

Justin Clarke-Doane

Columbia University

Draft: 6.5.21

[This is a draft of a paper that is forthcoming in *Analysis*, with replies from Smith.]

From Non-Usability to Non-Factualism¹

Holly Smith has done more than anyone to explore and defend the importance of usability for moral theories. In *Making Morality Work*, she develops a moral theory that is almost universally usable. But not quite. In this article, I argue that no moral theory is universally usable, in the sense that is most immediately relevant to action, *even by agents who know all the normative facts.*² There is no moral theory knowledge of which suffices to settle deliberation about what to do. However, this is unsurprising if the question of what to do is not a question of fact. One upshot of the discussion is that the search for a universally usable moral theory is misconceived. Another is that, contra Smith (341), agents who are radically uncertain need not lack autonomy.³

Smith's Theory

What ought we do when we are uncertain (or ignorant) of what to do? It does not help to be told that we ought to do whatever the true theoretical account of good and bad, right and wrong -- whatever that is -- prescribes that we do. There is a palpable sense in which we *cannot* do that in

¹ Thanks to Holly Smith and Jordan Walters for comments.

² The agents can also be assumed to know all the non-normative facts and have the ability to perform all of the prescribed actions. But this will not be important.

³ All references to page numbers without a corresponding year of publication refer to Smith (2018).

such a situation. Even if we knew what the true theoretical account was -- act utilitarianism, say -- we might lack a belief as to what act (non-morally characterized) counted as good by its lights. In order to be *usable*, a theory must satisfy *Guidance Adequacy*, according to which “for every occasion for decision, there is an appropriate [rule] which can be directly used by the agent for making a decision on that occasion (234).” The way to make a theoretical account of good and bad, right and wrong, usable is to supplement it with what Smith calls decision guides (DGs).

Traditionally philosophers sympathetic to such ‘hybrid’ theories have assumed that only one DG was needed. For instance, if the true theoretical account of goodness says that maximizing utility is good, then perhaps the corresponding decision guide would say that we ought to maximize *expected* utility. But one of the important upshots of Smith’s book is that any theory satisfying *Guidance Adequacy* will have to encompass an array of DGs tailored to different agents with different states of uncertainty. There is no one-size-fits-all DG. The DGs ascend from DG_0 , which is just the true theoretical account, to the degenerate DG_δ , which basically says ‘do something you can’. At first pass, Smith says, when we are uncertain of what to do, we ought to do what is licensed by the lowest ranked DG -- i.e., DG_β , where $\beta = \min(x)$: DG_x is usable by us.

However, another insight of Smith’s discussion is that this too cannot be the end of the story. What if we are *morally (or normatively)* uncertain? In particular, what if we are uncertain about the *correct* ranking of the DGs?⁴ Then none of them may be usable by us either. Continuing down this road, we will need to introduce a new hierarchy of (what Smith calls) *precepts* (PREs)

⁴ Smith assumes for simplicity that this counts as moral uncertainty since nothing turns on it (53, n. 53).

that plays the role *vis a vis* the correct ranking of the DGs that the DGs themselves play *vis a vis* the true theoretical account of good and bad. When we are uncertain of what the correct ranking of the DGs is, we ought to do what is licensed by PRE_β , where $\beta = \min(x)$: PRE_x is usable by us.

Of course, if we can be uncertain about the correct ranking of the DGs, then we can be uncertain about the correct ranking of the PREs. So, Smith ultimately postulates an unending series of maxims by which to rank the PREs, councils to rank the maxims, and so on. What we ought to do when we are uncertain about what to do, Smith concludes, is to do what is licensed by Γ_β , where $\Gamma = \min(x)$: x is some DG, PRE, maxim, council etc., and $\beta = \min(x)$: Γ_β is usable by us.

Universal Usability

An apparent problem, to which Smith is sensitive, is to say why such a Γ_β should exist at all.

Could not an agent be uncertain ‘all the way up’ the Γ s? Smith expresses the worry thus:

[T]he truly comprehensive version of the Hybrid theory attempts to remedy agents’ uncertainties by adding new components to the comprehensive moral theory, components that an agent can use to decide what to do when uncertain about lower layers in the theory. However, these new layers simply open new elements about which agents can also be uncertain, thus creating new problems as the price for solving the original one. The approach of adding more layers to the moral theory appears doomed to fall short of our goal of rendering any moral theory usable by any agent, however deep her uncertainty (337).

