

Virtues, Rights, or Consequences? Mapping the Way for Conceptual Ethics

MATTHIEU QUELOZ

Are there virtues that constitutively involve using certain concepts? Does it make sense to speak of rights or duties to use certain concepts? And do consequentialist approaches to concepts necessarily have to reproduce the difficulties that plague utilitarianism? These are fundamental orientating questions for the emerging field of conceptual ethics, which invites us to reflect critically about which concepts to use. In this article, I map out the ways in which conceptual ethics might take its cue from virtue-ethical, deontological, and consequentialist traditions of ethical thought while also flagging the main challenges facing each approach. I end by sketching how the various dimensions of evaluation singled out by these three traditions might be combined in a single approach.

Introduction

What makes a good concept? Unlike judgements or propositions, concepts are neither true nor false, neither warranted nor unwarranted. And yet some concepts evidently seem better than others. As the building-blocks of our thoughts, the concepts we possess delimit the horizon of our possibilities: what thoughts we can entertain, what desires we can form, and even what we can intelligibly do: without an entire conceptual edifice in place, squiggling something on a piece of paper could not count as the signing of a tax certificate; and a different set of concepts is required for it to count as the signing of a death warrant.

What is more, the concepts we employ inform our deliberation within that horizon of possibilities. They focus our attention, channel our emotions, direct our imagination, and marshal our memories; perhaps most notably, they determine what we recognize as a reason for or against what, weaving the web of reasoning pathways we consider correct, if

not necessarily the one that is objectively correct. In light of how profoundly our conceptual repertoire can shape our affairs, it seems worth thinking carefully about which concepts to live by.

The fast-growing literature on “conceptual ethics,” the evaluative reflection on which concepts to use, suggests that many are reaching the same conclusion.¹ But how exactly are we to evaluate concepts or conceptions? As I like to think of it, the newly invigorated enterprise of conceptual ethics invites us to seek *reasons for concept use*: reasons for a community to structure their affairs in terms of certain concepts rather than others. Where might we find such reasons? And how do they relate to that more immediate currency of ethics, reasons for action?

In what follows, I propose to map out the main ways in which conceptual ethics might hope to take its cue from the virtue-ethical, deontological, and consequentialist traditions of ethical thought. This mapping exercise seems called for by the fact that while the literature on conceptual ethics and conceptual engineering has devoted much ink to questions such as what concepts are, whether one might change concepts, and how one might do so without changing the topic, there have scarcely been any attempts to really put the *ethics* into conceptual ethics, and spell out on what sorts of ethical grounds one might evaluate concepts to begin with: are there virtues that constitutively involve using certain concepts? Does it make sense to speak of rights or duties to use certain concepts? And do consequentialist approaches to concepts necessarily have to reproduce the difficulties that plague utilitarianism?

¹ For overviews of this emerging field, see Alexis Burgess and David Plunkett: Conceptual Ethics I, in: *Philosophy Compass* 8/12 (2013) 1091–1101; Conceptual Ethics II, in: *Philosophy Compass* 8/12 (2013) 1102–1110, who coined the phrase “conceptual ethics,” as well as Herman Cappelen and David Plunkett: Introduction: A Guided Tour of Conceptual Engineering and Conceptual Ethics, in: *Conceptual Engineering and Conceptual Ethics*, ed. by Alexis Burgess, Herman Cappelen, and David Plunkett (Oxford: Oxford University Press, 2020) 1–26; and Steffen Koch, Guido Löhr, and Mark Pinder: Recent Work in the Theory of Conceptual Engineering, in: *Analysis* 83/3 (2023) 589–603.

In the limited space available, I can only gesture towards what seem to me like fruitful ways of tackling these questions. But my aim is not so much to offer a conclusive treatment as to open up avenues for further inquiry and flag some of the most interesting opportunities, complexities, and challenges facing attempts to bring the virtue-ethical, deontological, and consequentialist traditions of ethical thought to bear on conceptual ethics. To this end, I consider each of these three approaches in turn and highlight the main lines of inquiry as well as the chief difficulties with which they present conceptual ethicists. I end by sketching how the various dimensions of evaluation brought to light by these three traditions might be combined in a single approach.

1. Virtue-Ethical Approaches to Conceptual Ethics

We can begin with ethical thought in a virtue-ethical style: what might a virtue-ethical approach to conceptual ethics look like? This is a comparatively unexplored question in the literature on conceptual ethics. But it is safe to assume that a virtue-ethical approach to conceptual ethics would begin by inviting us to consider what concepts the *virtuous person* would use. After cataloguing the virtues possessed by the virtuous person, one could then examine what concepts each virtue requires for its realization.

