INQUIRY
2025, VOL. 68, NO. 7, 1501–1533
https://doi.org/10.1080/0020174X.2021.1930147

Routledge
Taylor & Francis Group

OPEN ACCESS Check for updates

Conceptual ethics, metaepistemology, and normative epistemology

Tristram McPherson[a] and David Plunkett[b]

[a]Philosophy, The Ohio State University, Columbus, OH, USA; [b]Department of Philosophy, Dartmouth College, Hanover, NH, USA

ABSTRACT

This paper advertises the importance of distinguishing three different foundational projects about epistemic thought and talk, which we call "systematic normative epistemology", "metaepistemology", and "the conceptual ethics of epistemology". We argue that these projects can be distinguished by their contrasting constitutive success conditions. This paper is motivated by the idea that the distinctions between these three projects matter for epistemological theorizing in ways that have been underappreciated in philosophical discussion. We claim that attention to the threefold distinction we advance allows us to better understand and evaluate existing views and debates in the field; identify and appreciate new or underexplored theoretical options in the field; and avoid important defects and ambiguities in our research on epistemic topics.

ARTICLE HISTORY Received 21 January 2021; Accepted 11 May 2021

KEYWORDS Epistemology; conceptual ethics; metaepistemology

Introduction

Contemporary philosophers widely agree that epistemology involves inquiry concerning (or connected to) *knowledge*, what *justifies* belief, what constitutes *good evidence* or *understanding*, and (at least in some sense) what it is *rational* to believe. Suppose we take such topics to mark the core of what "epistemology" is about. A question remains: what *kind* of inquiry about these topics do (or should) epistemologists engage in? Recent epistemologists offer a strikingly diverse range of answers to this question. In this paper, we argue that this diversity of views should be taken seriously. We think "epistemology" marks out –

CONTACT Tristram McPherson dr.tristram@gmail.com

© 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.

at a high level of abstraction – a relatively unified subject matter, or subarea of philosophy.[1] However, within this subarea, we argue that there are a range of importantly distinct projects that philosophers are (or could be) engaged in, even if we just focus on a given paradigmatic "epistemic" condition such as knowledge. More specifically, we argue for the importance of the contrasts among three projects concerning the epistemic which we call *systematic normative epistemology*, *metaepistemology*, and the *conceptual ethics of epistemology*.

We do not argue that these are the only three projects that philosophers could or should focus on in epistemology. Far from it. Rather, we focus on these specific projects for two interconnected reasons. First, while the distinction between "normative ethics" and "metaethics" is a familiar one, the parallel distinction has received much less sustained attention in epistemology.[2] In this paper, we aim to show that increased attention to this parallel distinction is useful in epistemology. Second, in recent years, the topics of "conceptual ethics" and "conceptual engineering" have received increasing philosophical attention.[3] In this paper, we aim to show that increased attention to these topics is useful in epistemology.

We break up our work in this paper as follows. We begin by putting forward our accounts of systematic normative epistemology, metaepistemology, and the conceptual ethics of epistemology (in §§2-4, respectively). We understand each of these projects as types of *inquiry* defined by their *aims* (as opposed to these projects being defined by distinctive collections of claims, or by sociological features of those that pursue these projects, etc.). We argue that this way of proceeding helps us identify philosophically fruitful cuts within inquiry in (or about) epistemology, while also avoiding certain apparent challenges to the coherence or significance of distinguishing between these three projects. In §5, we argue that the distinctions between these projects are philosophically deep. This is because the aims that are constitutive of each project are importantly different from each other. Following this, in §6, we show that the distinctions between these projects have important methodological implications for how we conduct inquiry in (or about) epistemology. In short, we argue that attending to these distinctions allows us to better evaluate existing views and debates, appreciate new or underexplored

[1] In this paper, single quotation marks (e.g. 'cat') are used strictly to mention linguistic items. Double quotation marks (e.g. "cat") are used for a variety of tasks including quoting others' words, scare quotes, and mixes of use and mention. Terms in small caps (e.g. CAT) pick out concepts.

[2] This is not to say that it has received no attention, however. For example, see (McHugh, Way, and Whiting 2019).

[3] See, for example, the papers collected in (Burgess, Cappelen, and Plunkett 2020).

theoretical options, better understand what kinds of arguments are needed to support claims, and avoid important kinds of unreliable inference. In §7, we briefly address several challenges to the depth and methodological importance of the distinctions among systematic normative epistemology, metaepistemology, and the conceptual ethics of epistemology that we have outlined here.

Before we begin, we want to flag a pair of connected questions that arise in each of the projects that we discuss. The first question is whether the range of "epistemic" phenomena constitutes a unified topic, and if so, what that unity consists in. For example: do these phenomena form a natural or normative kind, or are they unified only by idiosyncrasies of our classificatory practices? The second question is *which* words, concepts, relations, properties, etc. count as "epistemic", and why. Consider two examples of controversy related to this question. First, do questions about what we "really and truly" ought to believe count as properly "epistemic" questions, as opposed to ones in the "ethics of belief"? Second, is aversion to acquiring reliable evidence an "epistemic" vice? We explore these questions at length in a companion paper.[4] However, in this paper, we put these questions aside and simply focus on what are widely regarded as paradigmatic "epistemic" conditions, such as knowledge and justification.

1. The diversity of self-conceptions in recent epistemology

In this section, we motivate our project by briefly advertising the striking variety of ways that recent epistemologists describe either their work or the field as a whole. There are many such conceptions we could gloss on this front. We focus on five examples which illustrate the wide range of views on the table.

First, some philosophers characterize epistemology as aiming to provide a general theory of what it takes to satisfy some epistemic condition, or what a given epistemic condition consists in. For example, Richard Feldman proposes that "[i]t is the epistemologist's business to try to develop a general theory stating the conditions under which people have knowledge or justified belief".[5]

Second, many epistemologists claim that a (perhaps *the*) standard project in epistemology is to offer conceptual analyses of concepts like

[4]See (McPherson and Plunkett Manuscript).
[5]Feldman (2003, 1).

KNOWLEDGE or words like 'knowledge'. This is suggested, for example, by Frank Jackson's discussion of Gettier cases in *From Metaphysics to Ethics*, and by the orienting role Duncan Pritchard gives to the "analytical project" in his textbook *Epistemology*.[6] This kind of view also features as the standard foil for epistemologists defending alternative understandings of the project. For example, Hilary Kornblith writes that "the subject matter of epistemology is knowledge itself, not our concept of knowledge".[7] And Timothy Williamson states that the primary concern of epistemology is *not* "with the nature of the concept of knowledge. If knowledge were in fact identical with justified true belief, that would be what mattered epistemologically, irrespective of the conceptual possibility of their non-identity".[8]

Third, in his agenda-setting *Knowledge and Its Limits*, Williamson proposes that "knowing is a state of mind".[9] The study of knowledge is thus, for Williamson, a branch of the philosophy of mind, dedicated to understanding the nature of a type of mental state and its distinctive explanatory roles.

Fourth, in contrast to Williamson, Edward Craig (whose basic approach to epistemology has become increasingly influential in recent years) suggests that "[k]nowledge is not a given phenomenon, but something that we delineate by operating with a concept which we create in answer to certain needs, or in pursuit of certain ideals".[10] On Craig's view, our central goal in epistemology should be to understand the *function* of the concept KNOWLEDGE, which humans have introduced to do certain indispensable work in our social lives.

Finally, some philosophers frame their projects in epistemology as explicit attempts to *improve* how we form and assess our beliefs. For example, Alvin Goldman frames the goal of what he calls "veritistic epistemology" as identifying which practices promote true belief relative to error and ignorance.[11] Some such projects involve proposals for evaluating and improving the concepts we deploy in epistemology. For example, Sally Haslanger claims that, in theorizing about knowledge,

[6]Jackson (1998, 28) and Pritchard (2016, 1–2).
[7]Kornblith (2002, 1).
[8]Williamson (2007, 206).
[9]Williamson (2000, 21).
[10]Craig (1990, 3).
[11]Goldman (1999, 5). This is a classic ambition for epistemology, which (Kitcher 1992, 64) calls the "meliorative dimension" of the field; compare e.g. (Descartes 1641/2017), and, more recently, (Bishop and Trout 2004).

