consider *introducing* a new term. Moreover, related issues arise in many other areas of inquiry, including other parts of philosophy, intellectual history, history of science, and anthropology.

We do not expect the account we offer here to smoothly extend to all of these contexts. One reason is that one of our anchoring assumptions is less plausible when we move away from semantic change proposals. This assumption is that the relevant *topic* we are testing for continuity is fixed by reference to the word whose semantics a proposal suggests to change. We think that in other contexts, this assumption will be implausible, for two primary reasons. First, in conceptual engineering proposals that introduce a novel word, there may simply not be an obvious salient existing word that people will (either tacitly or explicitly) be using to use to fix

what the relevant topic is at hand. Second, in some contexts, it will be controversial how coarse-grained the relevant topics are. For example, we think this is plausible in the context of studying intellectual history. Consider the following case. In the late twentieth century, many philosophers followed John Rawls in thinking that there is a very close connection between the topic of *justice* and the topic of *how we should organize our basic social and political institutions*. Now suppose that an intellectual historian considers political thinkers who do not assume such a close connection, and asks whether they are addressing the same topic as Rawls. We might get very different answers depending on how coarse-grained we understand the relevant “topic” that Rawls’ work is engaged with. For example, contrast the answers we might get if we understand that topic to be the topic of “justice” as opposed to the topic of “how we should organize our basic social and political institutions” (which, it should be emphasized, not all thinkers throughout history tie in as closely to the topic of “justice” as Rawls does, let alone his specific view of “justice” in terms of a certain conception of “fairness”).

All of this means that we should proceed with caution in extending our proposal to other contexts. However, we think that the methodology we use in this paper *is* quite promising when considering how to engineer ‘topic continuity’ for these other contexts. The general methodology is this: seek to identify the dimensions of continuity or break that matter for the context at hand when talking about ‘topic continuity’, and then build a context-sensitive account like the one we have provided here, but tied to those dimensions in particular. That will yield different proposals about what we should mean by ‘topic continuity’ in these different contexts, but nonetheless ones that are closely related in their contextualist form, and also likely to have significant overlap in substance as well (e.g. in which dimensions of continuity matter).

While it is beyond the scope of this paper to develop this strategy in detail, we are optimistic that it will be generally fruitful. Further, even if this strategy turns out to be unfruitful in some contexts, we hope that the tools and distinctions that we have introduced in this paper provide a useful framework for developing competing ways of theorizing about “topic continuity” both in conceptual engineering and beyond.

Acknowledgements

Thanks to Derek Ball, David Braddon-Mitchell, Alexis Burgess, Herman Cappelen, John Cho, Jesse Ferraioli, Peter Galison, Zachary Lang, Amanda Li, Alison

McQueen, Jennifer Nado, Peter Railton, Adrian Russian, Tim Sundell, Amie Thomason, and Kenny Walden for helpful discussion and feedback. Thanks also to those who participated in discussions of earlier drafts of this paper at the University of Cambridge and the University of Hong Kong, and to the students in classes where we have talked about these issues (including a proseminar on concepts in philosophy at The Ohio State University, and a seminar on philosophical methodology at Dartmouth College).

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

- Ball, Derek. 2020. "Revisionary Analysis Without Meaning Change (Or, Could Women Be Analytically Oppressed?)." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 35–58. Oxford: Oxford University Press.
- Björnsson, Gunnar, and Tristram McPherson. 2014. "Moral Attitudes for Non-Cognitivists: Solving the Specification Problem." *Mind* 123 (489): 1–38.
- Blackburn, Simon. 1998. *Ruling Passions*. Oxford: Clarendon Press.
- Brandom, Robert. 1994. *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- Brandt, Richard B. (1979) 1998. *A Theory of the Good and the Right*. New York: Prometheus Books.
- Burgess, Alexis, and David Plunkett. 2013a. "Conceptual Ethics I." *Philosophy Compass* 8 (12): 1091–1101.
- Burgess, Alexis, and David Plunkett. 2013b. "Conceptual Ethics II." *Philosophy Compass* 8 (12): 1102–1110.
- Burgess, Alexis, and David Plunkett. 2020. "On the Relation Between Conceptual Engineering and Conceptual Ethics." *Ratio* 33 (4): 281–294.
- Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Cappelen, Herman. 2020. "Conceptual Engineering, Topics, Metasemantics, and Lack of Control." *Canadian Journal of Philosophy* 50 (5): 594–605.
- Cappelen, Herman, and Josh Dever. 2016. *Context and Communication*. Oxford: Oxford University Press.
- Cappelen, Herman, and David Plunkett. 2020. *A Guided Tour of Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen and D. Plunkett. Oxford: Oxford University Press.
- Carnap, Rudolf. (1947) 1956. *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago: University of Chicago Press.
- Carnap, Rudolf. (1950) 1962. *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Chalmers, David J. 2011. "Verbal Disputes." *Philosophical Review* 120 (4): 515–566.