Importantly, Smith concedes the point. She even allows that such systematically uncertain agents actually exist. Smith is forthright about the theoretical situation in Section 13.4.7.

I have difficulty seeing how [the theory] could be improved in ways that would render it invulnerable to the kinds of problems we have canvassed already (or their successors). If this pessimism is correct, then the...solution fails...to make moral codes usable, directly or indirectly, by every agent....[W]e have not found a way to preclude [such cases] (340).

Smith takes comfort in the sociological conjecture that such agents are unusual (340). Radical moral uncertainty only affects “[d]eeper thinkers...in the tiny minority” (343) - like philosophers (!), for example. But there is something unsettling about appeal to poll numbers in this setting. Is it really just dumb luck that we did not go down the rabbit hole of metaethics, only to surface unable to deliberate about what to do? Most of us can. But Smith allows that things could have easily been otherwise, even if we had lacked no ordinary agential or intellectual capacity. We could have been philosophers (gasp!), so reflective as to undercut our own autonomy (341).

What to Do and What we Ought to Do

The failure of Smith’s theory to be universally usable can seem like a serious problem. “The primary or even sole aim of a theory of right action”, Smith claims, “is to be usable by moral agents for making decisions” (354). However, I will argue that no theory is universally usable,

in the sense most immediately relevant to action, *even among those who know all the normative facts*.⁵ There is no theory knowledge of which suffices to settle deliberation about what to do.

Normative pluralism is the view, roughly, that any normative theory, syntactically individuated, that we might have adopted is true of a distinct subject matter (Clarke-Doane (2015), (2020, Ch. 6)).⁶ If this is hard to imagine, then note that it is a consequence of many prominent metaethical theories. Recall Horgan's and Timmons's 'moral twin earth' thought experiment. They write,

Suppose "that...human uses of 'good'...are regulated by certain functional properties; and that, as a matter of empirical fact, these are consequentialist...whose functional essence is captured by some specific consequentialist normative theory; call this theory Tc....Now consider Moral Twin Earth [where] Moral Twin Earthlings have a vocabulary that works much like human moral vocabulary....The properties tracked by twin English moral terms are...non-consequentialist moral properties, whose functional essence is captured by some specific deontological theory, call this...Td....[The problem is that m]oral and twin-moral terms do not [seem to] differ in meaning or reference, and hence...any apparent moral disagreements that might arise between Earthlings and Twin Earthlings [are] genuine...-- i.e., disagreements in moral belief and in normative moral theory, rather than disagreements in meaning" [1992, 460].

⁵ I focus on normative facts because those are the ones that are supposed to settle deliberation. The idea that non-normative facts do not settle it is just a variation on Hume's and Moore's points. But one is free to also assume that we know all the non-normative facts, and are able to do whatever the true normative theory prescribes, in what follows.

⁶ One must add additional constraints that are irrelevant here. See Clarke-Doane (2020, Chapter 6) for details.

Horgan and Timmons argue that moral naturalism is false because it has pluralist implications. It implies that two cultures whose use of moral terms is causally regulated by different properties do not disagree. More generally, their point is that normative pluralism predicts that people fail to have a normative disagreement when they really have one. So, normative pluralism is false.⁷

There is a different argument in the vicinity of Horgan's and Timmons's that does not depend on natural language semantics. If normative pluralism is true, then there is a gap between what we ought to do and what to do, in any sense of 'ought' (and similarly for what we have most reason to do, what is good to do, and so on). Even if we ought to kill the one to save the five, there is another property like the original that we could have latched onto -- call it $\text{ought}_{\text{Twin}}$ -- according to which we ought -- or rather $\text{ought}_{\text{Twin}}$ -- not kill the one.⁸ Recognizing this, the 'central deliberative question' remains -- *whether to do* what we ought, or $\text{ought}_{\text{Twin}}$, to do.⁹ It does not matter for this argument whether we would have latched onto $\text{ought}_{\text{Twin}}$ with 'ought' in some pertinent alternative scenario. All that matters is that, given that we ought to kill the one to save the five, and $\text{ought}_{\text{Twin}}$ not, the central deliberative question remains -- whether to do what we ought, or $\text{ought}_{\text{Twin}}$, to do.