> the task is not simply to explicate our ordinary concept of X; nor is it to discover what those things we normally take to fall under the concept have in common; instead we ask what our purpose is in having the concept of X, whether this purpose is well-conceived, and what concept (or concepts) would serve our well-conceived purpose(s).[12]

In recent work, Kevin Scharp argues for a view that, like Haslanger's, targets normative questions about the concepts we deploy in epistemology. As part of this, he argues that the concept KNOWLEDGE (among many central philosophical concepts we currently use) is defective and should be replaced.[13]

One possible reaction to these diverse self-conceptions is to conjecture that these philosophers are engaged in the same sort of inquiry, and simply have contrasting understandings of that shared inquiry. One might motivate this reaction by appeal to an observation from Richard Fumerton. Fumerton says that, despite their differing self-conceptions, epistemologists "seem to have relatively little trouble engaging in the give and take of argument over the specific answers to those questions that have been proposed".[14] This apparently smooth dialectical engagement might seem to be evidence that epistemic inquiry is unified. Another possible reaction is that, while many of these philosophers are engaged in different projects, the differences between these projects are *inconsequential*.[15]

We take these reactions to be mistaken. We offer a framework that suggests that, at least in many cases, the sorts of diverse self-conceptions sketched above reflect a range of importantly distinct projects. In particular, we take them to motivate distinguishing three projects within epistemology, which we call *systematic normative epistemology*, *metaepistemology*, and the *conceptual ethics of epistemology*. In characterizing these projects, our primary aim is to draw *theoretically useful* distinctions, rather than to capture existing usage of the terms we use to label those projects (e.g. 'metaepistemology').[16] We now turn to discussing these three projects.

[12]Haslanger (2000a, 468).
[13]See Scharp (2020) and Scharp (Manuscript).
[14]Fumerton (2006, 10).
[15]Compare (DeRose 2009, 19), who makes this sort of claim about the boundary between epistemology and the semantics of epistemic words.
[16]In the terms introduced in §4 below, our project in this paper is thus a kind of conceptual ethics project. Our proposed labels for these projects draw on existing strands of usage of the proposed terminology (and analogous terminology), but our goal is not to fully capture the exact contours of that usage. Instead, we use this terminology in the hopes that it is theoretically fruitful for philosophical inquiry, in our current social/historical context.

2. Systematic normative epistemology

In this section, we introduce the project that may be most familiar to many epistemologists – what we call *systematic normative epistemology* – and use it to illustrate important features of the sort of distinction among projects that we are offering in this paper.

To begin, consider a pithy statement of the project as we understand it:

> **systematic normative epistemology** is that inquiry that aims to identify and explain, in general terms, what it takes to satisfy an epistemic condition, and why.

To illustrate this formulation, consider a familiar kind of theory which can be treated as an exemplary contribution to this project. According to this theory, which we call *simple reliabilism about knowledge,* what it is for S to know that P is for S to truly believe that P, and for S's belief that P to have been formed by an appropriately reliable belief-forming mechanism. Consider how this theory illustrates the central features of systematic normative epistemology. First, simple reliabilism concerns knowledge, which is a paradigmatic epistemic condition. (Other paradigmatic epistemic conditions include epistemic justification, evidence, epistemic rationality, and understanding.) Next, as systematic normative epistemology requires, simple reliabilism is a theory that purports to tell us *what it takes* for S to know that P: namely, for S to have a true belief that P, which was formed by an appropriately reliable belief-forming mechanism. Finally, notice that simple reliabilism proposes a constitutive explanation: it purports to say *what it is* for S to know that P. It thus purports not merely to tell us what is required for knowledge, but to illuminate why that is required.[17]

Simple reliabilism is only one (helpfully simple) example of the sort of theory that, if correct, would partially satisfy the aims of systematic normative epistemology. Even if correct, it would not *completely* satisfy those aims. One reason is that it is silent concerning epistemic conditions other than knowledge. One could develop reliabilist accounts of these other conditions, but simple reliabilism itself is not such an account.

With respect to knowledge, simple reliabilism is perfectly general in scope: it purports to tell you, for example, what counts as knowledge about *any* subject. We take this generality to be partially definitive of

[17]It is also worth noting that our characterization of systematic normative epistemology is neutral concerning what exactly would constitute a full explanation of the issue at hand. Whether the constitutive explanation provided by simple reliabilism provides such a full explanation depends on how this controversy about explanation is resolved.

systematic normative epistemology. By contrast, consider inquiries that focus on (e.g.) specifically *moral* or *mathematical* knowledge. Because these inquiries focus on what it takes to have knowledge about a specific subject, or in a specific context, we classify them as instances of *applied normative epistemology*. We take this to be an instance of a useful general cut between systematic and applied normative inquiry, which also applies (for example) to inquiries in ethics and political philosophy.[18]

Simple reliabilism is in one respect strikingly ambitious: it purports to provide necessary and sufficient conditions for knowledge. Here, as elsewhere, many philosophers are suspicious about whether a tidy set of such conditions exist. We take systematic normative epistemology to aim to provide as much generality as possible, in identifying and explaining what it takes to satisfy an epistemic condition. On some views, what is possible may amount only to the provision of defeasible general principles, or to the identification of only some necessary or sufficient conditions. This explains why we can think of some skeptical theories as contributions to systematic normative epistemology, for example if those theories purport to identify a necessary condition for knowledge that entails that knowledge is never instantiated.

Our label 'systematic normative epistemology' encodes a controversial assumption: that inquiry into epistemic conditions such as knowledge is a kind of *normative* inquiry. Many epistemologists understand epistemology as "normative" in some sense.[19] This idea is hard to deny with respect to many paradigmatic epistemic conditions, such as epistemic *justification*, *reasons*, or *rationality*. Or, to take another example, we might think that knowledge provides an important sort of norm for believing, asserting, or inquiring, and thus marks a certain kind of epistemic achievement worth aiming at. Indeed, something like this idea is arguably what motivates much of the focus on knowledge in epistemology to begin with, whether in the attempt to analyze knowledge or in recent "knowledge-first" epistemology.[20] It may seem less clear, however, that certain other epistemic conditions, such as *evidence*, are normative. Whether they are or not will depend, among other things, on how one understands

[18]For connected discussion (about the parallel case of *applied ethics*), see (McPherson and Plunkett 2017).
[19]Consider two examples of strong statements of this idea. First, (Berker 2013) argues that just as the central question of ethics is "what should I do?", the central question of epistemology is "what should I believe?". Second, (Haslanger 2000a, 466, emphasis original) claims that epistemology's normative dimension involves the fact that it concerns such things as "how we *ought* to reason, on what basis we *ought* to form beliefs, and more generally what is epistemically *valuable*".
[20]E.g. in canonical texts like (Gettier 1963) and (Williamson 2000).

those conditions and what exactly it is for a condition to be "normative" (as opposed to, for example, being a condition that is often normatively relevant). We use "systematic *normative* inquiry" for two reasons. First, we are sympathetic to the idea that much of epistemology is normative in at least some sense of "normative". Second, we think that the analogy to "systematic normative ethics" is illuminating. However, it is worth emphasizing that the case we make for the significance of our framework does not depend on assuming the normativity of epistemology, let alone of any particular epistemic category. So, philosophers who disagree with us on the aptness of using the term 'normative' here can still accept the core of our arguments in what follows.

We now offer two connected clarifications of our inquiry-focused account of systematic normative epistemology, which also apply to the two other projects we introduce in the following sections. The first clarification is that we think of these projects as *collective* inquiries. This is important for the following reason. Suppose that Bakhita writes an incisive essay on the generality problem for reliabilism.[21] Bakhita may have as her goal in the essay *solving the generality problem*, not *identifying and explaining in general terms, what it takes to be (e.g.) knowledge*. And yet, plausibly, Bakhita has made a contribution to systematic normative epistemology. This is possible, we suggest, because systematic normative epistemology is a *collective* theoretical project: the kind of scope and unity suggested by our discussion of simple reliabilism need not be exemplified in the work of a single individual, but can instead be characteristic of a theoretical pattern that emerges from the work of many. An individual, like Bakhita, in turn, can count as doing systematic normative epistemology in virtue of the relationship of her work to this collective project. Because of this, someone could contribute substantively to this project without thinking of herself as contributing to it, or even conceiving of the project as such.

The second clarification arises because all three of the projects that we discuss have an important feature: they first and foremost classify *inquiries*, understood as sustained and ambitious patterns of collective investigation. They do not, in the first instance, classify specific claims, arguments, questions, theories, etc. This means that our characterizations

[21]Put roughly, this problem is that we can think of belief-forming mechanisms at varying levels of generality. For example: did Loren form the belief that there is a tree outside his window using *sense perception* or by *looking out his window at an object about forty feet away on a clear day with his glasses on while sober*, etc.? At least initially, the challenge for the reliabilist is to identify an appropriate level of generality in a non-ad hoc way that preserves the alleged explanatory power of the theory. Compare (Conee and Feldman 2004, 136ff).

of these projects leave open what it takes to count as a systematic normative epistemic "claim" (or, similarly, an epistemic "question", "issue", etc.). We take this to be a virtue of our view. This is because claims can bear a number of interesting relations to an inquiry, and we think that it is a context-sensitive matter *which* of those relations one can or should aim to highlight by using the label "systematic normative epistemic claim". For example, a claim might be *part of the content* of the correct account of what it takes to satisfy knowledge. Or it might be a central (or marginal) part of an *argument* for that account. Or it might bear one of these relations to some incorrect competitor to the correct account, etc. We think there are sensible uses of 'normative epistemic claim' that align with each of these options, and that none of them is always more theoretically useful than the others. For similar reasons, we endorse a context-sensitive account of what counts as a "normative epistemic question" or a "normative epistemic theory", etc.[22]

Given their content, and given *most* kinds of relations to systematic normative epistemic inquiry, certain theses, such as our example of simple reliabilism, will be sensibly classified as "normative epistemic theses" in most contexts. This is precisely what makes this a useful exemplar. In other cases, it may be useful to treat a claim as part of systematic normative epistemology in some contexts, but not in others. (We will offer some examples of such cases later in this paper.)