- Chrisman, Matthew. 2016. *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. Oxford: Oxford University Press.
- Cooper, John M. 1997. *Plato: Complete Works*. Indianapolis: Hackett.
- Dworkin, Ronald. 2011. *Justice for Hedgehogs*. Cambridge: Harvard University Press.
- Eklund, Matti. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Haslanger, Sally. 2000. "Gender and Race: (What) Are They? (What) Do We Want Them to Be?" *Noûs* 34 (1): 31–55.
- Haslanger, Sally. 2020. "Going On, Not in the Same Way." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen and D. Plunkett. Oxford: Oxford University Press.
- Hom, Christopher. 2008. "The Semantics of Racial Epithets." *Journal of Philosophy* 105 (8): 416–440.
- Jenkins, C. S. I. 2014. "Merely Verbal Disputes." *Erkenntnis* 79 (1): 11–30.
- Knoll, Viktoria. 2020. "Verbal Disputes and Topic Continuity." *Inquiry*.
- Kuhn, Thomas. (1962) 2012. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Lewis, David. 1989. "Dispositional Theories of Value." *Proceedings of the Aristotelian Society* 63: 113–137.
- McPherson, Tristram, and David Plunkett. 2017. "The Nature and Explanatory Ambitions of Metaethics." In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett, 1–25. New York: Routledge.
- McPherson, Tristram, and David Plunkett. 2020. "Conceptual Ethics and The Methodology of Normative Inquiry." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 274–303. Oxford: Oxford University Press.
- McPherson, Tristram, and David Plunkett. 2021. "Conceptual Ethics, Metaepistemology, and Normative Epistemology." *Inquiry*.
- Plunkett, David. 2020. "Normative Roles, Conceptual Variance, and Ardent Realism About Normativity." *Inquiry* 63 (5): 509–534.
- Plunkett, David, and Tim Sundell. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosophers' Imprint* 13 (23): 1–37.
- Prinzinger, Michael. 2018. "The Revisionist's Rubric: Conceptual Engineering and the Discontinuity Objection." *Inquiry* 61 (8): 854–880.
- Railton, Peter. 1986. "Moral Realism." *The Philosophical Review* 95: 163–207.
- Railton, Peter. 2003. *Facts, Values, and Norms*. New York: Cambridge University Press.
- Ridge, Michael. 2014. *Impassioned Belief*. Oxford: Oxford University Press.
- Schroeder, Mark. 2008. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- Schroeter, Laura, and François Schroeter. 2014. "Normative Concepts: A Connectedness Model." *Philosophers' Imprint* 14 (25): 1–26.
- Schroeter, Laura, and François Schroeter. 2020. "Inscrutability and Its Discontents." *Canadian Journal of Philosophy* 50 (5): 566–579.
- Smith, Michael. 1994. *The Moral Problem*. Cambridge: Blackwell.

- Sterken, Rachel Katharine. 2020. "Linguistic Interventions and Transformative Communicative Disruption." In *Conceptual Engineering and Conceptual Ethics*, edited by H. Cappelen, D. Plunkett, and A. Burgess, 417–424. Oxford: Oxford University Press.
- Strawson, Peter F. 1963. "Carnap's Views on Conceptual Systems Versus Natural Languages in Analytic Philosophy." In *The Philosophy of Rudolf Carnap*, edited by P. A. Schilpp, 503–518. La Salle: Open Court.
- Sundell, Tim. 2020. "Changing the Subject." *Canadian Journal of Philosophy* 50 (5): 580–593.
- Thomasson, Amie L. 2020. "A Pragmatic Method for Normative Conceptual Work." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 435–458. Oxford: Oxford University Press.
- Velleman, J. David. 1988. "Brandt's Definition of "Good"." *The Philosophical Review* 97 (3): 353–371.