Relaxing about Moral Truths

Christine Tiefensee

Abstract   As with all other moral realists, so-called relaxed moral realists believe that there are moral truths. Unlike metaphysical moral realists, they do not take themselves to be defending a substantively metaphysical position when espousing this view, but to be putting forward a *moral* thesis from *within* moral discourse. In this paper, I employ minimalism about truth to examine whether or not there is a semantic analysis of the claim 'There are moral truths' which can support this moral interpretation of one of moral realism's key theses. My results are both discouraging and encouraging: Whilst I will argue that the claim 'There are moral truths' cannot be shown to be both moral and capable of demarcating relaxed realism from irrealism on the basis of a convincing semantic analysis that would be compatible with relaxed commitments, the moral interpretation of moral realism can be secured by modifying our understanding of what distinguishes relaxed realism from error-theoretic irrealism. Yet, we will see that this moral interpretation of moral realism does not 'tumble out' of the semantics provided for its central claims. Rather, hard work needs to be done before we can fully relax.

Keywords   Relaxed moral realism • Minimalism • Semantics • Metaethics • Metaphysical moral realism

1 Introduction

As with all other moral realists, non-metaphysical moral realists—or 'relaxed' realists, as I will call them—submit that there are objective moral truths. In this vein, Scanlon (2014: 16) defends "the idea that that there are irreducibly normative truths about reasons for action", Dworkin (2011: 7) believes "that there are objective truths about value", Nagel (1986: 144) rejects arguments to the effect that "there are no irreducibly normative truths", Parfit (2011: 559) defends the view that "there are some irreducibly normative reason-involving truths" and Blackburn (1998: 319) replies "No, no, no" to the charge that his position on moral truths implies that "there aren't any really".

Unlike metaphysical moral realists, relaxed realists do not take themselves to be defending robustly metaphysical positions when espousing these views. Instead, they believe that they are putting forward *normative* or *moral* theses from *within* normative and moral discourse. As such, Scanlon (2014: 20) argues that "the truth values of statements in [the normative domain], including existential statements, are properly settled by standards internal to that domain", while Blackburn (1998: 311) explains that the claim 'Moral truths are mind-independent' "can *only* [involve giving] a moral verdict on [a particular] situation"; to Nagel (1986: 144), the "claim that certain reasons exist is a normative claim", whereas Dworkin (2011: 10) explicates that theories of moral knowledge and truth "are drawn from within morality—they are themselves moral judgments"; Kramer (2009: 1) states on the very first page of his book that "[m]oral realism … is a moral doctrine" and even Bart Streumer (2017: 127), himself an error theorist, is

"inclined to agree … that the sentence '[Normative] Realism is true' expresses a normative judgement …".

In a nutshell, then, relaxed moral realists submit that there are moral truths and that these truths are objective, whilst being committed to the moral, discourse-internal interpretation of these theses.[1] Surprisingly, though, whilst quite a bit has been written on the moral interpretation of morality's objectivity,[2] next to nothing has been said about a semantics of the more general thesis

(T)   There are moral truths

that would show (T) to be a discourse-internal, moral claim.[3]

There are two reasons why it is crucial for relaxed realists to make amends for this omission. The first is obvious: If (T) encapsulated one of moral realism's core theses, but there were *no* semantic analysis of (T) available according to which (T) would qualify as moral, relaxed realism would be a non-starter. The second uses the same thought to relaxed realists' advantage, stating that if (T) formed one central thesis of moral realism and (T)'s semantic analysis *did* classify (T) as moral, relaxed realists could pile up the pressure on metaphysical moral realists by urging them to explicate in which way their defence of this *moral* claim (T) could possibly amount to presenting a *metaphysical* position.[4]

---

[1] Relaxed moral realism as defined here thus includes the positions of Thomas Nagel (1986), Tim Scanlon (2014, 2017) and Derek Parfit (2011), the quasi-realism of Simon Blackburn (1998), Allan Gibbard (2003) and Huw Price (2013), together with anti-Archimedeanism as defended by Matthew Kramer (2009) and Ronald Dworkin (1996, 2011). Any disagreement that holds amongst these philosophers can be ignored for the purposes of this paper. For metaphysically substantive moral realism, see Enoch (2011), Brink (1989) and Shafer-Landau (2003).