3. Metaepistemology

We call the second project we aim to distinguish "metaepistemology" (or, equivalently, "metaepistemic inquiry"). We can characterize it as follows:

> **metaepistemology** is that inquiry which aims to explain how actual epistemic thought and talk – and what (if anything) that thought and talk is distinctively about – fits into reality.[23]

In what follows, we briefly unpack what is distinctive of metaepistemology, on this account, and draw out some of its implications. We also locate metaepistemology within metanormative inquiry more broadly.

[22]For a more detailed defense of a parallel context-sensitivity claim (about what counts as a "metaethical claim"), see (McPherson and Plunkett 2017).

[23]We develop a parallel account of metanormative inquiry in (McPherson and Plunkett 2017), where our focus is ultimately on meta*ethics* in particular. See also (Plunkett and Shapiro 2017), which develops the same basic account of metanormative inquiry, where the focus is ultimately on meta*legal* inquiry in particular. For more detailed discussion of the relevant moving parts in our account of metaepistemology below, see these other papers.

As we understand it, metaepistemology aims to explain *actual* epistemic thought and talk, and what (if anything) it is distinctively about. Metaepistemic inquiry thus aims to explain *descriptive* facts about how our actual epistemic thought and talk work, rather than to evaluate that thought and talk, or to advocate for alternatives to it.[24]

There are different ways to understand what "reality" amounts to. For our purposes here, we can take 'reality' to mean, roughly, the *totality of what there is*. Here, as elsewhere, our account of metaepistemology is intentionally schematic, leaving it entirely open *what* reality comprises. For example, it takes no stand on whether God exists, whether some form of physicalism is true of our actual world, or whether the real is identical to the fundamental. One reason that this neutrality is important is that competing metaepistemic theories sometimes differ in how they understand the nature of reality.

Epistemic thought and talk can appear to be distinctively *about* certain things, such as epistemic facts, properties, or relations. For example, thought and talk about (propositional) knowledge arguably appears prima facie to be thought and talk about a certain relation – the knowledge relation – that can hold between a thinker and a content.

On some views, there are such epistemic facts, properties, and relations. In these cases, understanding how these things fit into reality is part of metaepistemic inquiry. On other views, epistemic thought and talk is about such things only in the thin (intensional) sense that 'Pegasus' is about a winged horse.[25] On still other views, epistemic thought and talk is not about anything in even that thin sense.[26] For brevity, we will sometimes obscure these possibilities below, by glossing metaepistemic inquiry as the project of explaining how *epistemic thought, talk, and reality* fit into reality.

Next, consider our talk of "fitting in". The core idea here is that metaepistemic inquiry aims to explain how epistemic thought, talk, and reality *relate* to reality, including to specific salient aspects of it. For example: what kind of mental state is the thought that something is knowledge? What contribution do epistemic words make to the meanings of sentences? Can we offer conceptual analyses of central epistemic concepts

[24]It also contrasts with inquiry concerning the epistemic thought and talk of fictional or merely possible linguistic communities, such as Tolkien's elves, or the denizens of some epistemic twin earth. Such inquiry might of course be highly relevant to metaepistemology (just as, as we discuss later on, work on systematic normative epistemology and the conceptual ethics of epistemology might be relevant to metaepistemology).

[25]For example, consider the sort of epistemic error theory discussed in (Streumer 2017) and (Olson 2018).

[26]This would be true on an anti-realist version of epistemic expressivism discussed in (Chrisman 2012).

in non-epistemic terms? What is the distinction between epistemic and non-epistemic sentences? Do some or all epistemic sentences purport to be *about* a distinctive class of properties? Are any such properties instantiated? If so, how do we come to know about them? Etc.

We can think of metaepistemic inquiry as a branch of *metanormative* inquiry. G*eneral* metanormative inquiry concerns what is true of *all* of our normative thought and talk, from moral talk to thought about mafia norms. The thought animating our claim that metaepistemic inquiry is a distinctive branch of metanormative inquiry is that epistemic thought, talk, and reality comprises an apparently distinctive and unified subset of normative thought, talk, and reality, which is worth organizing a branch of metanormative inquiry around.

We have thus far characterized metaepistemology in very general terms. We can illustrate our account by noting that there are characteristic *research programs* in metaepistemology: collections of views that exemplify distinctive strategies for addressing the central metaepistemic task. Examples of such research programs include epistemic expressivism, epistemic error theory, and various forms of epistemic realism.

Briefly consider how one such research program – epistemic expressivism – addresses the explanatory challenge characteristic of metaepistemic inquiry. On one gloss, metaepistemic expressivism can be understood as the conjunction of three claims: (i) at the most basic explanatory level, epistemic judgments are some kind of non-cognitive attitude (e.g. desires or plans); (ii) epistemic statements consist of expression of the relevant non-cognitive attitude, rather than (e.g.) of the belief that one has that non-cognitive attitude; and (iii) the meaning of those epistemic statements is to be explained in terms of such expression.[27]

In light of these commitments, epistemic expressivism constitutes a distinctive kind of strategy for addressing part of the central explanatory challenge characteristic of metaepistemic inquiry: it offers a distinctive kind of account of how epistemic *thought and talk* fit into reality. In brief, on this account, epistemic thought consists of non-cognitive states, and epistemic talk consists of speech-acts which serve to express such states. However, the expressivist thesis by itself does not address a second part of the metaepistemic project: to explain how

[27]For a more detailed discussion of epistemic expressivism, see (Chrisman 2012). For an example of a view at least very close to expressivism concerning certain central kinds of epistemic thought and talk, see (Field 2000) and (Field 2018). See also (Gibbard 1990), whose expressivist account of 'rational', while focused on applications to ethics and moral philosophy, also applies to core parts of epistemic thought and talk.

what such thought and talk is about fits into reality. There are two central species of the expressivist research program, which address this challenge in different ways. On an "anti-realist" way of developing epistemic expressivism, epistemic thought and talk is not distinctively about *anything* (even in just a minimal, "intensional" sense of 'about'). Consider a familiar analogy: cheering "hooray!" expresses some positive attitude towards its object, but arguably does not make a claim about that object. Other expressivists are "quasi-realists", arguing that we can accept expressivism while still making sense of the thought that there are epistemic facts and properties which epistemic thought and talk are (in some sense) about. On a quasi-realist form of epistemic expressivism, the question of how epistemic facts and properties fit into reality remains intelligible.

4. The conceptual ethics of epistemology

This section introduces our third central project, which we call the *conceptual ethics of epistemology*.[28] On a first gloss, the core of this project involves normative inquiry about epistemic thought and talk. For example, the question of whether we ought to use epistemic concepts (and why) is a question in the conceptual ethics of epistemology. This example is potentially misleading however, because we understand the conceptual ethics of epistemology expansively, in each of three dimensions.

First, the term 'conceptual' here should be understood broadly, to signal focus on a cluster of related topics concerning thought and talk. There isn't a theory-neutral way to identify precisely *which* topics those are.[29] For our purposes here, we will take conceptual ethics to include the assessment of concepts, words, and pairings of concepts with words.[30]

Second, the term 'ethics' in 'conceptual ethics' is intended only to convey that conceptual ethics is a branch of normative inquiry. It is *not* intended to suggest that conceptual ethicists can engage in only moral

[28]The account we develop here draws on the account we develop for the conceptual ethics of *ethics* (as well as for the conceptual ethics of *normativity* more generally) in (McPherson and Plunkett 2021).

[29]One's views on this question will depend, in part, on which sorts of entities one countenances in one's theory of thought and talk, and what work those entities play in that theory. For example, some philosophers are skeptical that work in conceptual ethics (or connected work in conceptual engineering, which we will discuss below) really is about *concepts* as such. See (Cappelen 2018). For further discussion, see (Burgess and Plunkett 2020) and (Cappelen and Plunkett 2020).

[30]Our account of "conceptual ethics" here draws from (Burgess and Plunkett 2013a) and (Burgess and Plunkett 2013b).

or political evaluation of thought and talk.[31] Rather conceptual ethics can involve assessing concepts (etc.) against any number of different norms or standards, including the promotion of social justice,[32] epistemic fruitfulness,[33] or metaphysical joint-carvingness.[34]

So understood, work in conceptual ethics happens throughout philosophy, ranging from work in the philosophy of race to fundamental metaphysics.[35] The conceptual ethics of *epistemology* is the subset of conceptual ethics that focuses specifically on issues about epistemic thought and talk.

We now illustrate a final way in which our characterization of the conceptual ethics of epistemology is expansive: it concerns the assessment of both actual epistemic thought and talk and possible alternatives to that thought and talk.

