

Introduction

Metaethics takes up questions about the semantics, metaphysics, and epistemology of the ethical domain. Stepping back from engaged ethical practice itself, we ask higher-order questions about the sort of meaning ethical terms and statements have, about the existence and nature of ethical properties and facts, and about whether and how we can have knowledge of them. Ethical realism is the metaethical view that the ethical domain is to be understood on a *realist* model, providing realist answers to such questions. But what exactly counts as a “realist” construal of the ethical domain? What exactly must one be realist about and what does that involve? And why does it matter?

How we answer these questions will turn partly on what we think makes for the most illuminating classificatory scheme in metaethics, given our primary theoretical concerns. This leads to some variation in taxonomic choices, reflecting different priorities. Certain starting assumptions may make a difference too. Some might assume, for example, that realism in the ethical domain should closely mirror realism in other domains, such as mathematics or the sciences, with analogous commitments regarding meaning, truth, and objectivity. Such an approach might be misleading, however, if it turns out (as I will argue) that there are elements unique to ethics that require additional theoretical treatment by any non-deflationary realism *about ethics*. Ethical realism may therefore have a distinctive profile, building in more than other forms of realism do.

We'll begin, in Section 1, by laying out the broadest conception of ethical realism – what I'll call the “minimal core characterization” involving two fundamental claims – and then proceed to examine the reasons for adding two additional claims to arrive at a narrower characterization that is both more useful in dividing the conceptual space of metaethics and arguably more deserving of the realist label. This more filled-out characterization reflects the most common understanding of ethical realism in contemporary metaethics, given various complications introduced in recent decades. I will go on to argue in Section 2, however, that even this more refined conception remains too broad: for it is possible to meet those first four conditions while still failing to capture certain apparent elements of ethics in a way that a *realism about ethics* plausibly should. This position will be controversial, but I believe a strong case can be made for it.

In Section 2 we will go on to explore what seems to be missing so far and add further conditions to our characterization, arriving at a more adequate paradigm of ethical realism, as developed in Section 3. Given, however, that it would be revisionary to restrict our use of “ethical realism” to views that build in so much, we will continue to count the broader class of views satisfying the first four

conditions as forms of ethical realism, and refer to the more robust version as “ardent ethical realism.”¹ We will conclude, in Section 4, with an examination of support for and challenges faced by ethical realism, especially in its more robust form.

While this Element is designed to be accessible to a broad readership, without presupposing prior knowledge of metaethics, it goes into some depth on issues that will be of interest to readers already familiar with the field as well. Readers new to metaethics can safely ignore many of the footnotes and are encouraged to consult other Elements in this series for detailed treatments of rival metaethical views that can only be sketched here in the course of situating and elucidating ethical realism.

1 What Is Ethical Realism?

1.1 The Minimal Core Characterization of Ethical Realism

Consider a typical ethical belief you might hold, such as the belief that human trafficking is morally wrong. On the face of it, such beliefs are likely to strike you both as *meaningful* and as capable of being *true* in the same sense in which ordinary descriptive beliefs can be true. In this respect, ethical beliefs or claims appear to function very much like other beliefs or claims. Ethical realism begins with this plausible idea that we use ethical language to make claims that can be straightforwardly true or false, just as with claims in other spheres of discourse, and that at least some positive ethical claims are in fact true. On this approach, then, there are some ethical truths alongside truths in the sciences and mathematics, for example. We can express these ideas more precisely with the following two claims:

(1) **Representational Content:** Ethical sentences, such as “human trafficking is morally wrong,” express propositions that can be straightforwardly true or false, just as with ordinary descriptive sentences. More specifically, when construed literally they are true or false by virtue of their *representational ethical content* – that is, content that represents an ethical fact or state of affairs as obtaining, as by representing human trafficking as having the property of being morally wrong. Ethical sentences and propositions, just like descriptive ones, are *truth-apt* due to their representational content, and ethical claims purport to state ethical facts. Likewise, our ethical judgments are *cognitive* states – beliefs – with the same representational ethical content. Ethical claims express such beliefs, which can likewise be straightforwardly true or false, just as with other beliefs.