Since I will focus exclusively on moral claims and moral discourse, in what follows I will use the terms 'discourse-internal' and 'moral' interchangeably. Importantly, this is not to deny that we find different kinds of claims within moral discourse, with some being more abstract and others being more concrete. See Kramer (2009: 11) for more on this point.

[2] Here, I think of moral objectivity in terms of mind-independence (see Blackburn 1998: 311, Kramer 2009: ch. 2 and Dworkin 1996). Given these prior studies about the moral reading of mind-independence, in this paper I will focus exclusively on the thesis that there are moral truths, rather than the thesis that there are mind-independent moral truths. As such, I am primarily interested in moral realism's opposition to error-theoretic irrealism, which rejects the existence of moral truths, rather than its opposition to versions of constructivism, which accept the existence of moral truths but deny that these are mind-independent. My project can thus also be understood as asking in which way anyone who endorses the existence of moral truths—including realists *and* constructivists—can be understood as putting forward a discourse-internal view, *irrespective* of whether they conceive of these truths as mind-dependent or mind-independent. Showing that views on mind-independence are also discourse-internal is thus an additional project as pursued by, for instance, Blackburn, Dworkin and Kramer.

[3] Simon Blackburn (1998: 76) has come closest when suggesting a disjunctive analysis of a claim such as 'What Johnny said this morning was true'. I will build on and expand this disjunctive analysis when examining the more general thesis 'There are moral truths'.

[4] It might be argued that metaphysical moral realists could respond that their endorsement of (T) amounts to *both* a metaphysical *and* a moral position (Berker 2017: 771). Nonetheless, if there is to be a meaningful contrast between relaxed and metaphysical moral realism, and if metaphysical realists now hold that (T) is both moral and metaphysical, they will still have to make explicit in which way this metaphysical interpretation *adds* to or *differs* from relaxed realists' moral interpretation.

At first sight, providing a semantics for (T) which supports (T)'s moral classification does not seem too difficult a task. After all, relaxed realists are generally also minimalists about truth,[5] and thus hold that there is no more to understanding the concept of truth than to grasp that

(T*)   It is true that snooping is wrong

and asserting

(M)   Snooping is wrong

are conceptually equivalent,[6] in that whenever we can assert (T*), we can also assert (M) and *vice versa*. To the thesis that there is no substantive difference between (T*) and (M), add the information that (M) is moral and we arrive at the result that (T*) too is most plausibly understood as a moral claim—just as relaxed realists intend. As Blackburn (1998: 78, 296) appositely puts it, minimalism thus allows us to add as many addenda to a claim such as (M) as we like—that (M) is true, that it states fact, that (M) really and truly represents reality, etc.—and yet never leave the moral ground from where we started, namely: (M).

Given this fast-track, minimalism-based path to (T*)'s moral classification, we might expect minimalism to provide the same kind of straightforward support to relaxed realists' moral interpretation of (T):

(T)   There are moral truths.

For instance, Kramer (2017: 206-207) appears to do so when declaring that "[o]f central importance to [relaxed moral realism] … is its minimalist account of truth and reality and facts and properties. Precisely because those notions and cognate notions are all to be construed minimalistically, debates over the existence or nature or objectivity of moral principles and properties are substantive ethical debates."[7] However, things are not that simple. Two reasons explain why.

Firstly, in contrast to (T*), (T) does not state which specific proposition is true and thus falls into the context of indirect, rather than direct truth-ascriptions.[8] Indirect truth-ascriptions are, in turn, to be understood on grounds of the logical function that the truth-predicate plays in our language, which is to make possible semantic ascent by allowing us to generalise over and indirectly refer to other claims. Given as much, minimalists will tell us that the thesis 'There are moral truths' is no more than a different

---

[5] As far as I am aware, all relaxed realists adopt minimalism about truth in some shape or form. For instance, see Blackburn (1998: 78), Price (2013: 55), Scanlon (2014: 43), Parfit (2011: 756, n. 295), Kramer (2009: 261) and, more hesitantly, Dworkin (2011: 173). For more on minimalism, see Horwich (1998) and Soames (2003).