The first strand of the conceptual ethics of epistemology focuses on epistemic words and concepts that are (or have been) in use. It asks questions like: is this concept *defective* in some way? Is it *good to use*? Consider, for illustration, a familiar view that can be read as a contribution to the conceptual ethics of *morality*: you might think that Friedrich Nietzsche's genealogy of morality involves a critical evaluation of the use of the modern concept MORALITY, and related concepts.[36] Some philosophers, such as Kevin Scharp, argue that our existing epistemic concepts might similarly be defective (but for different reasons) in ways that make them bad (or at least suboptimal) to use.[37]

A second strand of the conceptual ethics of epistemology begins from the initially plausible thought that the actual epistemic concepts people employ (or have employed in the past) are not the only possible ones we could use. There might be *alternative* epistemic concepts that we could use instead, giving rise to different kinds of epistemic (or at least "epistemic-ish") thought and talk.[38] This strand of the conceptual ethics of epistemology investigates and evaluates these alternatives. For example, we can ask: are any of these alternative concepts *better* than our current concepts, and *should* we therefore be using them instead?

[31]For more detailed discussion of this issue, see (Burgess and Plunkett 2013a) and (Burgess and Plunkett 2013b).
[32]See e.g. (Haslanger 2000b).
[33]See e.g. (Pérez Carballo 2020).
[34]See e.g. (Sider 2011).
[35]For further discussion, see (Burgess and Plunkett 2013a) and (Cappelen and Plunkett 2020).
[36]See (Nietzsche 1887/1994). For a brief proposal on how Nietzsche's arguments here can fruitfully be read as ones in conceptual ethics, see (Plunkett 2016).
[37]See (Scharp 2020) and (Scharp Manuscript).
[38]We say "at least epistemic-ish" to signal that, on some views of categorizing what counts as "epistemic", this category might be defined in terms of our current concepts. For connected discussion of some of the delicate issues here about topic and concept individuation, see (Cappelen 2018).

This last set of questions – questions about the reform or replacement of our current concepts – brings up an important thought: if we think that we should reform (or replace) our current concepts, how do we actually go about doing that? This thought suggests a connection between conceptual ethics and *conceptual engineering*. As we see it, paradigmatic projects in conceptual engineering draw on work in conceptual ethics to improve on the concepts (or other broadly "representational" or "inferential" devices) that we use, either by making new ones, or reforming the ones we have, and then trying to implement the use of those new (or reformed) concepts in practice.[39]

Our discussion so far has emphasized the *diversity* of the conceptual ethics of epistemology. Despite this diversity, we take the conceptual ethics of epistemology to be a unified project in the following sense: whichever epistemic words or concepts one examines, and whichever standards one uses in doing the conceptual ethics of epistemology, one is engaged in the assessment of epistemic words and/or concepts.[40] This underwrites a clear contrast with both the project of systematic normative epistemology, and the project of metaepistemic inquiry, as we now explain.

5. The depth of the distinctions between these projects

Thus far, we have sketched three projects (systematic normative epistemology, metaepistemology, and the conceptual ethics of epistemology), and characterized each project in terms of its distinctive aims or success conditions. In this section, we emphasize the depth of the distinctions between these projects in two ways: by explaining the significance of the contrasts between their aims, and by contrasting the characteristic motives for engaging in each project. We begin by illustrating the significance of the contrast in aims via an example.

Recall our exemplary theory in systematic normative epistemology: simple reliabilism, according to which knowledge consists in reliably formed true belief. This theory has some claim to be the outline of a complete systematic normative theory of knowledge. This is because, if it were

[39]Our use of the term 'conceptual engineering' in this way draws broadly on the uses put forward by Scharp (2013), Eklund (2017), and Cappelen (2018). Our gloss above is only a rough characterization of how we think of conceptual engineering, and its relation to conceptual ethics. For further discussion of this issue, see (Cappelen and Plunkett 2020) and (Burgess and Plunkett 2020). For a sample of the variety of ways in which different philosophers use the terminology of 'conceptual ethics' and 'conceptual engineering', see the papers collected in (Burgess, Cappelen, and Plunkett 2020).

[40]For defense of the idea that conceptual ethics is a unified branch of normative inquiry, see (Burgess and Plunkett 2013a).

correct, it would fulfill (in outline) the aim of systematic normative epistemology, at least for knowledge: it would identify and explain in general terms what it takes to have knowledge.

By contrast, even if correct, simple reliabilism would not be anything like a complete metaepistemic theory. One reason for this is that simple reliabilism says nothing about how epistemic thought and talk fit into reality. For example: is the thought that I have knowledge a belief? Or is it rather a desire-like or plan-like state (as an epistemic expressivist would argue)? Simple reliabilism is also silent on many important questions about how "epistemic reality" fits into reality. For example: does the relation between knowledge and reliably-formed true belief mentioned in simple reliabilism constitute a metaphysical reduction of knowledge, or does it mark an informative relationship between distinct properties, or does it have no metaphysical implications whatsoever? How (if at all) can we come to know that simple reliabilism is true, if it is? Simple reliabilism itself is silent on these questions. But these are central questions that naturally arise in metaepistemic inquiry.

Finally, simple reliabilism fails to address questions at the heart of the conceptual ethics of knowledge. For example, suppose simple reliabilism were true. This would not by itself settle whether we should use the word 'knowledge' and/or the concept KNOWLEDGE in central ways in our practices of belief regulation and interpersonal evaluation. Nor does it speak to the possibility of fruitful reform of the concept, or whether it would be better to replace it with some alternative concept.

We can step back from this example to consider the depth of the contrasts among the three projects in general terms. Note first that systematic normative epistemology and the conceptual ethics of epistemology are (on their face) both parts of *normative* inquiry. They are nonetheless distinct, because they address *different* normative questions. The normative epistemologist asks, for example, what it takes for a person to be epistemically rational, or for a belief to be epistemically justified. The conceptual ethicist asks (e.g.) whether it is good to use certain epistemic words (e.g. 'epistemically rational') or concepts (e.g. EPISTEMICALLY JUSTIFIED) in certain ways.[41] In contrast to both of these projects, metaepistemology is on its face a *descriptive* project: it aims to explain how a branch of normative thought, talk, and

[41]It is worth noting that *some* ways of engaging in the conceptual ethics of epistemology can be usefully classified, in at least some contexts, as a part of *applied* epistemology. This would be so if one *used* epistemic norms in that project. For example, a conceptual ethicist might ask which epistemic concepts are such that their widespread use in inquiry maximizes certain broadly "epistemic" goods, such as the spread of knowledge.

reality fits into reality. The aim of metaepistemology, then, insofar as it relates to issues of thought and talk, is not to prescribe or evaluate (or to identify what counts as correct prescription or evaluation). Rather, it is to explain descriptive facts about a part of thought and talk (namely, epistemic thought and talk) that arguably involves some prescription or evaluation.

Next, consider the relation between central *motivations* a person can have for engaging in each of the three projects. Engagement with each of the three projects is in part motivated by the intuitive *importance* of epistemic categories. We *care* about whether we (or our interlocutors, testimonial sources, etc.) know certain claims, whether our beliefs are *justified* or *rational*, whether we *understand* certain topics, etc. Further, epistemic categories appear important in practice: they can play characteristic important roles in our reasoning, and in our evaluation of other people (e.g. as trustworthy or not).

Given this observation, one characteristic motivation for engaging in the project of systematic normative epistemology is to better understand these seemingly important epistemic conditions. A historically influential (but more controversial) motivation for engaging in systematic normative epistemology has been to understand whether (and why) these epistemic conditions are attainable. This motivation has loomed large, for example, in discussions of skeptical challenges. Another controversial motivation for engaging in systematic normative epistemology is that it might be *practically useful*, in providing guidance for inquiry and the management of belief.

Next, consider motivations for engaging in metaepistemology. Our epistemic *thoughts* seem to play distinctive roles in our reasoning and thinking, which contrast with the significance of our thoughts about (at least intuitively) "less normatively important" norms, such as the rules of a board game, or the legal code of a long-dead empire. One central motivation for engaging in metaepistemology is to better understand what in the nature of epistemic thoughts (or possibly epistemic reality) allows these thoughts to play these distinctive roles. Alternatively, a metaepistemic project could be motivated precisely as seeking to debunk these apparently distinctive roles of epistemic thought and talk. Notice that, although this motivation begins with the intuitive importance of epistemic categories, it is quite different from the canonical motivations for systematic normative epistemology.

Engagement in the conceptual ethics of epistemology can also be motivated by reflecting on the key roles of epistemic words and concepts in the thought and talk that we use to regulate our beliefs, interpersonal assessment, and inquiry. However, here the leading thought is that, given

these roles, it becomes important to investigate whether the words and concepts that play these roles are optimal, and (if not) whether they could then be improved in some way, or replaced, in order to better serve these (or other) roles.