¹ I borrow the expression “ardent realism” from Matti Eklund (2017), with qualifications discussed later. Eklund does not himself endorse the view.

(2) **Truth:** At least some positive ethical sentences, and the propositions expressed by them, are in fact *true*, and straightforwardly so, in the way ordinary descriptive sentences or propositions are: they are true because the world is in fact as it is represented to be by the propositional content in question. The difference between ethical and ordinary descriptive sentences or propositions is that what is represented by the content of the former is an *evaluative* or *normative* state of affairs (e.g., that human trafficking is morally wrong) and not simply a descriptive state of affairs (e.g., that human trafficking causes suffering). Since some positive ethical claims are true, there are likewise *ethical properties and facts*. It is in virtue of such properties and facts that some of our positive ethical claims succeed in being *true*, that is, when they accurately represent a state of affairs involving the instantiation of ethical properties.

These two theoretical commitments constitute what we may call the “minimal core characterization” of ethical realism, and they already set realism apart from several competing metaethical views.

The first commitment, when properly understood, excludes traditional *non-cognitivism* and its more recent outgrowth, non-realist (or anti-realist or quasi-realist) *expressivism*. According to such views, ethical language and concepts have a fundamentally different function from ordinary descriptive language and concepts. Instead of being used to represent ethical facts or states of affairs – such as the fact that human trafficking is wrong – they are primarily used *to express certain attitudes or commitments or plans* – such as an attitude of disapproval toward human trafficking or a commitment to avoid it – *without* thereby purporting to represent any ethical state of affairs as obtaining (Gibbard 1990, 2003, 2006; Horgan and Timmons 2000, 2008). The first realist claim rejects this construal of how ethical language and concepts work, insisting that ethical judgments have representational content just as other familiar judgments do (even if they can obviously *also* be used to express attitudes or commitments).²

² This point is complicated by the fact that sophisticated non-realist expressivists have adopted minimalist construals of “truth,” “proposition,” “property,” “belief,” and perhaps even “representation” that *derivatively* enable them to *say*, in the end, most of the same things realists say about these things, muddying the difference between the positions (Dreier 2004). For example, on a minimalist conception of truth, to say that it is *true* that murder is wrong (or that “murder is wrong” is *true*) is nothing more than to say that murder is wrong. So anyone who judges that murder is wrong is entitled to say that it is *true* that murder is wrong – even on an expressivist account of moral judgment, where it is simply a matter of having or expressing a certain attitude or commitment. And if murder’s having the *property* of being wrong is understood to mean no more than that wrongness is truly predicated of murder, then again judging that murder is wrong (even under an expressivist construal) entitles us to speak of murder’s “having the property” of being wrong. So the expressivist can say, in a sense, that there are moral properties and moral truths. There can similarly be minimalist talk of “believing” murder to be wrong insofar as one takes it to be wrong, and even of thereby “representing” murder as having the property of being wrong. The difference, however, is that while the expressivist may be able to say these realist-sounding things in a minimalist sense at the end of the day, their *route* to this is very

The second realist claim, about the existence of ethical truths, excludes *error theory* or *metaethical nihilism*. According to such views, all positive ethical claims – such as the claim that human trafficking is wrong – are false. Unlike the expressivist views mentioned previously, error-theoretic views agree with the realist’s first claim about the representational function of ethical language and concepts. The problem, they maintain, is that positive ethical claims are defective insofar as they presuppose the existence of certain properties that (according to error theorists) do not in fact exist, such as objective values, thus making all such claims false (Mackie 1977).³ The realist’s second claim rejects this error-theoretic assessment, claiming instead that at least many positive ethical claims are straightforwardly true rather than saddled with error. This claim also likely excludes crude forms of subjectivism or relativism involving such distorted interpretations of ethical claims that they flout platitudes associated with the meaning of ethical concepts, thus failing to make any ethical claims come out as true when construed *literally* (Sayre-McCord 1988).