[6] I loosely follow Soames (2003: 372, 382, fn. 4) here, who explains that the propositions 'It is true that *p*' and '*p*' are conceptually equivalent, meaning that they "are trivial, necessary and apriori consequences of one another" in that "there is a way of apprehending the proposition that p is true … such that when the proposition that p is true is apprehended in this way, the agent is able, apriori, to infer p."

[7] Tellingly, the examples Kramer lists in support of this statement all fall into direct truth-contexts such as (T*), not indirect truth-contexts such as (T) in which I am interested here.

[8] For the difference between direct and indirect truth-contexts, see Soames (1997: 3).

formulation of the claim 'At least one moral claim is true' which, in turn, can be taken as equivalent to an infinite disjunction comprising only moral disjuncts, such as:

(D1) Snooping is wrong or snooping is not wrong, *or*

Betraying someone's confidence is permissible or betraying someone's confidence is not permissible, *or*

Caring for others is right or caring for others is not right, *or*

…

Since it is moral claims that can make (T) true, (T) must generalise over all claims p_m that qualify as moral. And since two relaxed moral realists can agree that moral truths exist whilst disagreeing about which moral truths these are, (T) must be analysed disjunctively. (D1) captures both these insights.

Analysing (T) in terms of (D1), then, has great *prima facie* plausibility. At the same time, it remains obscure in which way (D1) is to classify as moral. After all, (D1) appears to be nothing more than a collection of instances of the law of the excluded middle ($p_m \lor \neg p_m$), and thus a logical rather than a moral truth.[9] And since (D1) is supposed to be equivalent with (T), the same holds for (T). Consequently, in contrast to

---

[9] I should address two worries at this point. Firstly, could a truth such as ($p_m \lor \neg p_m$) not be both logical and internal to moral discourse? Yes, and nothing I will say here will deny this. Nonetheless, even if it were agreed that ($p_m \lor \neg p_m$) is both logical and discourse-internal, we can and should still ask *on which grounds* ($p_m \lor \neg p_m$) can indeed be classified as a discourse-internal truth. As I explain in §2.1, I interpret this question as asking in which way ($p_m \lor \neg p_m$) can be supported on the basis of moral considerations and provide an answer by presenting two moral arguments to this effect. There, we will also see that although ($p_m \lor \neg p_m$) is usually categorised as a logical truth, it is not trivial in the same sense as ($p_m \to p_m$) might be. Rather, I will argue that since it is not trivial whether or not moral claims are determinately true, it is not trivial either whether or not ($p_m \lor \neg p_m$) holds within the moral domain.

Secondly, as stated in footnote 2, I am primarily interested in moral realists' thesis that there are moral truths, rather than their thesis that these truths are objective and mind-independent. However, it might be argued that if we merged both these theses into the claim (T_M) 'There are mind-independent moral truths', worries about its logical classification would no longer arise, as (T_M) might now be analysed along the lines of disjunctions such as (D_M) 'Either it is objectively the case that snooping is wrong, or it is objectively the case that snooping is not wrong'. (D_M), though, is not an instance of the law of the excluded middle; indeed, constructivists might want to reject both its disjuncts. Accordingly, it might be thought that by shifting focus from (T) to (T_M), my challenge to relaxed realism collapses. My response to this suggestion is twofold. Firstly, even if this analysis of (T_M) were correct, asking in which way the more general thesis (T) 'There are moral truths' can be seen as discourse-internal becomes no less significant for relaxed realists. Rather, as set out in footnote 2, since relaxed realists are committed to (T)'s moral interpretation, showing (T) to be discourse-internal remains one objective that they must achieve in order for their project to be successful. Secondly, it is not clear that (D_M) is the correct analysis of (T_M). Rather, it could be suggested that (T_M) must be understood along the lines of (D_M*) 'Either, it is the case that snooping is wrong and that snooping's being wrong is mind-independent, or it is not the case that snooping is wrong and that snooping's being wrong is mind-independent'. Admittedly, this analysis of (D_M*) is more complicated than (D1)'s. Still it is an instance of the law of the excluded middle, the negative disjunct of which would be endorsed both by error-theoretic irrealists *and* constructivists, albeit for different reasons. Accordingly, if (D_M*) were the correct analysis of (T_M), moving from (T) to (T_M) would still require us to ask in which way (D_M*) can be seen as moral. Given my interest in moral realism's opposition to error-theoretic irrealism,

(T*)   It is true that snooping is wrong,

it is by no means clear in which way

(T)   There are moral truths

qualifies as a moral, discourse-internal claim, as relaxed realists have us believe.