The difficulty, however, is that there is no one-to-one correspondence between virtue possession and concept possession. Some virtues do not require possession of the *concept* of the virtue in question. They are compatible with what we might call *conceptual innocence*, where one possesses a virtue while remaining innocent of the concept thereof. It is possible to possess the virtue of temperance without possessing the concept of temperance; or to be benevolent without ever thinking in terms of benevolence.

Even if one possesses the concept of a virtue, one may not have to apply it to oneself when acting *from* that virtue. The person who acts from benevolence is typically not moved by the consideration that *this is the benevolent thing to do*, but by considerations that make no mention of benevolence, such as “It will comfort him,” or “She has more use

for it.” This is to display what Bernard Williams calls *intelligent innocence*.² One possesses the concept of a virtue, but does not apply it to oneself.

A particular problem for a virtue-ethical approach to conceptual ethics is that while some virtues merely *allow* for intelligent innocence, other virtues *demand* it.³ If the Spartan king Leonidas had chosen to fight the Battle of Thermopylae primarily *because it was the courageous thing to do*, he would have displayed not so much courage as an oddly reflexive preoccupation with his image of himself as a courageous person. At most, the virtue of courage allows self-conscious awareness of one’s own courage to act as a secondary consideration fortifying one’s resolve. But if it is to be *courage* one exhibits, one’s principal motivation had better not be: “This is the courageous thing to do.” And with virtues such as modesty, the demand for intelligent innocence is uncompromising: awareness of one’s own modesty under that description destroys the virtue. To say something *because it is modest* is precisely not to be modest, but to exhibit what we call “false modesty.”

The relations of virtues to concepts are therefore complex. Far from always *requiring* the use of specific concepts, virtues may merely *allow* for their use, or even *preclude* it. Even when a virtue requires certain concepts under which to conduct the deliberation that the virtue expresses itself in, there is typically not just *one* concept involved. Yes, the truthful person minimally requires the concept of truth; but so does the liar; and the truthful person also has to be conceptually sensitive to a wider range of considerations articulated in terms of such concepts as honesty, accuracy, sincerity, transparency, or deception.

Due to these complexities in the relations that virtues bear to concepts, a systematic

² Bernard Williams: *Utilitarianism and Moral Self-Indulgence*, in: *Moral Luck* (Cambridge: Cambridge University Press, 1981) 40–53 (46).

³ In this paragraph and the next, I develop some suggestive remarks that Williams makes about virtues in Bernard Williams: *Ethics and the Limits of Philosophy* (London: Routledge, 2011) 11; *Utilitarianism and Moral Self-Indulgence*, in: *Moral Luck* (Cambridge: Cambridge University Press, 1981) 40–53 (46).

virtue-ethical approach to conceptual ethics will have to operate on a case-by-case basis, without much by way of a readily applicable, off-the-shelf virtue-theoretical framework to guide one, because the range of considerations that the virtuous person should be sensitive to is too various, and the conceptual exigencies of one virtue do not necessarily carry over to the conceptual exigencies of other virtues.

Nevertheless, a virtue-ethical approach does hold promise for conceptual ethics. Insofar as one has reason to cultivate a certain virtue and that virtue requires one to possess certain concepts, reasons of virtue will provide reasons for concept use: reasons to cultivate the concepts themselves as integral constituents of virtue.

2. Deontological Approaches to Conceptual Ethics

What about deontological approaches? Can we identify reasons for concept use that take the form of duties to use concepts?

It quickly becomes apparent that there are bound to be at least *some* duties to use certain concepts. Kantian theory, so far as it demands that one be motivated by a sense that it is one's moral duty to do something, would seem to imply a duty to use the concept of *moral duty*, for example, though its linguistic expression might also operate through equivalent terms, such as "moral obligation," or a special sense of "ought." This presents us with a salient example of a deontological reason for concept use—the moral duty to be conceptually sensitive to moral duty.

But must one's sense of duty always register explicitly under the concept of moral duty? Or can the Kantian demand that one be motivated by a sense of duty be relaxed somewhat, so that it admits motives that register one's duties under more particularized descriptions? The moral agent might then be motivated not necessarily by the all-purpose thought "Because this is my moral duty," but by any of a variety of more particularized thoughts, such as "Because I promised," or "Because he needs it."

Indeed, it would seem that one *needs* to be conceptually equipped to register one's duties under more particularized descriptions. Otherwise, one would find it impossible to discriminate between situations in which the thought "This is my moral duty" applied and situations in which it did not. It is only thanks to one's possession of concepts such as *promise* or *need* that one is in a position to recognize that a moral duty results from the applicability of these concepts. Likewise, one will require whatever concepts articulate *what it is* that it is one's duty to do if one is to be in a position to recognize *that* it is one's duty to do it. The overarching duty to use the concept *moral duty* thus ramifies into a wider array of duties to use whatever concepts equip one to recognize *when* one has a duty and *what* one's duty is.