Now those who are primarily concerned with the issues in claims one and two, and so with excluding expressivism, error theory, and crude forms of subjectivism or relativism, might be happy to stop with this minimal core characterization of ethical realism. Notice, however, that nothing has yet been said about the *source* of ethical truth. In particular, there is so far no requirement that it be *objective* rather than subjective (as long as the account respects platitudes constitutive of the meanings of ethical concepts). On the minimal core characterization, ethical realism is ecumenical on this front.⁴ This approach therefore still leaves us with an extremely heterogeneous collection of metaethical views under the umbrella of ethical realism – including many that are commonly placed into distinct metaethical categories typically viewed as *rivals* to realism, such as various forms of ethical constructivism.⁵

different from the realist’s. The nonrealist expressivist says these things in a minimalist sense after providing an account of the semantics of moral discourse radically different from that of ordinary descriptive discourse, whereas the realist takes the chief semantic role of moral language to be the same as with ordinary descriptive language, such that moral language is used in the first instance to make assertions about the instantiation of properties that are as metaphysically robust as any others (Copp 2007, ch. 5). Or as Dreier (2004) emphasizes, whereas nonrealist expressivists do not *need* to make any appeal to any notion of moral properties in explaining what it is to make a moral judgment, realists must do so.

³ Although Mackie (1977) is standardly understood to be the paradigm error theorist, Selim Berker (2019) argues that Mackie himself did not in fact espouse error theory as formulated earlier.

⁴ Sayre-McCord (1988, 16) holds that “realism is not solely the prerogative of objectivists,” and takes any independence of ethical truth to be relevant only insofar as it might bear on the second claim (since accounts proposing certain kinds of dependence on us would fail to secure the truth of ethical claims construed literally).

⁵ See, for example, Sharon Street (2006, 2008b), who explicitly presents her subjectivist metaethical constructivism as an *anti*-realist view that avoids the Darwinian Dilemma she raises against realist views.

For example, *subjectivist constructivist* views allow for positive ethical truths but take them to be grounded in facts about hypothetical attitudes, desires, or other responses we would have if we applied some sort of idealizing procedure to our existing psychology; ethical truths are constructed from or derived *via* some such procedure applied to our mental states or activities. The procedure in question is meant to be specifiable in an ethically neutral way – for example, involving procedural deliberative exercises starting from our existing attitudes with improved empirical information – so that ethical truth can be fully explained without circularity, appealing only to psychological and procedural facts.⁶ Similarly, ideal observer or appraiser views seek to derive ethical truths from the responses or judgments of a hypothetical agent occupying some ideal perspective (again characterized without ethical assumptions). And neo-Kantian constructivist or *constitutivist* views take ethical truths to be a function of principles bound up with procedures constitutive of the exercise of rational agency as such. On this approach, ethical truths are *constructed* through such practically necessary procedures rather than being already *there to be discovered* in the world in the way that biological or physical truths are (Korsgaard 1996, 2009).

These views seem to satisfy the minimal core characterization of ethical realism: they are compatible with the representationalist semantics and cognitivism given in the first claim and they allow for the existence of positive ethical truths as posited by the second claim.⁷ But again, they are also generally viewed today as alternatives to ethical realism rather than as paradigms of it, and for good reason. If we construe ethical realism as broadly as the minimal core characterization proposes then we wind up capturing everything from subjectivist constructivism (with ethical facts constrained by our desires or attitudes) or neo-Kantian constitutivism (with ethical facts arising only through procedures associated with our exercise of agency) to a robust platonic objectivism (with ethical facts rooted in transcendent, timeless values). So unless our principal theoretical concerns in metaethics are just with the issues laid out in the minimal core characterization, it is doubtful that such a heterogeneous collection of views makes for a particularly useful metaethical category.

For many of us at least, issues concerning the source of ethical truth or the nature of ethical facts are no less central to our metaethical concerns than the semantic issues emphasized in the minimal core characterization. This is reflected in the association of ethical realism with views that used to be called

⁶ See Williams (1981) for the classic employment of this strategy in accounting for normative reasons.