Recall our initial foil: the idea that while (as we saw in §1) epistemologists sometimes *describe* their projects differently, all of them might nonetheless be engaged in substantively the same project. We take the important contrasts we have marked here to render this idea very implausible. We have argued that the distinctions between the success conditions of these projects are deep, where this depth is illustrated by the contrasts among the characteristic motivations people have for engaging in each of the three projects. In our view, these facts constitute a strong prima facie case for the theoretical importance of attending to these distinctions. The next section bolsters this initial case.

6. The methodological value of distinguishing the projects

In this section, we illustrate several ways that attention to the distinctions between the three projects we discuss can be methodologically valuable. These include improving our interpretation of claims in epistemology, illuminating interesting theoretical options that might otherwise remain invisible, clarifying the sorts of arguments that can support claims, and highlighting certain important kinds of unreliable inference. We take these methodological payoffs to be relevant across many central topics in epistemology. To make the discussion manageable, however, we will organize our discussion around a single topic in epistemology: encroachment.

Encroachment theses concern a distinctive way that certain non-epistemic norms can (allegedly) bear on the epistemic, encapsulated in the following schematic thesis:

> **Encroachment**: Facts about a person's *non-epistemic normative circumstances* (e.g. the practical or moral stakes of decisions that they might make on the basis of a belief) can affect (e.g.) how much evidence is required for them to count as having knowledge or epistemically rational belief.

One important question about encroachment which has recently begun to receive sustained attention is *which norms* might encroach on the epistemic in this way. For example, are they prudential norms, moral norms, or any norm the agent happens to care about?[42] We want to focus on a

[42]For discussion of some of the most well-known discussions about "pragmatic" encroachment, see (Fantl and McGrath 2009) and (Hawthorne and Stanley 2008). See (Bolinger 2020) for helpful

different question, however: which *project* we are engaged in when we are considering encroachment theses. For example, one might consider an encroachment thesis as a claim in systematic normative epistemology, about the conditions for (e.g.) knowledge. Or one could introduce it as a metaepistemic claim, for example, as a putative conceptual truth about the concept KNOWLEDGE, or as a claim about the metaphysical grounds of the knowledge relation. Finally, one could introduce an encroachment thesis as a conceptual ethics claim about the epistemic (or at least "epistemic-ish") concepts that we *should* be using, or the property that those concepts pick out.[43]

Next, consider why we might *care* about a claim like Encroachment. This can again vary depending on which project is our focus. For example, if we focus on Encroachment as a systematic normative thesis, we might think that it simply helps us to better understand the conditions on knowledge. If we think of it as a metaepistemic claim about our normative concepts, we might think that it would help illuminate how our epistemic thought works: for example, why people take certain inferences from knowledge to the practical permissibility of action to be obvious or intuitive. If we think of it as a claim in conceptual ethics, then we might understand it as a claim about concepts that it would be in some way *good* to use to regulate our beliefs.

Next, consider how locating an encroachment thesis within each project can affect how it makes sense to argue for (or against) that thesis. What is arguably the standard form of argument for encroachment theses begins with a claim like the following:

Conditional: If S knows that P, then S can rationally proceed as if P.[44]

general discussion and overview of the different varieties of "moral encroachment" theses in recent epistemology, and a number of important ways in which they differ from each other (and from many of the classic "pragmatic" encroachment views). See (Fritz 2017) and (Fritz Forthcoming) for discussion of issues involved in distinguishing the kinds of norms that (purportedly) "encroach".

[43]One might wonder whether the *same* claim could really play the relevant sort of role in each project. For example, one might think that Encroachment is a different claim from the claim that Encroachment is a conceptual truth, or ought to be. We want to make three points. First, a thesis like Encroachment might play an important role in metaepistemology even if it was not itself put forward as a conceptual truth. For example, the truth of Encroachment might be a central premise in an argument for the conclusion that it *is* a conceptual truth. Second, what we should say about this issue will depend importantly on how we individuate claims. For example, if claim-identity can be preserved over changes in semantic content, the conceptual ethicist might want to argue that they are arguing for the truth of Encroachment itself. (For connected discussion on this kind of point, see (Cappelen 2018) and (Plunkett 2015)). Third, even if the claims being made in the different projects are strictly different, it may be that inquirers are not sharply distinguishing them in their reasoning. Thanks to Pekka Väyrynen for pressing us on this question.

[44]For example, see (Fantl and McGrath 2009) and (Hawthorne and Stanley 2008).

This sort of claim is supposed to support encroachment theses for the following reason: sometimes (e.g.) the practical stakes that face an actor can affect how much evidence they need in support of a proposition P, in order to rationally proceed as if P.

Consider the standard sort of case used to illustrate this idea: suppose that it is Friday, and I want to make a deposit in person at the bank before Sunday. If I know that the bank will be open Saturday, then Conditional suggests that I can rationally proceed on this basis, waiting until Saturday to make my deposit. Now consider two variants of the case. In the first, failing to make a deposit before Sunday would only be a small inconvenience. Here it would not take much evidence to warrant waiting to make my deposit. In the second, the bank will repossess my house unless I make a payment at the bank before Sunday. Given these stakes, I would need to have extremely good evidence that the bank is open in order to wait until Saturday. And according to Conditional, this means that I would need extremely good evidence in order to count as knowing the bank is open.

Crucially, which project we are engaged in affects how we should understand Conditional, and thus what sort of evidence is apt to support it. If we are engaged in a conceptual ethics project, one exemplary way to understand Conditional is as a claim about the knowledge (-ish) concept it would be good for us to use. To defend Conditional understood in this way would require normative argument, for example, perhaps an argument that given the sorts of creatures we are, it would be *valuable* to be able to rely in our practical deliberation on the truth of theses that we treat as known. Suppose instead that we are considering Conditional as a metaepistemic thesis about the concept KNOWLEDGE. This would require very different evidential support: whichever sort of evidence is required to support conceptual truths (perhaps introspection via the method of cases, or perhaps well-constructed x-phi experiments on a representative sample of those competent with the concept KNOWLEDGE). Finally, suppose that one seeks to defend Conditional as a claim in systematic normative epistemology. Here one would need to use whichever sorts of evidence are appropriate to that project. What sort of evidence that is depends on what sort of metaepistemic view is correct. For example, if knowledge is a natural kind, empirical investigation into that kind might be called for. By contrast, if knowledge is a non-natural property, perhaps the only direct evidence for this sort of claim would be some sort of intuition or rational insight.

Similar points arise for two other important ways of seeking to defend encroachment theses. First, some have sought to argue that an encroachment thesis is made true by facts about the *function* of epistemic discourse (à la Craig).[45] Here, we can helpfully distinguish *metaepistemic* versions of this view, which claim that the relevant function is realized by our actual epistemic discourse, and conceptual ethics versions of the view, which claim that it would be *good* (in some way) for our epistemic discourse to realize the relevant function. Second, Jamie Fritz has recently argued that denying certain encroachment theses can commit the epistemologist to the rationality of a distinctive sort of akrasia.[46] Fritz treats this as an argument for encroachment. But again, we can ask: how is the argument from this conclusion to encroachment supposed to go? If it rests on the idea that it would be *bad* for central epistemic concepts to license this sort of akrasia, then the thesis is most clearly a conceptual ethics argument (and would need to be defended as such). If it instead rests on the idea that it is a conceptual truth about RATIONALITY that rationality does not license such akrasia, then evidence for this alleged conceptual truth would need to be adduced. Finally, suppose the function and akrasia arguments are intended to support an encroachment thesis in systematic normative epistemology. As with the argument from Conditional, how to complete these arguments will depend on what the correct metaepistemic theory is.[47]

Notice three kinds of contexts where attending to the sorts of possibilities that we have just been emphasizing can be valuable. First, when one interprets another's epistemic work, identifying which project that person is engaged in is crucial to interpreting her work charitably. This is because doing so will help us to better understand her claims, arguments, and underlying motivations. Because the distinctions we have been exploring here are not part of the standard toolbox in contemporary epistemology, the interpretive questions here may sometimes be quite challenging. Further, even in the comparatively rare case where someone is explicit that she is engaging in one or another of these projects, this does not settle the matter. As we emphasized in §2, which of these projects a piece of work contributes to is in our view less a matter of authorial intention (which may sometimes involve a misunderstanding of one's own

[45]For example, see (Hannon 2019).

[46](Fritz 2021). For examples of other recent discussion on epistemic akrasia, see (Horowitz 2014) and (Greco 2014).

[47]For argument that the true metanormative theory will play a central role in explaining what sorts of evidence and argument we need to support a normative claim, see (McPherson 2012). For connected discussion, see (Darwall 1998).

work), and more a matter of how that work, in fact, fits into the broader social activity of inquiry.[48]

Second, when one is making one's own arguments concerning topics in epistemology, being clear for oneself about which project one is engaged in will make it easier to ensure (i) that one's claims reflect one's central motivations in engaging in the relevant inquiry, and (ii) that one's arguments are apt to support the claims that one is making. Further, making clear to one's audience which project one is engaged in makes it easier for one's audience to charitably interpret one's claims.