Exploiting this type of ramification, one might offer deontological reasons for concept use according to the following schema: it is one's duty to do X; using concept F is necessary to the fulfilment of that duty, because it is necessary either to recognizing when the duty arises or to recognizing what it requires; therefore, it is one's duty to use concept F. This way, our various duties might spawn a corresponding variety of rationales for concepts.⁴

A deontological approach to conceptual ethics might then be developed further by also considering *rights* to use concepts. This is the path advocated by Nicholas Smyth.⁵ He argues that people have rights to concepts such as *gay* or *black* if these concepts play an indispensable role in making people who they are. "This means much more than just that they have the right to *think of themselves* as falling under the concept," Smyth explains. "They have the right to live in a surrounding community which generally respects the classificatory boundaries under which they have formed their identities and which counts

⁴ I am grateful to Claus Beisbart for pressing me on this point.

⁵ Nicholas Smyth: The New Philosopher-Kings: Conceptual Engineering and Social Authority, in: PhilArchive (2023) 1–34.

them as falling under those boundaries.”⁶ Accordingly, he calls these concepts *identity-constituting* concepts.

Such a rights-based development of a deontological approach to conceptual ethics seems particularly well suited for defending concepts against proposals for conceptual abandonment. Sally Haslanger⁷ or Herman Cappelen,⁸ for example, advocate the abandonment of concepts such as *woman* or *democracy*.⁹ Against this, a rights-based approach could be used to argue that these proposals threaten to violate the rights of individuals to concepts that are essential to the way they think of themselves and constitute their identity. Deontological reasons for concept use can thus stand in the way of proposals to eliminate certain concepts from our repertoire.

As Smyth observes, such proposals exacerbate the already salient question of the *social authority* of conceptual engineers: *by what right* do some individuals purport to legislate what concepts others should or should not use? Though this question may have acquired a new urgency in view of the rise of conceptual engineering, the problem itself is not new. Nietzsche once remarked that “a certain legislation of *concepts*” was in fact what “Plato and fundamentally all the post-Socratics did.”¹⁰

It is also possible to envisage such a rights-based approach being extended beyond the retention of identity-constituting concepts, moreover. There may be concepts that citizens have a right not just to retain, but to *acquire*, because their ability to exercise their political rights depends on it: for example, popular consent can only function as an effective check

⁶ Op. cit., 14.

⁷ Sally Haslanger: Gender and Race: (What) Are They? (What) Do We Want Them to Be?, in: *Noûs* 34/1 (2000) 31–55.

⁸ Herman Cappelen: *The Concept of Democracy: An Essay on Conceptual Amelioration and Abandonment* (Oxford: Oxford University Press, 2023).

⁹ For a more recent reassessment of her earlier proposals, see Sally Haslanger: *Going On, Not in the Same Way*, in: *Conceptual Engineering and Conceptual Ethics*, ed. by Alexis Burgess, Herman Cappelen, and David Plunkett (Oxford: Oxford University Press, 2020) 230–260.

¹⁰ Friedrich Nietzsche: *Digital Critical Edition of the Complete Works and Letters* (Berlin: De Gruyter, 2009) 1885:34[84].

on state power if the populace is conceptually equipped to recognize abuses of power as such. The idea of *rights to concept acquisition* would then carry implications for education, public discourse, and censorship. Similarly, ideals of freedom of thought or freedom of religion could conceivably be shown to imply rights to certain concepts.

But even if the deontological approach is creatively extended in these ways, it is doubtful that it could become applicable to all kinds of concepts. Does it make sense to say that we have a *duty* or a *right* to use the concept *meeting point*, or the concept *chair*, or the concept *uncanny*? It seems that many of our concepts resist rationalization in deontological terms.¹¹ To say anything about what makes these concepts better or worse, we have to look beyond deontological approaches.

3. Consequentialist Approaches to Conceptual Ethics

This brings us to the last of the three main traditions of ethical thought, the consequentialist approach. It invites us to seek reasons for concept use in the consequences of using a concept: what difference does the use of a concept make to our lives? Do things go better or worse once this concept is introduced into our conceptual repertoire? These questions can be asked of any concept or conception—they are as applicable to rival conceptions of genes as they are to different value concepts; indeed, it was arguably a version of this approach that Hume and Nietzsche applied to their critiques of certain Christian values.¹²

¹¹ Moreover, there are good reasons not to couch every consideration bearing on conceptual ethics in terms of rights; see Matthieu Queloz: *The Dworkin–Williams Debate: Liberty, Conceptual Integrity, and Tragic Conflict in Politics*, in: *Philosophy and Phenomenological Research* Early View (2023) 1–27; *A Shelter from Luck: The Morality System Reconstructed*, in: *Morality and Agency: Themes from Bernard Williams*, ed. by András Szigeti and Matthew Talbert (New York: Oxford University Press, 2022) 182–209.