⁷ The situation with neo-Kantian constructivism is actually a bit tricky: it may wind up violating the second claim, as discussed below.

“ethical objectivism,” embodying a commitment not only to the basic claims about meaning and truth but also to some notion of *objective ethical reality*. While it is complicated what exactly this means, as we’ll see, the rough idea is that there exists a set of ethical properties and facts that are *there to be discovered or recognized* – not necessarily in the world *apart* from human beings and relevant facts about human life, but at least *there in the human-involving world* in a way that does not involve being constructed from elements or acts of human psychology in the ways described earlier.

The difference we’re after is clearest in the case of neo-Kantian constructivism. On this view, while we may speak of ethical truth, it is *derivative*. We begin with certain procedures or deliberative rules that are allegedly necessary for exercising agency at all, and then we identify principles allegedly bound up with those procedures (such as Kant’s categorical imperative), which then count as “correct” or “true” *not* because they track existing ethical facts but simply because they spring from those practically necessary procedures; such principles and their implications (e.g., that we have a duty not to lie) can then be regarded as “ethical truths or facts” in this derivative sense rooted in necessary agential procedures. By contrast, an ethical realism that incorporates objectivity begins by positing *nonderivative* ethical facts that serve as the truth-makers for true ethical claims. If there are correct procedures for answering ethical questions, then those procedures will count as correct because they or the principles they yield *accurately track* those existing ethical facts (not because they play a practically necessary or constitutive role in exercising agency, for example).⁸

The situation is a little different with subjectivist constructivist views that instead apply idealizing procedures, such as deliberation with full empirical information, to desires or attitudes. Here the issue isn’t that the notion of truth itself is taken to be derivative in ethics, as on the neo-Kantian approach mentioned earlier, but just that the various ethical truths have as their truth-makers a complex set of facts involving our subjective states and the results of

⁸ Korsgaard (1996, 36) characterizes her neo-Kantian constructivism as “procedural realism,” since it allows for ethical truths or facts, though only in this procedurally-based way; and she contrasts this with “substantive realism,” which takes ethical facts to exist in a more straightforward (nonderivative, procedurally independent) way. Ethical realism is generally understood as substantive realism, which means that we need to add a further condition to the minimal core characterization to exclude merely procedural views. It might be argued, however, that the minimal core characterization already excludes neo-Kantian procedural realism: the second realist claim may eliminate such views already on the grounds that they fail to make some ethical claims *literally and straightforwardly true or false*, since they employ a nonstandard, derivative notion of truth. Indeed, Korsgaard’s view makes the notion of ethical truth a theoretical *afterthought*, given the view’s radically practical orientation in reducing normativity to a form of practical necessity for exercising agency (FitzPatrick 2005). In any case, our characterization of ethical realism needs to exclude such approaches, either with the second claim or through the independence condition added later.

hypothetical procedural exercises. This isn't offering a revised notion of *what it is* for certain principles or claims to count as *true*, but just a subjectively oriented claim about what *grounds* ethical truths or facts, based on what ethical claims are allegedly really about – namely, some idealized function of our mental states. Still, this kind of dependence of ethical truths or facts on contingent facts about our desires and attitudes stands in stark contrast with views that posit ethical facts that are not dependent on our psychologies in that way.

For all these reasons, it is theoretically useful to build more into our characterization of ethical realism than we get from minimal core characterization, adding constraints on the source or grounds of ethical truth to capture some idea of *objective ethical reality*. We'll take this up in the following subsection, though I'll argue in Section 2 that even this is insufficient to capture the objectivity that ultimately matters to non-deflationary realism about ethics. But this is still an important step along the way.

1.2 Adding Objectivity/Independence

I have proposed adding to our characterization of ethical realism an objectivity or independence condition to capture the idea that the source or grounds of ethical truth are in some interesting way *independent of us*. How exactly to characterize this condition, however, is complicated.