⁷ It is useful to bring out the generality of their objection because some versions of non-naturalism have pluralist implications too. For instance, Scanlon's realist Carnapianism says that "as long as some way of talking [is] well defined, internally coherent, and [does] not have any presuppositions or implications that might conflict with those of other domains, such as science" such talk is true (2014, 27, original emphasis). See also Huemer (2005) and Cuneo and Shafer-Landau (2014).

⁸ For simplicity, I speak as though the semantic value of 'ought' were a property, even though it is more likely to be an operator. Nothing in what follows turns on what kind of thing one takes its semantic value to be.

⁹ See Lord (2017) for this terminology.

To be clear: there is no question that we really ought to kill the one (in whatever sense of ‘ought’ you like). (Nor is there any question of whether we really have most reason to kill the one, whether this would be good, whether it is the thing to do, and so on for any other normative terms.)¹⁰ We conceded that we ought. Nor is there any doubt that we ought to *use* ‘ought’ rather than ‘ought_{Twin}’. Any ought-like property will be self-sanctioning, in that we ought, in that sense, to use it (so we ought to use ‘ought’ and ought_{Twin} to use ‘ought_{Twin}’). Nor, finally, is there any question that we ought in some non-moral sense do what we ought_{Twin} rather than ought to do, since we are conditionalizing on pluralism about all normative properties, moral and otherwise. The question that remains open is what to do, not what we ought to do, in any sense of ‘ought’.

Nevertheless, as stated, this argument does depend on there being normative-like properties. It is still an argument against normative pluralism. It shows that *if* normative pluralism is true, *then* knowledge of what we ought to do (in any sense of ‘ought’) fails to do the primary thing it was supposed to do -- settle what to do.¹¹ This makes it sound like one could avoid the conclusion by simply rejecting normative pluralism. But this turns out to be incorrect, for two reasons.

First, the component of normative pluralism that one would need to reject is the *metaphysical* rather than metasemantic component. One would have to deny that there are normative-like

¹⁰ I will not continue to add this qualification.

¹¹ It is tempting to help one’s self to some special ‘ought’ (Vermaire (Forthcoming)). One can stipulate that the relevant (all-things-considered, or subjective, or whatever) ‘ought’ is the one ‘most immediately relevant to action’ (Jackson (1991, 472)). But if this just means that it is the ‘ought’ people actually tend to focus on, then knowing what we ought to do in that sense does not resolve our deliberation for familiar is/ought reasons. If it means the *right* kind of ought to settle deliberation, then we face the problem again at the level of ‘right’ (Risberg (Manuscript)).

properties, whether or not our language would have latched onto them in different circumstances. But it is hard to see how there could *fail* to be normative-like properties, in whatever sense there are normative properties. Who denies that there is a property of maximizing utility (Bentham (1789, ch. 1), Mill (1863)), causally regulating our use of the word ‘good’ (Boyd (1988)), or satisfying our ‘folk moral theory’ (Jackson (1998)) (even if people do deny that our use of ‘ought’ happens to latch onto them)? And while philosophers do deny that there are ‘non-natural’ properties like those postulated by Moore (1903), Huemer (2005), or Enoch (2012), even this is suspect (Clarke-Doane (2020, Sec. 1.5)). On typical formulations of platonism about properties, properties are as abundant as can be. They are the semantic values of predicates (barring pathological ones like Russell’s). So, there are properties ‘up there’ conforming to Moore’s or Huemer’s or Enoch’s conceptions right alongside Boyd’s and Jackson’s natural surrogates.¹² On the other hand, the standard formulation of nominalism about properties (Quine (1948)) says that the question of whether, e.g., Moore’s goodness ‘exists’ is just that of whether to accept some ideology into our theory of the world. One could reasonably wonder how the predicate ‘is good’ could be more objectionable than ‘is a bicycle’.