¹² On Hume, see Rachel Cohon: *Hume’s Moral Philosophy*, in: *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta (2010); *Hume’s Morality: Feeling and Fabrication* (New York: Oxford University Press, 2008). On Nietzsche, see John Richardson: *Nietzsche’s Values* (New York: Oxford University Press, 2020), and Matthieu Queloz: *Nietzsche’s Conceptual Ethics*, in: *Inquiry* 66/7 (2023) 1335–1364.

The best-known elaboration of consequentialism, which has, unhelpfully, become almost synonymous with it, is utilitarianism. Utilitarianism evaluates consequences based on how they affect a single metric, utility, which measures welfare or preference-satisfaction in one way or another, depending on what version of utilitarianism one uses. For the purposes of evaluating concepts, the most suitable form of utilitarianism is not *act* utilitarianism, which directly measures the value of particular actions according to the amount by which they increase utility, but rather *rule* utilitarianism, which judges particular actions according to whether they conform to rules that, if consistently followed, would tend to maximize utility. Instead of directly evaluating particular actions by their consequences, rule utilitarianism thus focuses on the consequences of adopting certain rules under which actions are selected.

That the latter form of utilitarianism is better suited to conceptual ethics comes out if we use the alternative nomenclature of *direct* and *indirect* utilitarianism. This has the advantage of acknowledging, first, that it is not only *rules* that can be evaluated in terms of their consequences, but also other indirect determinants of action, including notably the *concepts* under which an agent deliberates; and second, the direct/indirect nomenclature makes explicit that there is an element of indirection involved in indirect utilitarianism, because instead of directly evaluating actions in terms of their consequences for utility, it evaluates the rules, concepts, or dispositions that inform the agent's choices. This element of indirection makes room for the possibility that the concepts whose use tends to produce the best consequences overall may not advert to consequences at all, and may indeed recommend actions that agents thinking in purely utilitarian terms would never have selected. In this way, indirect utilitarianism can yield consequentialist reasons to embrace and live by virtue-ethical or deontological concepts, for example.

Direct utilitarianism, by contrast, is less well positioned to provide reasons for concept use. It is firmly focused on providing reasons for action, which it locates in the fact that a

particular action is optimific: it produces the best consequences in terms of utility. In special cases, the optimific action might itself happen to consist *in using a certain concept*, so that the resulting reasons for action in effect amount to *reasons to use a concept*.

But it is important to realize that these reasons to use a concept are *not* the same as the *reasons for concept use* that we presented conceptual ethics as seeking. While reasons to use a concept remain, in line with direct utilitarianism's focus, reasons for a particular action, reasons for concept use are not reasons for action, but rather reasons to cultivate the *dispositions* to use a concept and treat certain types of considerations *as* reasons for action.

Unlike direct utilitarianism, indirect utilitarianism can systematically be focused on reasons for concept use. By stepping back from particular actions and considering the wider pattern of consequences that the general use of a concept tends to have across a variety of situations, the indirect approach can validate the dispositions to be responsive, in one's day-to-day deliberation, to the reasons articulable under that concept, even if these reasons do not themselves refer to consequences. This results in a two-level structure, with openly consequentialist reasons for concept use at the more reflective level, and various kinds of reasons for action, which may or may not be consequentialist, at the first-order level.

However, even though indirect utilitarianism accommodates non-consequentialist considerations, the ultimate source of rational authority remains the single metric of utility. Conceptual sensitivity to other normative considerations is admitted only insofar as it still serves the overall maximization of utility in a roundabout way—and these other considerations can never truly override that overarching consideration of utility. As comfortable as indirect utilitarians may be with recommending that we relax into heeding virtue-ethical or deontological reasons when time is short and decisions need to be made, they still regard these reasons, at a more reflective level, as rule-of-thumb-like, *presumptive* reasons we let ourselves be guided by *faute de mieux* in the heat of action. The *real* reasons

are still those arising directly from what truly matters, namely utility. We may only be able to work back to these real reasons in the cool hour of reflection, when we have the opportunity to tweak what presumptive reasons we rely on the rest of the time. But it is ultimately these consequentialist considerations that constitute the reasons to be conceptually sensitive to certain presumptive reasons rather than others. Thus, both kinds of reasons derive whatever rational authority they possess from the practical aim of maximizing utility.