Finally, it might well be that for many topics in epistemology, no one has in fact ever engaged in discussion of that topic as a thesis in (e.g.) the conceptual ethics of epistemology. This helps to bring out a third way that explicit attention to the distinction among these projects can be valuable: it can help to illuminate interesting and underexplored sorts of inquiry. For example, even if no one ever has explicitly and systematically examined Encroachment as a thesis in the conceptual ethics of epistemology, we think it would be extremely interesting to do so.

So far we have focused on the constructive value of attention to the distinctions we have advertised. There is a critical flipside to this, however: failure to attend to these distinctions raises the danger of making bad arguments. Many of the relevant vices are just the flipsides of the virtues we have advertised. For example, one might easily misinterpret another philosopher's work if one implicitly miscategorizes the project it mainly contributes to.

Another danger of ignoring the distinctions we are highlighting is that doing so makes one vulnerable to certain kinds of *unreliable inferences*. To appreciate the kinds of unreliable inferences we have in mind, consider that this section has highlighted the differences in evidence and argument that may be relevant depending on which project one is engaged in. If a philosopher isn't sensitive to these differences, she might well implicitly switch between the kinds of evidence that are relevant to one

[48]One way to argue for the idea that some philosophical work may be best understood as *implicitly* engaged in conceptual ethics, even absent authorial intention, relies on the idea of "metalinguistic negotiation" discussed by (Plunkett and Sundell 2013a). In a metalinguistic negotiation, speakers put forward views about how terms should be used – including, for example, putting forward views about which concept(s) a term should be used to express, in a given context – by (at least seeming to) use, rather than mention, those very terms. On this picture, speakers might well be unaware that this is what they are doing, even if this is in fact what they are doing. (See (Plunkett and Sundell 2021) for an extended discussion of this point). If philosophical work involves more metalinguistic negotiation than most philosophers appreciate, then that means there is more *tacit* argument about conceptual ethics going on in the field than most appreciate. (See (Plunkett 2015)). For closely connected ideas, see (Ludlow 2014) and (Thomasson 2016).

project but another, and thus end up making unreliable inferences based on a given claim within epistemology. For example, a philosopher might implicitly make an inference from a claim in conceptual ethics about how we should use the term 'knowledge' to a normative claim about what the conditions on knowledge are. The possibility of such "project-switching" unreliable inferences may exacerbate the danger of other sorts of unreliable inference. For example, it is a familiar fact that terms that are widely used in epistemology, such as 'rational', are frequently used in ways that associate them with widely divergent sorts of ideas, including coherence, normative significance, epistemic justification, and epistemic virtue. In other work, we've explored how philosophers might well tacitly switch from some such associations to others, leading to bad forms of argument or unreliable inference using a term such as 'rational'.[49] Some of these associated ideas may be more salient when an epistemologist is implicitly engaging in (e.g.) a conceptual ethics project than they would be when they are implicitly engaged in metaepistemology, meaning that implicit shifts across projects may also raise the possibility of unreliable inferences that arise from the multiple ideas associated with epistemic terms. Insofar as the threefold distinction between projects we have discussed is on point, it can help illuminate important differences between kinds of claims, thereby helping to make salient and reduce unreliable inferences of these kinds.

In this section, we laid out our central case for the methodological importance of attending to the contrasts between the three projects we have introduced. The heart of the case is the insight that which project a thesis contributes to can affect both the significance of that thesis, and what sorts of arguments are apt to support that thesis. This in turn makes attending to the contrast between the projects crucial to properly interpreting others' ideas, clarifying one's own arguments, avoiding unreliable inferences, and illuminating neglected theoretical options.

7. Complications and clarifications

In making our case for the importance of distinguishing systematic normative epistemology, metaepistemology, and the conceptual ethics of

[49]See McPherson and Plunkett (2020) for further discussion of related issues. Note that this issue here might also arise for the category of the "epistemic" itself. For example, some might take "epistemic" conditions to necessarily entail normative facts about what one "really and truly" should believe, whereas others might take them to be tied to constitutive standards of belief where there isn't necessarily such entailment. We take up these issues in McPherson and Plunkett (Manuscript).

epistemology, we have presented these projects (and the distinctions between them) in a streamlined way. In doing so, we have largely ignored several possibilities that might be seen to complicate or undercut the case we have made. In this section, we discuss four such complications. These arise from the possibility of overlapping content among the projects, of methodological entanglement among the projects, of metasemantic hypotheses that complicate the distinctions among the projects, and of certain hypotheses on which meaningful distinctions among the projects might seem to disappear. In each case, we suggest that the possibility in question does not undercut the positive case we have made for the depth and methodological significance of the contrasts that we have made between the projects. However, in many cases, these possibilities enrich and complicate the basic story we have put forward thus far.

The first complication is a point that we already made in introducing our projects: our account does not entail that the sets of claims (or even the *true* claims) that are appropriately classified as belonging to each of these three projects are wholly distinct from each other. As we have emphasized, on our view it is a context-sensitive matter whether it makes sense to classify a given claim as a "normative epistemic claim", a "metaepistemic claim", or a "claim in the conceptual ethics of epistemology".

One might think that this vitiates the depth or importance of the contrast between the projects. But, as we have emphasized in the previous two sections, this is not so. We have explained the *depth* of the contrast between the projects in terms of contrasts in their aims and not in terms of specific types of claims that come up in pursuing those projects. And we have just explained the methodological significance of the distinctions among the projects in a way that emphasizes the very different role a single claim (like Conditional) might make within the three projects. We take it to be a virtue of our view that it can vindicate the plausible thought that the claims made within these three projects can overlap, while nonetheless also vindicating the idea that there are deep theoretical cuts between the projects.

The second complication is that the three projects are plausibly *methodologically* entangled, rather than distinct. For example:

- Some metaepistemic conclusions might affect the sorts of evidence that it makes sense to appeal to in systematic normative inquiry. For example, as we noted in the preceding section, the metaepistemic

thesis that knowledge is a natural kind might have very different implications for how to engage in systematic normative epistemology than a non-naturalistic metaphysics of epistemology would have. More generally, our best metaepistemic theorizing might decrease the plausibility of certain otherwise attractive normative epistemic theories, and vice versa.[50]

- Both metaepistemic and systematic normative epistemic theory could play a central role in informing arguments in the conceptual ethics of epistemology. For example, if our current concept KNOWLEDGE is in perfectly good working order, then there isn't any obvious motivation for conceptual ethics projects seeking to reform or replace it.[51] By contrast, such a project could be motivated by arguments that certain existing epistemic concepts are incoherent or defective, or that there are in fact many distinct concepts that philosophers pick out with terms like 'epistemic justification'.[52]
- Some conclusions in each project might have implications for our ability to cogently engage in one of the other projects. For example, skeptical views in systematic normative epistemology might serve to undermine the other projects, and metaepistemic error theory might undermine the project of systematic normative epistemology.
- Views in the conceptual ethics of epistemology that vindicate the use of at least some (even if not all) of our current epistemic concepts could serve to bolster the motivations for engaging in systematic normative epistemology and metaepistemology.
- Conversely, conceptual ethics views according to which our existing epistemic words and concepts are *bad to use* could undercut the motivations for engaging in the other projects.
- Arguably certain specific metaepistemic views and views in the conceptual ethics of epistemology are *presuppositions* of certain work in systematic normative epistemology. For example, some such work presupposes that knowledge is in some sense *more valuable* than true belief, others that the concept or property of *justification* is constitutively connected to truth, etc.

We again accept these possibilities while denying that they make trouble for our account. Notice first that these methodological connections do not cast doubt on our distinctions between the projects, which, as we

[50]For connected discussion, see (McPherson 2012) and (Darwall 1998).
[51]For further discussion of these general issues, see (McPherson and Plunkett 2020).
[52]Compare (Alston 2005) and (Cohen 2016).

have emphasized, are a matter of contrasting distinctive aims. Notice next that these methodological connections do not threaten the idea that these distinctions are theoretically deep and important. Finally, we also think that they illustrate the importance of attending to the distinctions we are discussing. This is because we can most clearly understand, investigate, and evaluate hypotheses like the ones just mentioned within a framework that clearly distinguishes the three epistemic projects.

The third complication we want to examine arises from the possibility that certain ambitious claims about mental or linguistic content could undercut the depth of the distinction that we have advertised between descriptive inquiry into our concepts and words (of the sort involved in metaepistemology) and normative inquiry into them (of the sort involved in the conceptual ethics of normativity).[53]

Suppose that facts about the meaning or content of our current epistemic words or concepts are partly grounded in normative facts. One might defend such a view either generally, or only for certain kinds of content. For example, one might accept a general metasemantic account that gives a central explanatory role to facts about rationality.[54] Or one might argue that epistemic concepts are among what Ronald Dworkin calls *interpretive* concepts. According to Dworkin, an interpretive concept is one whose content is determined by the best overall justification for our practices involving those concepts.[55] Views of either of these kinds might seem to muddy the contrast we drew between "descriptive" metaepistemic inquiry, and "normative" conceptual ethics inquiry.