It is true that a very unusual kind of Aristotelian realism about properties, according to which properties are sparse (Armstrong 1980) but include normative ones, *could* be true. But it is hard to believe that normative properties’ practical clout turns on a *recherché* solution to the problem of universals! And, indeed, the second point is that even if there were *only* normative properties *per se* (and no merely normative-like ones), the question would remain whether to regulate our

¹² Moreover, since properties’ identity conditions entail instantiation conditions, there is no doubt about their being instantiated if they exist.

behavior by consulting them (rather than consulting non-normative properties).¹³ *But this question cannot be the one of whether we ought to, in any sense of ‘ought’, on pain of triviality.* We ought to regulate our behavior by consulting ought, for any self-respecting notion of ‘ought’. This banality does not help us. Even if our question is *easy*, it is not contentless.¹⁴

So, knowledge of the normative facts turns out not to settle deliberation about what to do, *whether or not normative pluralism is true.* Consequently, no normative theory is usable, in the sense most immediately relevant to action, even by those who know all the normative facts.

Non-Factualism and Uncertainty

The failure of any normative theory to be usable in this sense is a peculiar problem for normative theories. Again, Smith has demonstrated better than anyone that we do not determine what we ought to do for the sake of accumulating normative theorems. We do so in order to act. So, if knowledge of the normative facts fails to settle what to do, then this severely undercuts the import of those facts.¹⁵ By contrast, nobody would have suggested that, e.g., mathematical facts would settle—all by themselves—what to do, even in the arena of pure mathematics. They would not even settle—by themselves—what to prove, what logic to accept, or what axioms to use. That is an easy application of Hume’s point that one cannot derive an ‘ought’ from an ‘is’, as well as an illustration of Moore’s conclusion (Moore (1906)) that one can know that something is F, for any *descriptive* property, F, while competently wondering whether it is good.

¹³ That is, we can avoid ‘star philosophy’ as my colleague, Chris Peacocke, affectionately put it.

¹⁴ Indeed, it is apparently Nietzsche’s when he advocates “critiqu[ing]...the value of these values” (1887, Preface, § 6).

¹⁵ This is slightly tricky, since ‘import’ is a normative term. I will spare you a precise statement of the point.

What are we to make of this failure? We need to revisit the reasons for expecting some moral theory to be universally usable. The culprit is the assumption that when we deliberate, our conclusion is a belief. Although Smith claims that “Nothing in the present argument seems to turn on this issue (38, n. 20)”, *the failure of any theory to be universally usable is unsurprising if it is false*. It is the truism that belief is one thing and intention (or surrogate non-cognitive attitude) is another. “Even if...[a] belief were settled, there would still be issues of what importance to give it, what to do, and all the rest...For any fact, there is a question of what to do about it (Blackburn (1998, 70)).” From a ‘non-factualist’ perspective, the search for a universally usable moral (or even normative) theory is quixotic, indeed hopeless, from the start.

Of course, there is a mundane sense in which a non-factualist about what to do will allow that some -- actually, *any* -- moral theory is universally usable. We can all *decide* to use it! But this *practical* decision is not something that can be ‘packed into’ the theory itself. An agent-theory pair can satisfy Smith’s criteria of usability (16), and more -- so that the agent truly believes the theory’s verdicts and is able to do what it prescribes -- while she is unable to use the theory in the sense that matters. She may have no idea *whether to* regulate her behavior by consulting it, even though she knows that she *ought to* (and is able to perform the actions that it prescribes).¹⁶

¹⁶ What about a ‘theory’ that issued only imperatives? Such a ‘theory’ would tell us not that we *ought* to use it (or that this would be good, or that we have most reason to, etc.). It would say: ‘use this!’ Such a device -- indeed, infinitely many of them -- would be usable in roughly the present sense. But this is only because we cannot have the attitude corresponding to an imperative -- e.g., intention -- toward incompatible imperatives simultaneously. We cannot fully intend to kill the one (in a fixed context) and fully intend not to at the same time, since we can only do one or the other. However, we *can* fully believe that we ought to kill the one and fully believe that we ought not, at the same time. This is the well-known problem with relativism, construed as a factual doctrine (Hare (1997, 30)).