The resulting consequentialist approach to conceptual ethics holds considerable appeal, not least because it can straightforwardly be applied to just about any concept. But many of the familiar difficulties that plague utilitarianism also threaten to carry over to this kind of approach to conceptual ethics. In particular, I see three main difficulties with such a consequentialist approach to conceptual ethics.

Perhaps the most obvious difficulty is that in evaluating concepts *solely* with regard to their utility, we commit ourselves to a stingy axiology. Are we really willing to subscribe to the view that the only thing that truly matters is whatever utility quantifies, and that the value or importance of absolutely everything else can be adequately captured in terms of its consequences for utility? The controversial assumption is not so much the positive claim that what is quantified through the notion of utility matters, but the implicit negative claim that *nothing else* matters.

A second difficulty is the instability involved in trying to live by some considerations while believing, at a more reflective level, that they do not really matter. On the consequentialist approach, non-consequentialist ways of thinking are reflectively understood as mere tools for the maximization of utility rather than as articulating considerations that seriously compete with consequentialist considerations. Adverting to non-consequentialist considerations is simply part of the rational mechanism by which the concepts achieve their desired effect, while the *real* rationale or justification for thinking this way is that it maximizes utility. At the reflective level, utilitarianism thus

invites us to withdraw our evaluative commitment from the panoply of non-consequentialist thick concepts that guide our everyday deliberation, and refocus that evaluative commitment entirely on the thin master concept of utility. All that really matters is utility; but in order to maximize it, we occasionally have to regard other considerations *as if* they mattered.

This is at best an uneasy combination. If the only real reason to heed a non-consequentialist consideration ultimately lies in the consequentialist rationale for doing so, then one's commitment to taking that consideration seriously is bound to unravel under reflective scrutiny.¹³ Distinguishing between different levels of thought or different time-points in the agent's stream of consciousness cannot indefinitely stave off the reflective instability. Unless one opts for some form of "Government House utilitarianism,"¹⁴ whereby everyone except a small elite is systematically kept in the dark about the consequentialist rationales for non-consequentialist reasoning, both types of consideration have to cohabit within a single mind, at the end of the day, and they cannot realistically be so fully insulated from each other as to prevent the rational tension between them from alienating the agent from his or her non-consequentialist concerns.

The third and least well understood difficulty is that it is not at all clear that the reasons that can be advanced *for* the use of a concept should always count for more than the reasons adverted to *in* its use. Why should the reasons for concept use derivable from considerations of utility be granted authority over deeply entrenched reasons that have at least as much force with us? By what right do these reasons for concept use invalidate considerations that strike many as more immediately compelling? Surely, the reasons that guide and flow from our everyday judgements sometimes simply count for more than any

¹³ See Bernard Williams: The Structure of Hare's Theory, in: *Philosophy as a Humanistic Discipline*, ed. by A. W. Moore (Princeton: Princeton University Press, 2006) 76–85 (80), and Matthieu Queloz: The Self-Effacing Functionality of Blame, in: *Philosophical Studies* 178/4 (2021) 1361–79 (1371–1376).

¹⁴ B. Williams: *Ethics and the Limits of Philosophy*, 120.

reasons that could be offered for or against cultivating the dispositions to judge in these terms.¹⁵

If we learn that there is a human being trapped in a burning building, for example, then, on the strength of that thought alone, we mobilize enormous resources to save them.¹⁶ Now imagine that the utilitarian tells us: “The fact that this is a human being provides no real reasons for action, because the consequences of using the concept *human being* are not sufficiently conducive to the maximization of utility to justify our thinking in these terms. There may be better concepts in the vicinity. But we are better off ignoring this particular concept and its concomitant reasons for action.” Is the concept’s poor score on the utility metric really more compelling than the thought that there is a human being trapped in the burning building? That presupposes a degree of authority on the part of utilitarian reasons for concept use which it is not clear they possess.

Though powerful and well suited to the evaluation of concepts in its indirect utilitarian manifestation, the consequentialist approach to conceptual ethics is thus not without its difficulties. Those not willing to rely exclusively on this form of evaluation will therefore want to look further. Is there perhaps a way of obtaining the universal applicability of the consequentialist approach without surrendering the sensitivity to other dimensions of evaluation embodied by the virtue-ethical and deontological approaches?

¹⁵ Bernard Williams: *op. cit.*, 127, makes a related point in passing, in connection with proposals to adopt a more sharply delimited concept of person.