A first response is just to note that such metasemantic theories are intensely controversial, so it is not clear how pressing this complication is. If we suppose that some such theory is correct, however, it is still possible to contrast the metaepistemic and conceptual ethics projects, in two ways. First, such metasemantic theories typically give a privileged role to

[53]The kinds of claims we are interested in here are often characterized as "metasemantic". Note that this use of 'metasemantic' doesn't line up with our own preferred way of talking about "meta-level" projects. This is because, on our way of thinking about things, "metasemantics" would most plausibly be understood as the project that aims to explain how actual *semantic* thought and talk – and what (if anything) it is distinctively about – fits into reality. Some philosophical work on semantics (or more broadly on meaning, or linguistic or mental content) fits smoothly into this project. (See, for example, (Gibbard 2012)). In contrast, the standard way of using 'metasemantics' concerns questions about what explains why particular words or thoughts have the meaning (or content etc.) that they do. (See for example the essays in (Burgess and Sherman 2014)). These explanatory questions – which often involve issues about the *grounds* of semantic meaning – are most naturally thought of as part of the explanatory component of something parallel to systematic normative epistemology, or systematic normative ethics, rather than at the core of a hermeneutic "meta-level" project.

[54]For example, consider the metasemantic views defended in (Williams 2018) and (Schroeter and Schroeter 2014).

[55]See (Dworkin 2011). For critical discussion, see (Plunkett and Sundell 2013b).

specific norms (such as rationality or Dworkinian practice-justification). By contrast, a conceptual ethics project might evaluate our practice using a variety of standards, from joint-carvingness to the promotion of social justice. Second, even holding fixed the norm in question, in these metasemantic theories such norms typically play a specific and somewhat constrained role. In contrast, in conceptual ethics, the relevant norm could play a much broader role, such as potentially being the main (or even sole) norm one uses to rank candidate concepts.

A final complication arises if we imagine a certain optimistic take on our current epistemic concepts. Consider the following hypothesis:

Convergence: our actual epistemic concepts are:
(i) apt to be investigated using "standard philosophical methodology" and
(ii) optimal with respect to any plausible standard that one might apply in conceptual ethics (e.g.: non-defective, pragmatically useful to use, authoritatively normative, joint-carving, useful to liberatory political projects, etc.).

If Convergence turned out to be true, it would arguably undercut much of the significance of our threefold distinction. In particular, it would suggest that epistemologists could simply deploy our current epistemic concepts to pursue systematic normative epistemology, without worrying much about issues stemming from the conceptual ethics of epistemology or metaepistemology potentially mattering for their work. Moreover, against one of our ongoing strands of argument, it might seem that they could advance their theses, arguments, etc. without worrying much about the distinction between the projects in general.

Our first response to this possibility is that we doubt Convergence is correct for the epistemic concepts currently employed by any living epistemologists. Our epistemic concepts might well be in fairly good working order, but it would be surprising if any contemporary philosopher (let alone all of us) employed concepts that were optimal along so many dimensions.

More importantly, consider that, in order to *establish* Convergence, we would need to engage carefully in each of the three projects we have sketched. We think that clearly distinguishing the projects would be valuable to investigating and evaluating Convergence, just as it is elsewhere in our investigation in (or about) epistemology. Of course, having established Convergence, the methodological interest of the distinctions between the

three projects might fall away – just as it might if we conclusively established the truth of some other currently highly controversial thesis in (or about) epistemology. But we suspect that establishing Convergence would require such extraordinary advances in each of the three projects that we have discussed that this is a somewhat distant worry.

The complications we have discussed in this section help to clarify the claims we make on behalf of the distinctions between projects in epistemology that we have advertised. We take those distinctions to be both theoretically deep and methodologically important. However, this does not mean that each of the projects is isolated from the others. To the contrary, we think that part of the value of attending to the distinctions we have introduced is that they help us to think more clearly about the sorts of overlap among claims and methodological entanglements among the projects that are possible. Further, as we have emphasized, there are certain hypotheses, such as the metasemantic hypotheses we have discussed and Convergence, that would at least to some extent mitigate the methodological importance of focusing on the distinctions. However, as we have also emphasized, this mitigation is plausibly not total, and the distinctions we have mapped are important for clearly investigating the plausibility of theses like Convergence.

8. Conclusion

In this paper, we've proposed a way of understanding three contrasting projects in epistemology: systematic normative epistemology, metaepistemology, and the conceptual ethics of epistemology. We have advertised the importance of the distinctions among these projects, and considered hypotheses that complicate those distinctions. In conclusion, we want to make four observations about what follows if our argument in this paper is successful.

First, if we grant that there is a deep and methodologically important distinction among these projects, this raises a natural question: should we think of these projects as *complementary,* or as *competitors*? Some of the rhetoric around conceptual ethics and conceptual analysis suggests that some philosophers think of these as competing projects.[56] On our way of thinking about things, it's often most natural to think of work in "conceptual analysis" as being pursued as part of a "meta-level" project (such as metaepistemology or metaethics). Thus, this rhetoric suggests a kind of

[56]See, for example, (Scharp 2020) and (Haslanger 2000b).

competition between metaepistemology and the conceptual ethics of epistemology. Furthermore, as we noted in the previous section, there are certainly *hypotheses* within these projects that imply tension. For example, it would be hard to motivate systematic normative epistemology if one accepted the conceptual ethics view that we should eliminate the use of epistemic concepts. Despite this, we think it is more plausible to (at least initially) think of all three projects as *complementary*. In §5 we identified central motivations for engaging in each project, which each strike us as prima facie reasonable and non-competing motivations for inquiry. Further, as we have emphasized at several points in this paper, work in each project can potentially illuminate inquiry in the other projects.

If we take seriously this idea that these projects can be complementary, rather than competitors, it suggests that we can understand the diversity of conceptions of epistemology mentioned in §1 as conceptions of different *parts* of epistemology. Consider some brief examples. Feldman's gloss is very similar to our account of systematic normative epistemology. The analysis of epistemic words or concepts (along the lines pursued by Jackson) is plausibly a central metaepistemic task. Similarly, both Craig's account of the function of epistemic concepts, and Williamson's stated project of studying the state of mind of knowledge, are naturally understood as projects in metaepistemology. Finally, Haslanger and Scharp frame their projects in ways that suggest they are primarily engaged in conceptual ethics projects – in their case, in the service of broader ones in conceptual engineering. This taxonomy comes with an important caveat, however. We have suggested that many of these *conceptions* can be plausibly classified. However, especially because explicit attention to these distinctions is rare, it is plausible that when we dig into the work of these epistemologists, we would find contributions to more than one of these projects. For example, we find it plausible that Williamson's work on "knowledge-first" epistemology involves contributions to all three projects we have discussed (even if he himself might well resist that characterization). If we are on the right track, then our framework should allow us to more carefully think about how the different parts of his work interact, as well as open up underexplored ways of further developing ideas in that work.

Second, we have argued in §5 that it can be useful for epistemologists to attend to the distinctions we have sketched here. In short, we aim to get our threefold distinction on the map – and underscore its importance – so that different philosophers, coming from a range of different

perspectives, can wrestle to some degree with our arguments, and then see where that takes them in future work. It does not follow from this, however, that every epistemologist needs to stop what they are doing and focus on attending to this threefold distinction. Far from it. There are obviously many other worthwhile things to do in epistemology, and any reasonable view about the social organization of inquiry should support a division of labor here, with different epistemologists attending to different issues.

That being said, we do think that this paper can help to focus epistemologists' attention on a range of important questions that we think have been comparatively neglected. This is true both of work in metaepistemology and the conceptual ethics of epistemology.

Start with the case of metaepistemology. This brings us to our third and fourth points. At the beginning of this paper, we noted that the distinction between metaethics and normative ethics has received more sustained attention than the parallel distinction in epistemology between metaepistemology and normative epistemology. One upshot of this, we think, is that the range of possible positions in metaepistemology, and their relative theoretical virtues and vices, are underexplored in comparison to those in metaethics. This has started to change in recent years, which is a development that we enthusiastically welcome.

Parallel points hold for the recent growth of work in the conceptual ethics of epistemology. This area is even less systematically explored than metaepistemology. We think there is rich terrain here for epistemologists, as we've discussed at length in this paper. Beyond the reasons for this that we have focused on, it is also worth noting that some topics in conceptual ethics themselves raise interesting epistemological challenges. Consider an example from our own work. In a recent paper, we explore what we call "the vindicatory circularity challenge" to the conceptual ethics of normativity.[57] The core issue is this: in conceptual ethics, we might hope to *improve* on our existing normative concepts. But when assessing whether some concept might count as an improvement, we inevitably deploy one of our own normative concepts. This raises an epistemological issue: given that our existing normative concepts might be imperfect (or perhaps even seriously defective), why should we think that a conceptual ethics project will help us to identify normative

[57]See (McPherson and Plunkett Forthcoming). This epistemological challenge is closely related to the *metaphysical* one discussed in (Eklund 2017), which we discuss in (McPherson 2020) and (Plunkett 2020).

concepts that are better in (to put it intuitively) the *most normatively important* ways?