Although non-factualism tends to undercut the search for a universally usable theory, it also lets us avoid Smith's conclusion that "[d]eeper thinkers" (343) -- like philosophers! -- have lost their autonomy (341). If non-factualism is true, then we can settle what to do *despite having no idea what we ought to do* (just as, I argued, we can have no idea what to do while *knowing* what we ought to do). Let us return to an agent who is uncertain 'all the way up' the DGs, PREs, maxims, councils, and so on. Will she continue to act? Of course she will. Every moral philosopher knows a moral skeptic (not to mention a nihilist (Sec. 3.4.2)), and they behave remarkably like the enlightened. It is not just that their bodies move in ways that we would expect them to if they were deliberating and intending. They deliberate! Indeed, most of us who work in moral epistemology have first-personal acquaintance with deliberation despite uncertainty about every (non-logical) moral claim. Most of us have at least temporarily indulged in moral skepticism.

How do we do it? It is as if we ask ourselves what we ought to do, given that we do not know what to do. But that cannot be right if 'ought' is some kind of moral ought. By assumption, we do not know what we ought to do in any moral sense. So, perhaps we ask ourselves what we ought to do in some *non-moral* sense of 'ought'.¹⁷ Maybe we ask what we ought *rationally* to do, or *prudentially* out to do, or *all-things-considered* to do, for example. But that cannot be right either. For just as we can deliberate despite uncertainty about all moral claims, we can deliberate despite uncertainty about all normative claims of any kind. Indeed, skepticism about moral claims is commonly alleged to engender skepticism about other normative claims (Cuneo

¹⁷ This interpretation is similar to the one delivered by Sepielli's 'dividers' strategy. See Sepielli (Forthcoming).

2007). So, if we really became convinced of moral skepticism, and then became convinced of this ‘companions in guilt’ reasoning, then we would become convinced of normative skepticism.

Maybe our question is ineffable, albeit factual?¹⁸ Even if it were, that would not be the end of the story. For a parallel puzzle would arise *vis a vis* the ineffable fact in question. If the proposal is coherent, then we must be able to refer to the proposition we ponder, even if we cannot express it with a sentence. So, let us name the proposition, P. Then whether or not ‘P’ and ‘P is true’ mean the same thing, they are true in the same circumstances. So, if we are uncertain as to whether P is true, then we will be uncertain as to whether P. But we can certainly be uncertain as to whether an ineffable proposition is true! It is not even uncontroversial that the idea of an ineffable truth *makes sense*. Yet, if we were so uncertain, we would continue to deliberate as before.

Is there any other account of the question that we resolve when we are uncertain about all the normative (and even ineffable) truths? Of course. As before, the question could be the non-factual one of what to do (Clarke-Doane (2015, 2020), Risberg (Manuscript)).¹⁹ We may be uncertain about what we ought to do, what is good to do, what is the thing to do, and so on in any sense of these terms you like. (We may even be uncertain about whatever ineffable truths there are supposed to be.) We are still able to determine *what to do*. This is because the ‘central

¹⁸ Eklund considers this possibility in connection with (something like) the problem of normative pluralism, discussed previously, in his (2017). For further critical commentary, see Clarke-Doane (2020, Sec. 6.6).

¹⁹ In previous work, I have focused on deliberation under the assumption of normative pluralism (discussed above), rather than uncertainty. Risberg (manuscript) focuses on deliberation under uncertainty, and helpfully discusses the connection between this and deliberation under the assumption of pluralism.

deliberative question' is "not truth-apt, so that it is misleading to describe the agent's 'conclusion' as 'false' (38, n. 20)." This is so *even if there are normative (and ineffable) truths*.

Note that if this argument is sound, then normative truths are not 'deliberatively indispensable', contra Nagel (1997) and Enoch (2012). Moreover, if the argument from pluralism is sound, then they are also not 'deliberatively sufficient'. In other words, if what I have argued here is sound, then knowledge of the normative facts is *neither necessary nor sufficient* for settling what to do.

Objections and Replies

Although the above argument is in keeping with Smith's hope that "thinking about the problem of human cognitive deficiencies...is an illuminating way of testing and...developing one's views about...morality", she is apt to reject it. It implies that an agent may determine what to do despite systematic skepticism. More troubling for her project, it implies that an agent may be uncertain what to do *even though she knows exactly what she ought to do*. She may be uncertain whether to regulate her behavior by consulting the property she actually consults with 'ought'.