¹⁶ I take the example from Williams’s reflections on the concept *human being* as a weighty ethical concept in its own right. See B. Williams: *The Human Prejudice*, in: *Philosophy as a Humanistic Discipline*, ed. by Moore (Princeton: Princeton University Press, 2006) 135–154 (142). See also Cora Diamond: *Bernard Williams on the Human Prejudice*, in: *Philosophical Investigations* 41/4 (2018) 379–398.

4. Combining the Three Approaches

In previous work,¹⁷ I have advocated a “concern-satisfaction” approach to conceptual ethics that invites us to consider whether the effects of using a concept tend to be conducive to the satisfaction of our concerns—the needs, desires, projects, and aspirations that animate our use of anything like that concept to begin with.¹⁸

This sounds, certainly at first pass, like a variation on a consequentialist approach, which raises the question of whether it can avoid the difficulties plaguing the indirect utilitarian approach we considered. Let me therefore end by sketching how the concern-satisfaction approach avoids the three problems I identified for the indirect utilitarian approach. This will reveal that the concern-satisfaction approach is only superficially consequentialist, and in fact promises to offer a way of *synthesizing* the various dimensions of evaluation singled out by the three approaches we considered.

A key feature of the concern-satisfaction approach is that while it does *begin* by considering the consequences of using a concept, it does not *evaluate* these consequences merely on the basis of a single metric such as utility or some other measure of welfare, but rather draws on the entire stock of evaluative considerations adverted to by the rest of our concepts. This means that while it is methodologically consequentialist, in that it pays particular attention to the consequences of using certain concepts, it is firmly non-consequentialist in its axiology: the concern-satisfaction approach might evaluate a concept according to its tendency to help us *honour certain rights*, or to *realize certain virtues*. As one might also put it, the concern-satisfaction approach is consequentialist in form, but not in substance.

¹⁷ Matthieu Queloz: Function-Based Conceptual Engineering and the Authority Problem, in: *Mind* 131/524 (2022) 1247–1278.

¹⁸ This account was itself a simplified sketch of a more complex, “needs-based” approach developed in Matthieu Queloz: *The Ethics of Conceptualization: A Needs-Based Approach* (Oxford: Oxford University Press, 2024).

The concern-satisfaction approach asks whether our use of a given concept helps or hinders the promotion of the concerns we have independently of the concept—but there is no limitation on what these concerns might be concerns *for*. They might be a concern for truthfulness—a virtue—or a concern for people’s right to freedom of thought.

This makes the standard of evaluation richly multidimensional. In place of the monism of utilitarianism, the concern-satisfaction approach puts a pluralism of values and projects: the dimensions along which we will evaluate concepts will be as irreducibly various as our concerns are. And instead of treating a commitment to the importance of utility as an incontrovertible bedrock from which the rest of our evaluative commitments must be reassessed, the concern-satisfaction approach treats any concepts and concerns *from* which we evaluate a given concept as subject to being evaluated in turn, based on the rest of our concepts and concerns. No concept or concern is invariably beyond question in the way that, on the utilitarian picture, the concept of *utility* and the concern with it are. On the concern-satisfaction picture, the only constraint is that we should draw on the concepts and concerns we are comparatively *more* confident in to evaluate the concepts we are *less* confident in.

This enables us to draw simultaneously on all the different dimensions of evaluation singled out by the three approaches we considered, and bring them jointly to bear on the evaluation of a concept. The concern-satisfaction approach thus promises to synthesize what is best and most relevant in each of the three traditions of ethical thought while retaining indirect utilitarianism’s wide applicability to all sorts of concepts.

At the same time, the concern-satisfaction approach avoids the three difficulties facing indirect utilitarianism. First, it abandons utilitarianism’s stingy axiology, since there is no commitment to recognizing only a single currency of value. Any kind of evaluative input is admissible in principle. We might criticize concepts for having consequences that render it unfair, or inimical to liberty, or deceptive, or dishonourable. Even the evaluative perspective provided by the concept under evaluation could, though barred from

informing that particular evaluation, nonetheless inform the evaluation of other concepts once it had been approved in light of the rest of our concepts.

Second, the concern-satisfaction approach avoids the problem of instability under reflection, because the considerations acting as reasons for concept use carry no implication that they matter to the exclusion of any other consideration: whether things go better or worse if we use a concept is a consideration we can heed alongside the considerations that more immediately present themselves in the use of that concept. It is merely that when reflecting on whether we should be using a concept, we need to temporarily disengage from the reasons the concept itself immediately provides in order to gain some critical distance from it. There is no doubt a sense in which the widespread presence of integrity in the world itself provides a reason to use the concept of integrity; but to leave it at that would be to evaluate the concept by an insufficiently independent standard. If the fact that there is a lot of *F* in the world, which we become sensitive to by using the concept *F*, were itself already a decisive reason to use the concept *F*, then any non-empty concept could legitimate its own use. The question must rather be whether we have any independent reasons to demarcate and attach a certain significance to what we subsume under *F*.