The second reason why (T)'s moral categorisation is not as straightforward as (T*)'s is that minimalism supports the moral classification of claims only when combined with an independent criterion that secures the moral categorisation of the claims which are called true. This is why I said earlier that (T*)'s conceptual equivalence with

(M)   Snooping is wrong

supports (T*)'s moral classification *when supplemented* with the information that (M) is moral. Now, (M)'s moral classification is a clear-cut case. (D1)'s classification and, as we will see, that of other claims is not. Hence, although Kramer is absolutely right in emphasising minimalism's importance for relaxed realism, minimalism by no means settles all questions in relaxed realists' favour.

Accordingly, much more needs to be said before we can arrive at a minimalism-friendly, relaxed semantic analysis of (T) that shows (T) to be moral. As I will suggest in the next section, three such analyses are conceivable. However, as I will also argue, none of them takes relaxed realists to where they want to be: Whilst the first no longer draws the correct demarcation line between realism and error-theoretic irrealism (§2.1), the second is incompatible with relaxed realists' commitments about moral discourse and thus unavailable to them (§2.2). And whereas the third is indeed compatible with relaxed realism, it must be rejected on independent grounds since it relies on unsustainable semantic assumptions about moral concepts (§2.3).

As I will suggest in §3, though, the conclusion to be drawn from this negative result is *not* to give up relaxed realism. Rather, it is to modify our understanding of what this relaxed position and its opposition to irrealism involves. Accordingly, enquiring into the moral classification of (T) forces us not only to explicate relaxed realism's semantic commitments, but also to recalibrate its opposition to irrealism.

Before I start, though, let me briefly address a worry you may have about my focus on (T)'s semantic analysis. That is, it might be thought that even if I were right about the unexpected difficulties that a minimalism-friendly semantics of

(T)   There are moral truths

encounters, relaxed realists need not be overly concerned by this finding as long as they can simply blank out (T) and instead focus exclusively on specific moral truth-ascriptions such as

(T*)   It is true that snooping is wrong.

Since, as I have indicated above, relaxed realists' path to (T*)'s moral categorisation is straightforward, shifting focus from (T) to (T*) would secure the moral categorisation of their view without having to become entangled in any hitches surrounding (T).

---

though, I will keep concentrating on (T) and (D1) rather than on (T_M) and (D_M^*). I am grateful to an anonymous referee from *Ergo* for pressing me on both worries.

Yet, this shift from (T) to (T*) will not do. For, although endorsing a specific truth-ascription such as (T*) is sufficient to qualify as a relaxed realist, it is not necessary. After all, whilst Anna, a relaxed realist, might hold that snooping is wrong, Bella, also a relaxed realist, might think that doing so is not wrong and instead believe that it is permissible. Indeed, Charlotte might even be a relaxed realist *without* endorsing *any* specific moral truth-ascription such as (T*), since although she is sure that some moral claim is true, she does not know which one this is. To understand what makes all three of them relaxed realists, then, we need to know which claim it is that Anna, Bella and Charlotte equally *endorse* despite their *disagreement* on specific truth-ascriptions such as (T*). Of course, we might rightly respond at this point that all three are clearly committed to holding *some* moral claim to be true, no matter which one this is. It is just as clear, though, that by giving this response, we have simply arrived back at their commitment to (T)! Accordingly, in order to understand what is involved in being a relaxed realist and how relaxed realism is itself a moral position, there is no way around providing a semantics of (T) which shows (T) to be moral. Let us examine next, then, what such a semantics could look like.

2 The Troubles Surrounding (T)

Our first stab at a minimalism-friendly semantics of (T) took (T) to be equivalent with (D1) which, in turn, looked just like a collection of instances of the law of the excluded middle, and thus a logical rather than a moral truth.[10] Hence, for a relaxed realist to resist this result, she must be able to show that (T) is discourse-internal and do so in a way that is compatible with her relaxed realism. There are two ways of establishing as much. Firstly, she can focus exclusively on (T) analysed along the lines of (D1) and show that (D1) does, after all, qualify as a moral truth. I will examine this approach in §2.1. Alternatively, she can attempt to modify (T)'s analysis in such a way that it no longer features instances of the law of the excluded middle ($p_m \vee \neg p_m$) as contained by (D1), and instead show (T) to be equivalent with a disjunction of the form ($p_m \vee q_m \vee r_m \vee \ldots$), where p_m, q_m, r_m etc. are moral claims. The general thought behind this strategy, then, is to argue that just as (T*) must count as moral because it is conceptually equivalent with a specific moral claim, (T) must qualify as moral because it is conceptually equivalent with an infinite *disjunction* of moral claims. The suggestions discussed in §2.2 and §2.3 follow this second path.