We think this is an important epistemological problem that only clearly comes into view when thinking about the (relatively neglected) project of the conceptual ethics of normativity. It thus illustrates one way that explicit attention to these three projects can reveal new epistemological questions and puzzles. We expect attention to the distinctions we have advertised in this paper to frequently be fecund in this way. In §6 we illustrated how our theorizing about encroachment theses can change when we consider them as contributions to each of the three projects we discussed. In our view, this is the tip of the iceberg: explicit attention to these distinctions helps to create fertile soil for potentially fruitful new ways of developing existing views, or for developing new views altogether. We hope our work in this paper can help spur such developments.

Acknowledgements

We are indebted to audiences at the Vienna Unity of Normativity Conference, Leeds, the Arché research centre at the University of St. Andrews, Toronto, and the Oslo Conceptual Engineering Online Workshop for helpful comments on ideas in this paper. Thanks also to Alexis Burgess, Herman Cappelen, Jorah Dannenberg, Tyler Doggett, Rachel Fraser, Céline Henne, Jonah Hirsch, Nadeem Hussain, Nithya Kasarla, Zachary Lang, Amanda Li, Errol Lord, Kate Nolfi, Mark Pinder, Paul-Mikhail Catapang Podosky, Joey Pollock, Adrian Russian, Scott Shapiro, Kevin Scharp, Dan Singer, Declan Smithies, Tim Sundell, and Pekka Väyrynen for helpful discussion of ideas in this paper.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

Alston, William. 2005. *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.

Berker, Selim. 2013. "Epistemic Teleology and the Separateness of Propositions." *Philosophical Review* 122 (3): 337–393.

Bishop, Michael A., and J. D. Trout. 2004. *Epistemology and the Psychology of Human Judgment*. New York: Oxford University Press.

Bolinger, Renée Jorgensen. 2020. "Varieties of Moral Encroachment." *Philosophical Perspectives* 34 (1): 5–26.

Burgess, Alexis, Herman Cappelen, and David Plunkett. 2020. *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
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.
Burgess, Alexis, and Brett Sherman. 2014. *Metasemantics: New Essays on the Foundations of Meaning*. Oxford: Oxford University Press.
Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
Cappelen, Herman, and David Plunkett. 2020. "A Guided Tour of Conceptual Engineering and Conceptual Ethics". In *Conceptual Engineering and Conceptual Ethics,* edited by A. Burgess, H. Cappelen, and D. Plunkett, 1–26. Oxford: Oxford University Press.
Chrisman, Matthew. 2012. "Epistemic Expressivism." *Philosophy Compass* 7 (2): 118–126.
Cohen, Stewart. 2016. "Theorizing About the Epistemic." *Inquiry* 59 (7–8): 839–857.
Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
Craig, Edward. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press.
Darwall, Stephen L. 1998. *Philosophical Ethics*. Boulder: Westview Press.
DeRose, Keith. 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context*. Oxford: Oxford University Press.
Descartes, René. 1641/2017. *Meditations on First Philosophy*. Cambridge: Cambridge University Press.
Dworkin, Ronald. 2011. *Justice for Hedgehogs*. Cambridge: Harvard University Press.
Eklund, Matti. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.
Fantl, Jeremy, and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
Feldman, Richard. 2003. *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
Field, Hartry. 2000. "A Priority as an Evaluative Notion." In *New Essays on the A Priori*, edited by P. A. Boghossian, and C. Peacocke, 117–149. Oxford: Oxford University Press.
Field, Hartry. 2018. "Epistemology from an Evaluativist Perspective." *Philosophers' Imprint* 18.
Fritz, James. 2017. "Pragmatic Encroachment and Moral Encroachment." *Pacific Philosophical Quarterly* 98 (S1): 643–661.
Fritz, James. 2021. "Akrasia and Epistemic Impurism." *Journal of the American Philosophical Association* 7 (1): 98–116.
Fritz, James. Forthcoming. "Knowledge and the Many Norms on Action." *Erkenntnis*.
Fumerton, Richard. 2006. *Epistemology*. Oxford: Wiley-Blackwell.
Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: 121–123.
Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge: Harvard University Press.

Gibbard, Allan. 2012. *Meaning and Normativity*. Oxford: Oxford University Press.
Goldman, Alvin I. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.
Greco, Daniel. 2014. "A Puzzle About Epistemic Akrasia." *Philosophical Studies* 167 (2): 201–219.
Hannon, Michael. 2019. *What's the Point of Knowledge? A Function-First Epistemology*. Oxford: Oxford University Press.
Haslanger, Sally. 2000a. "Defining Knowledge: Feminist Values and Normative Epistemology." In *The Proceedings of the Twentieth World Congress of Philosophy, Volume 8: Contemporary Philosophy*, edited by Daniel Dahlstrom. Charlottesville: Philosophy Doc Ctr.
Haslanger, Sally. 2000b. "Gender and Race: (What) Are They? (What) Do We Want Them to Be?" *Noûs* 34 (1): 31–55.
Hawthorne, John, and Jason Stanley. 2008. "Knowledge and Action." *Journal of Philosophy* 105 (10): 571–590.
Horowitz, Sophie. 2014. "Epistemic Akrasia." *Noûs* 48 (4): 718–744.
Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon.
Kitcher, Philip. 1992. "The Naturalists Return." *The Philosophical Review* 101 (1): 53–114.
Kornblith, Hilary. 2002. *Knowledge and its Place in Nature*. Oxford: Oxford University Press.
Ludlow, Peter. 2014. *Living Words: Meaning Undetermination and the Dynamic Lexicon*. Oxford: Oxford University Press.
McHugh, Conor, Jonathan Way, and Daniel Whiting. 2019. *Metaepistemology*. Oxford: Oxford University Press.
McPherson, Tristram. 2012. "Unifying Moral Methodology." *Pacific Philosophical Quarterly* 93 (4): 523–549.
McPherson, Tristram. 2020. "Ardent Realism Without Referential Normativity." *Inquiry* 63 (5): 489–508.
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. "Metaethics and the Conceptual Ethics of Normativity." *Inquiry*.
McPherson, Tristram, and David Plunkett. Forthcoming. "Evaluation Turned on Itself: The Vindicatory Circularity Challenge to the Conceptual Ethics of Normativity." In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau, Vol 16., 207–232. Oxford: Oxford University Press.
McPherson, Tristram, and David Plunkett. Manuscript. "On the Nature and Foundations of Epistemic Normativity."
Nietzsche, Friedrich Wilhelm. 1887/1994. *On the Genealogy of Morality*. New York: Cambridge University Press.

Olson, Jonas. 2018. "Moral and Epistemic Error Theory: The Parity Premise Reconsidered." In *Metaepistemology*, edited by C. McHugh, J. Way, and D. Whiting, 107–121. Oxford: Oxford University Press.

Pérez Carballo, Alejandro. 2020. "Conceptual Evaluation: Epistemic." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 304–332. Oxford: Oxford University Press.

Plunkett, David. 2015. "Which Concepts Should We Use?: Metalinguistic Negotiations and The Methodology of Philosophy." *Inquiry* 58 (7-8): 828–874.

Plunkett, David. 2016. "Conceptual History, Conceptual Ethics, and the Aims of Inquiry: A Framework for Thinking About the Relevance of the History/Genealogy of Concepts to Normative Inquiry." *Ergo* 3 (2): 27–64.

Plunkett, David. 2020. "Normative Roles, Conceptual Variance, and Ardent Realism About Normativity." *Inquiry* 63 (5): 509–534.

Plunkett, David, and Scott Shapiro. 2017. "Law, Morality, and Everything Else: General Jurisprudence as a Branch of Metanormative Inquiry." *Ethics* 128 (1): 37–68.

Plunkett, David, and Tim Sundell. 2013a. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosophers' Imprint* 13 (23): 1–37.

Plunkett, David, and Timothy Sundell. 2013b. "Dworkin's Interpretivism and the Pragmatics of Legal Disputes." *Legal Theory* 19 (3): 242–281.

Plunkett, David, and Tim Sundell. 2021. "Metalinguistic Negotiation and Speaker Error." *Inquiry* 64 (1–2): 142–167.

Pritchard, Duncan. 2016. *Knowledge*. New York: Palgrave.

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

Scharp, Kevin. 2020. "Defective Concepts." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett. Oxford: Oxford University Press.

Scharp, Kevin. Manuscript. "Replacing Knowledge."

Schroeter, Laura, and Francois Schroeter. 2014. "Normative Concepts: A Connectedness Model." *Philosophers' Imprint* 14 (25): 1–26.

Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.

Streumer, Bart. 2017. *Unbelievable Errors: An Error Theory About All Normative Judgments*. Oxford: Oxford University Press.

Thomasson, Amie L. 2016. "Metaphysical Disputes and Metalinguistic Negotiation." *Analytic Philosophy* 57 (4): 1–28.

Williams, J. Robert G. 2018. "Normative Reference Magnets." *The Philosophical Review* 127 (1): 41–71.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, Timothy. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.