And third, we need not treat reasons for concept use as carrying more authority than reasons in concept use, on the concern-satisfaction approach. We can grant that reasons *for* concept use remain orthogonal to the business of justifying action in terms of reasons *in* concept use. And we can also grant that reasons for action often count for more than any reasons that might be advanced for the disposition to recognize them as reasons.

Indeed, I think it best to conceive of the reasons for concept use offered up by the concern-satisfaction approach not as reasons guiding normative deliberation at all, but rather, in the first instance, as reasons providing a form of explanatory understanding: in particular, *self-understanding*, of the kind that helps us to see why it *makes sense* for us to think in certain terms.

However, the kind of self-understanding worth having does not leave everything where it was. The idea is thus not that there is an impermeable barrier between the reasons that provide explanatory understanding and the reasons that guide normative deliberation. We should resist the overdrawn contrasts between explanation and justification that philosophers introduced to demarcate themselves from psychology.¹⁹ *Responsible* deliberation had better be informed by explanatory understanding, and worthwhile explanatory understanding can properly inform normative deliberation.

What acts as the conduit between explanatory understanding and deliberation is our *confidence* in concepts and the reason statements they enable us to articulate. Understanding how the use of a concept relates to our concerns puts us in a position to see whether and why it makes sense for us to cultivate the dispositions to treat certain types of considerations as justifying certain types of action.²⁰ This form of understanding should in turn affect our confidence in the concept and its concomitant reasons: our confidence should be *strengthened* by the realization that the concept ties in with concerns we critically endorse.²¹ It shows that the practice of justifying action in these terms is not just a fetish, or a product of deception, or an archaic holdover, but something that it *makes sense* for us to cultivate, because it helps us to live, in particular by serving concerns we have independently of that concept.

¹⁹ See Matthieu Queloz: Genealogy, Evaluation, and Engineering, in: *The Monist* 105/4 (2022) 435–451 (443); Martin Kusch: *Psychologism: A Case Study in the Sociology of Philosophical Knowledge* (London: Routledge, 1995).

²⁰ This distinction between justifying individual thoughts and explaining ways of thinking can be traced to Wittgenstein; see Matthieu Queloz: Wittgenstein on the Chain of Reasons, in: *Wittgenstein-Studien* 7/1 (2016) 105–130; Two Orders of Things: Wittgenstein on Reasons and Causes, in: *Philosophy* 92/3 (2017) 369–397.

²¹ An example of this in Williams's own work is his genealogical explanation of the intrinsic value of truth. See Matthieu Queloz: *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering* (Oxford: Oxford University Press, 2021) ch. 7; . for complementary discussions of how genealogical explanation can have evaluative import.

But equally, our confidence should be *weakened* by the realization that the concept does nothing for us, and that we accordingly have no reason to heed the reasons it offers up to deliberation. This may be because the concept has ceased to be conducive to the satisfaction of our concerns, having effectively been rendered obsolete by changing circumstances; or, more alarmingly, it may turn out that the concerns served by the concept are concerns of a sinister kind that one does not want to see served. In such a case, it would be one's failure to endorse the concerns providing reasons for concept use that would result in an erosion of confidence in the concept. This indicates an important complication that utilitarianism's monistic axiology has difficulty accommodating, namely that when the concepts that some people have reason to use are not the same as those that other people have reason to use, this does not necessarily reflect a difference in their respective circumstances. It may, more fundamentally, reflect a difference in their values. Some concepts align with, and thereby express, one's own values in a way that other concepts do not.

Conclusion

The tentative finding of this high-altitude survey of possible ways forward for conceptual ethics is, on the whole, encouraging: all three of the main traditions of ethical thought suggest fruitful and distinctive ways of approaching the evaluation of concepts. What is more, the rough sketch of a concern-satisfaction approach gives us reason to hope that these different traditions, though incompatible when interpreted as providing exhaustive ethical theories, might even be made to join forces when brought to a common focus on the evaluation of concepts.

Just because these ethical traditions differ so starkly in the aspects and nuances they each latch on to, they complement one another well, highlighting the variety of complexities that conceptual ethics must contend with. And conceptual ethics in turn provides a fresh lens through which to view these familiar ethical resources. This can

reinvigorate them, and perhaps even provide an opportunity to reassess whether they really are as mutually exclusive as they are often made out to be. When engaging in an exercise as abstractly demanding as that of bending the critical gaze afforded by our concepts back on these concepts themselves, we may well find that we need to tap into all the resources we can muster.