Ethical truths are, after all, largely *about us* (how it is good and right for us to live and to treat each other) and on any plausible view their grounds will often involve facts about us. Part of what *makes* human trafficking wrong (and hence makes true the claim that it is wrong) is the fact that it causes suffering in its victims – a fact obviously not independent of human beings. Similarly, the truth of ethical claims about moral blameworthiness or desert will depend partly on facts about agents' epistemic states (knowledge or ignorance of relevant circumstances) and intentions and motives – again, all facts about human beings and so hardly “independent of us.” Many particular ethical truths will even depend on social conventions. If someone does something wrong in making a gesture because this amounts to a gratuitous insult, then part of what makes the giving of that gesture wrong will be facts about its conventional meaning in the social context.

The independence we're seeking, therefore, requires careful unpacking. As a first pass, we might characterize it by saying that despite the aforementioned forms of dependence, there is a true or correct set of *ethical standards*, from which particular ethical truths are derived, and these standards will not themselves depend on such things as our ethical beliefs or conventions, or on our contingent desires, attitudes, commitments, choices, and so on – either directly

or indirectly *via* idealizing procedures of the sort described earlier. Russ Shafer-Landau helpfully formulates this as a *stance-independence* condition, which we may add to our initial two claims:

(3) **Stance-Independence:** There are positive “[ethical] truths that obtain independently of any preferred perspective in the sense that *the [ethical] standards that fix the [particular ethical] facts are not made true or correct by virtue of their ratification from within any given actual or hypothetical perspective,*” such as one employing idealizing conditions or procedures (Shafer-Landau 2003, 15).

With the addition of this condition, ethical realism takes on a commitment to a set of *objective ethical standards* in this stance-independence sense. We’ve said that a central wrong-making feature of human trafficking – a feature that makes it wrong and hence makes true the claim that it is wrong – is the fact that it causes human suffering. This latter, empirical fact about the effects of human trafficking – call it “CHS” – is obviously not something that is independent of humans. But the stance-independence condition is drawing attention to a different fact posited within this picture: namely, the *evaluative or normative meta-fact that CHS is wrong-making* (cf. Dancy 2006). Our third claim, about stance-independence, is concerned with such ethical meta-facts, because such meta-facts (suitably generalized and qualified) are what constitute the ethical standards referred to in our third claim.

To keep things simple, let’s focus on ethical standards concerning *rightness* of action. An account of these standards will articulate a set of claims about which features of actions tend to be right-making and which tend to be wrong-making, how various circumstances might affect these tendencies (e.g., making for exceptions to typical wrong-makingness in certain circumstances), and how to weigh complex sets of factors to arrive at an overall assessment of an action as right or wrong based on all its features and circumstances. Such an account of the ethical standards thus tells us what acts must be like – the conditions they must meet with regard to their relevant features – in order to qualify as right and avoid coming out as wrong. On some views, the ethical standards might be given with a single ultimate rule or principle, while on others – such as a virtue-theoretic account – they may be far more complex and less codifiable, involving a plurality of considerations and a need for educated judgment to determine how exactly they interact to yield all-things-considered verdicts in complex cases.⁹

⁹ We may here set aside debates over the degree to which the ethical standards are codifiable in terms of one or more principles, though I will tend to discuss them using a pluralistic, virtue-theoretic framework, which I favor.

A nonethical analogy may help to illustrate the idea of evaluative standards and their role in connection with evaluative properties (further developed in Section 2). Consider the standards for being a good move in chess. An account of this will articulate a set of claims about features of a chess move that tend to count as good-making (e.g., increasing control of the center squares) and features that instead tend to count as bad-making (e.g., sacrificing material without positional or tactical gain), along with the differences that various circumstances can make and how to weigh complex sets of factors to arrive at an overall assessment. The standards constituted by such propositions tell us what a chess move needs to be like in order to constitute a good move and avoid being a bad one (allowing room for gradations and qualifications). With such a set of standards in hand, we are then in a position to evaluate a given chess move as a good or bad one based on whether its overall features, in the circumstances, make it meet or violate those standards.