There are two ways that Smith might challenge this argument. The first is to claim that even if the conclusion were true, it is just a roundabout way of stating *motivation externalism*, a view to which Smith herself may subscribe. Motivation externalism says that an agent may sincerely judge that she ought to X (in whatever sense of 'ought') despite failing to be even defeasibly motivated to X (Brink 1986). Motivation is one thing, and deliberation is another. Smith is concerned with deliberation. She nowhere denies that we can be weak in will. But my point is

about deliberation too. My point is that knowing that we ought to do something does not settle whether to do it, regardless of whether we will go on to do it. Nor is it even necessary for settling this. What to do and what we ought to do, in any sense of 'ought', are independent.

To repeat, the question of what to do is not *factual*. If Gibbard (2003) is right, only intention settles it. Alternatively, maybe we can intend to do what we believe (to use the vulgar) we ought not to do. Whatever attitude we appeal to, our account is of deliberation, not motivation.

A different objection is that I have played fast and loose with intelligibility. The question of what to do *just is* the question of what we ought to do in some or another sense of 'ought' (moral, prudential, all-things-considered, or some such). So, if we are really sure that 'ought' out of our mouths picks out, e.g., consequentialist properties, then it is of no relevance that the same word, syntactically individuated, picks out deontological properties out of a possible community's mouth. The question of whether to kill the one to save the five is settled in the only sense that matters. Conversely, if we are truly uncertain about what we ought to do, in *every* sense, then we can only 'settle' what to do in bad faith -- by bracketing our doubts. But it is just intellectual doublethink (to recall Quine's phrase) to be skeptical of every normative claim while deciding to do something. To decide to do something is to deem it the *thing to do!*

However, there are three problems with this response. The first is that, even if it were right, that may simply be because non-factualists are correct about natural language semantics. If 'what ought we do?' out of our mouths is just another way of expressing practical indecision, then of

course we cannot be systematically ‘uncertain’ about what we ought to do while ‘knowing’ what to do. Likewise, we cannot really wonder what to do if we ‘know’ we ought to kill the one (even if we ought_{Twin} not). The ‘unintelligibility’ of the scenarios above is no argument for factualism.

But, second, it is hard to think of an argument that the scenarios really are unintelligible. The claim that there is no intelligible question of what to do when we are skeptical of what we ought to do attributes incoherence to agents unnecessarily. The only reason to do this is to preserve the assumption that the question of what to do is the question of what we ought to do, factually construed. Conversely, if there were really no intelligible question of what to do, given knowledge that we ought to kill the one (but ought_{Twin} not), then we might as well dismiss what are clearly legitimate epistemological worries too. For one can ‘transpose’ them to the key of normative pluralism. Most moral epistemologists would concede that evidence that we could have easily had very different, and so false, moral beliefs, threatens our moral deliberation. This is evidence that our moral beliefs are not *safe*. But now suppose, as, again, some realists maintain (Cuneo and Shafer-Landau (2014)), that we cannot change our moral beliefs very much without changing the subject. Then if we could have easily had systematically different moral beliefs, we really could have easily had moral_{Twin}, instead of moral, ones. This is no longer an *epistemological* problem. Had we had moral_{Twin} beliefs, our beliefs may not have been false, since they may have been true of the moral_{Twin} facts. But, insofar as we were rational_{Twin}, we still would have done what we ought (morally) not have. We would have been using the ‘wrong’ moral-like concepts to guide our behavior (Eklund 2017).²⁰ (This is *not* to say that we would

²⁰ I assume here that such realists must agree that moral_{Twin} properties are *like* moral ones, since they are just the transposition of different moral views to the level of different concepts.

have been using concepts that fail to be ‘metaphysically privileged’, contra Enoch & McPherson (2017). If metaphysical privilege has ramifications for *good* theorizing, then it is just another normative concept (Dasgupta (2018)), and the argument from pluralism reapplies to it. If it does not have such ramifications, then it is neither here nor there from the standpoint of which concepts we ought to use.) The *practical* problem of safety is the same whether it is epistemological or not.

The final problem is that even if the questions were unintelligible, we still need to answer them! Maybe we have taken too much philosophy, and really have talked ourselves into incoherence. Still: what to do? This question, even if unintelligible, is pressing, and, evidently, non-factual.