Bibliography

- 1 Alexis Burgess, and David Plunkett, 'Conceptual Ethics I', *Philosophy Compass*, 8 (2013), 1091–1101.
- 2 ———, 'Conceptual Ethics II', *Philosophy Compass*, 8 (2013), 1102–1110.
- 3 Herman Cappelen, *The Concept of Democracy: An Essay on Conceptual Amelioration and Abandonment* (Oxford: Oxford University Press, 2023).
- 4 Herman Cappelen, and David Plunkett, 'Introduction: A Guided Tour of Conceptual Engineering and Conceptual Ethics', in *Conceptual Engineering and Conceptual Ethics*, ed. by Alexis Burgess, Herman Cappelen and David Plunkett (Oxford: Oxford University Press, 2020), pp. 1–26.
- 5 Rachel Cohon, 'Hume's Moral Philosophy', in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta (2010).
- 6 ———, *Hume's Morality: Feeling and Fabrication* (New York: Oxford University Press, 2008).
- 7 Cora Diamond, 'Bernard Williams on the Human Prejudice', *Philosophical Investigations*, 41 (2018), 379–398.
- 8 Sally Haslanger, 'Gender and Race: (What) Are They? (What) Do We Want Them to Be?', *Noûs*, 34 (2000), 31–55.
- 9 ———, 'Going On, Not in the Same Way', in *Conceptual Engineering and Conceptual Ethics*, ed. by Alexis Burgess, Herman Cappelen and David Plunkett (Oxford: Oxford University Press, 2020), pp. 230–260.
- 10 Steffen Koch, Guido Löhr, and Mark Pinder, 'Recent Work in the Theory of Conceptual Engineering', *Analysis*, 83 (2023), 589–603.
- 11 Martin Kusch, *Psychologism: A Case Study in the Sociology of Philosophical Knowledge* (London: Routledge, 1995), pp. 439–443.
- 12 Friedrich Nietzsche, *Digital Critical Edition of the Complete Works and Letters*. ed. by Paolo D'Iorio (Berlin: De Gruyter, 2009).
- 13 Matthieu Queloz, 'The Dworkin–Williams Debate: Liberty, Conceptual Integrity, and Tragic Conflict in Politics', *Philosophy and Phenomenological Research*, Early View (2023), 1–27.

- 14 ———, *The Ethics of Conceptualization: A Needs-Based Approach* (Oxford: Oxford University Press, 2024).
- 15 ———, 'Function-Based Conceptual Engineering and the Authority Problem', *Mind*, 131 (2022), 1247–1278.
- 16 ———, 'Genealogy, Evaluation, and Engineering', *The Monist*, 105 (2022), 435–451.
- 17 ———, 'Nietzsche's Conceptual Ethics', *Inquiry*, 66 (2023), 1335–1364.
- 18 ———, *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering* (Oxford: Oxford University Press, 2021).
- 19 ———, 'The Self-Effacing Functionality of Blame', *Philosophical Studies*, 178 (2021), 1361–79.
- 20 ———, 'A Shelter from Luck: The Morality System Reconstructed', in *Morality and Agency: Themes from Bernard Williams*, ed. by András Szigeti and Matthew Talbert (New York: Oxford University Press, 2022), pp. 182–209.
- 21 ———, 'Two Orders of Things: Wittgenstein on Reasons and Causes', *Philosophy*, 92 (2017), 369–397.
- 22 ———, 'Wittgenstein on the Chain of Reasons', *Wittgenstein-Studien*, 7 (2016), 105–130.
- 23 John Richardson, *Nietzsche's Values* (New York: Oxford University Press, 2020).
- 24 Nicholas Smyth, 'The New Philosopher-Kings: Conceptual Engineering and Social Authority', *PhilArchive* (2023), 1–34.
- 25 Bernard Williams, *Ethics and the Limits of Philosophy* (London: Routledge, 2011).
- 26 ———, 'The Human Prejudice', in *Philosophy as a Humanistic Discipline*, ed. by Adrian W. Moore (Princeton: Princeton University Press, 2006), pp. 135–154.
- 27 ———, 'The Structure of Hare's Theory', in *Philosophy as a Humanistic Discipline*, ed. by A. W. Moore (Princeton: Princeton University Press, 2006), pp. 76–85.
- 28 ———, 'Utilitarianism and Moral Self-Indulgence', in *Moral Luck* (Cambridge: Cambridge University Press, 1981), pp. 40–53.