In the case of chess, the standards that fix the particular facts about good or bad chess moves are grounded in – or made true or correct by virtue of – the rules and aims internal to the game. Obviously, the model for ethics will be different, and our characterization of ethical realism has not yet provided one. What we are adding with the third claim, however, is at least a negative point about how *not* to understand the source of the ethical standards: they are not grounded in (or made true or correct by virtue of) being ratified from within some actual or hypothetical perspective, as on constructivist views. It tells us, for example, that whatever exactly it is that grounds the correct set of standards for ethically evaluating human action, it is not to be found in any such fact as that we would endorse these standards if we underwent certain deliberative exercises starting from our existing attitudes with full empirical information. According to the realist's third claim, whatever grounds ethical standards, it is not any sort of *stance-dependent underwriting*, as by deriving the standards from the application of ethically neutrally specifiable sets of conditions or procedures to our beliefs, desires, attitudes, agential capacities, conventions, and so on.¹⁰

Importantly, this is not to deny that ethical standards might still be grounded largely in facts about human beings or human life. It is consistent with the third

¹⁰ We noted earlier that neo-Kantian constructivism might in a way be further from ethical realism than subjectivist constructivism, insofar as it employs a derivative notion of truth. There is also, however, a sense in which it may be *closer* to ethical realism in connection with the objectivity or independence issue: for although neo-Kantian constructivism violates the third condition we've now laid down, just as subjectivist constructivism does, the former at least avoids making certain core ethical facts (such as moral duties) dependent on the details of our contingent psychological states: we will (allegedly) wind up with the categorical imperative regardless of what contingent desires or attitudes we start out with, which thus secures *independence* from at least the contingent details of human psychology (Korsgaard 1996).

claim, for example, that there is a correct set of ethical standards consisting partly in requirements to respect human dignity, where human dignity is an evaluative or normative status grounded in facts about human rational capacity. To be consistent with the third claim, we simply have to hold that this fact about the evaluative or normative ramifications of human rational capacity is itself an objective evaluative or normative fact about it. All that is denied is that this fact itself is somehow dependent on being ratified from within some actual or hypothetical perspective. It is thus worth emphasizing again that the third claim does not commit ethical realism to any transcendental or radically independent grounding of the ethical standards wholly apart from facts about human life: it simply insists that the grounding of the ethical standards is a stance-independent matter in the sense described.¹¹

1.3 Metaphysical Matters

Assuming that ethical realism posits some sort of objective or stance-independent ethical reality, we may now ask whether this brings with it *metaphysical* commitments. In particular, does ethical realism require us to add such things as ethical properties and facts to our ontology, as part of the fabric of (relevant parts of) the world or reality? It is natural to think that it does, and that when our ethical claims are true they are *made* to be true by accurately representing that ethical reality. For example, a realist might think that part of what grounds the ethical standard forbidding cruelty is the fact that the suffering of sentient beings is intrinsically bad – an evaluative fact about the world consisting in the possession of an evaluative property (badness) by a worldly phenomenon (suffering). In this picture, the ethical truth that causing suffering for amusement is wrong has this worldly evaluative fact as (at least part of) its *truth-maker*: the ethical claim is true because of how the world is, ethically speaking – specifically, because of a real evaluative property possessed by suffering.

Some metaethicists, such as Derek Parfit, however, embrace objective ethical truths but deny that they have any such ontological implications for the world or reality, or indeed that they have (or need) any truth-makers at all, any more than logical or mathematical truths do. More precisely, these theorists will allow that some nonbasic ethical truths might have “truth-makers” in the sense that they are derived from more basic ethical truths: for example, the truth of “causing

¹¹ Note that the ethical standards might still perfectly *correspond* to what would be ratified from some ideal perspective: the claim is just that this is not what would *make* the standards the true or correct ones (cf. Shafer-Landau 2003, 16). It would be the other way around: the perspective that wound up endorsing those standards would count as ideal because it accurately tracked the true or correct standards and their implications.