Conclusions

Smith’s book is a classic. Like most classics, it fails to accomplish what it sets out to. But also like most classics, the reasons that it fails are revealing. I have argued that no theory is universally usable, in the sense that is most immediately relevant to action, *even by agents who know all the normative facts*. No normative theory is such that knowledge of it suffices to settle deliberation. But that is simply because the question of what to do is not a question of fact. This means that the search for a universally usable theory is misconceived. But it also means that, contra Smith, agents who are normatively uncertain -- like us, perhaps -- need not lack autonomy.

Bibliography

Armstrong, D.M. [1980] *Universals and Scientific Realism*. Cambridge: Cambridge University Press.

Bentham, Jeremy. [1789] *An Introduction to the Principles of Morals and Legislation*. London: T. Payne & Son. Available online at <http://www.koeblergerhard.de/Fontes/BenthamJeremyMoralsandLegislation1789.pdf>.

Blackburn, Simon. [1998] *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Clarendon.

Boyd, Richard. [1988] "How to Be a Moral Realist," in Geoffrey Sayre-McCord (ed.), *Essays on Moral Realism*. Ithaca: Cornell University Press. 181–228.

Brink, David. [1986] "Externalist Moral Realism." *Southern Journal of Philosophy*. Vol. 24. (Supplement): 23–40.

Clarke-Doane, Justin. (2015) "Objectivity in Ethics and Mathematics," Ben Colburn (ed.), *Proceedings of the Aristotelian Society, The Virtual Issue, No. 3 (Methods in Ethics)*. Available online: <https://www.aristoteliansociety.org.uk/the-virtual-issue/the-virtual-issue-no-3/>

----- (2020) *Morality and Mathematics*. Oxford: Oxford University Press.

Cuneo, Terrence. (2007) *The Normative Web*. Oxford: Oxford University Press.

Cuneo, Terrence, and Russ Shafer-Landau. [2014] "The Moral Fixed Points: New Directions for Moral Nonnaturalism." *Philosophical Studies*. Vol. 171. 399–443.

Dasgupta, Shamik. [2018] "Realism and the Absence of Value." *Philosophical Review*. Vol. 127. 279-322

Eklund, Matti. [2017] *Choosing Normative Concepts*. Oxford: Oxford University Press.

Enoch, David. [2012] *Taking Morality Seriously*. Oxford: Oxford University Press.

Gibbard, Alan. [2003] *Thinking How to Live*. Cambridge: Harvard University Press.

Hare, R. M. [1997] *Sorting Out Ethics*. Oxford: Clarendon Press.

Horgan, Terence, and Mark Timmons. [1992] "Troubles on Moral Twin Earth: Moral Queerness Revived." *Synthese*. Vol. 92. 221–60.

Jackson, Frank. [1991]. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101: 461–482.

----- [1998] *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. Oxford: Clarendon Press.

Lord, Erol. (2017). "What You're Rationally Required to Do and What You Ought to Do (Are the Same Thing!)", *Mind* 126: 1109–1154.

McPherson, Tristram, and David Enoch. [2017] "What do you mean 'This isn't the Question'?" *Canadian Journal of Philosophy*. Vol. 47. 820–40.

Mill, John Stuart. [1863] *Utilitarianism*. London: Parker, Son, and Bourne, West Strand. Available online at <https://archive.org/stream/a592840000milluoft#page/n3/mode/2up>

Nietzsche, Friedrich. [1887] *On the Genealogy of Morality*. Available online at:

Nagel, Thomas. [1997] *The Last Word*. New York: Oxford University Press.

Risberg, Olle. [Manuscript] "Ethics and the Question of What to Do."

Scanlon, T.M. [2014] *Being Realistic about Reasons*. New York, NY: Oxford University Press.

Sepielli, Andrew. [orthcoming] "Subjective and Objective Reasons," in *Oxford Handbook of Reasons and Normativity*, ed. Star, D., Oxford University Press.

Smith, Holly. (2018) *Making Morality Work*. Oxford: Oxford University Press.

Streumer, Bart. [2013] “Can We Believe the Error Theory?” *Journal of Philosophy*. Vol. 110.
194–212.

Vermaire, Adam. [Forthcoming] “Against Schmought.” *Journal of Philosophy*.