2.1 (D1) as a Moral Truth

Let us consider, then, in which way (T) analysed along the lines of (D1), and thus disjunctions of the form ($p_m \vee \neg p_m$), could be shown to be moral.[11] In doing so, we first

---

[10] As many readers will have noticed, this problem is the mirror image of the inconsistency challenge faced by moral error theorists: If negations of moral claims are themselves moral, moral error theorists cannot consistently endorse the truth of these negations and simultaneously hold all moral claims to be false. Hence, the same argument that accuses moral error theories of inconsistency can be used to accuse relaxed realism of stating a logical truth. Recent discussions of the inconsistency challenge thus also strongly inform the moves discussed in this section (Sinnott-Armstrong 2006: 32-37; Pigden 2007: 450-455; Olson 2014: 11-15; Streumer 2017: §54).

[11] You might have spotted a problem with regard to (T)'s analysis in terms of (D1), namely: error-theoretic irrealists can endorse (D1) by accepting its negated disjuncts. This worrying

need to take a step back and make explicit what has so far remained implicit, namely: What would have to be the case for (D1) to qualify as moral? More generally: What is it that grounds the moral classification of a claim *p*?

Here is one idea: *p* counts as moral iff it uses (rather than mentions) moral vocabulary. Let us call this the 'use criterion'. This criterion sounds *prima facie* plausible—after all, all paradigmatically moral claims, such as 'Snooping is wrong', 'Helping others is fair', etc. fulfil it. Still, it will not do for our purposes, as using moral terms is neither necessary nor, arguably, sufficient for moral categorisation. It is not necessary because there are claims that we would want to categorise as moral although they do not use moral vocabulary. For instance, the claim 'Everything you said is true' does not use moral terms but should be classified as moral if what you said included only paradigmatically moral claims. Using moral terms is, in turn, not sufficient for moral categorisation since there are claims whose moral categorisation is contested even though they do feature moral vocabulary. 'Snooping is wrong and snooping is not wrong' is one example of a claim that falls into this contested category; crucially, the disjunctions featuring in (D1), such as 'Snooping is wrong or snooping is not wrong', are further examples. The 'use criterion', then, is far too coarse-grained to help settle controversies about moral categorisation.

Here, then, is another proposal: Following Streumer (2017: 2-4) and Olson (2014: 15), relaxed realists could declare that *p* counts as moral iff it conceptually entails that some moral predicate is satisfied. Let us call this the 'conceptual criterion'. Again, this criterion has great *prima facie* plausibility, as all paradigmatically moral claims do entail satisfaction of some moral predicate. At the same time, by demanding conceptual entailment relations, this criterion sets the bar very high for moral classification. Indeed, for the purposes of relaxed realists, it sets it *too* high. For, at least at first sight, (D1) does not clear it: Since (D1) contains negated disjuncts, it does not entail that some moral predicate is satisfied and thus will not qualify as moral according to the conceptual criterion. Now, it might be thought that closer inspection could overturn this result, namely by showing that a negated disjunct such as 'Snooping is not wrong' does conceptually entail a non-negated claim such as 'Snooping is permissible', in which case the conceptual criterion would, after all, be fulfilled. Yet, as I explain in §2.3, there are good reasons to believe that no such conceptual entailment relations between negated and non-negated moral claims obtain. Accordingly, the conceptual criterion will not do for relaxed realists either.

Instead, I believe that relaxed realists should adopt a criterion along the following lines: A claim *p* counts as moral iff we find out about *p*'s truth or falsity by engaging in moral reasoning.[12] As such, the claim 'Killing is wrong' qualifies as moral because we can

---

observation is correct and I will return to it at the end of this section. Over the next few pages, though, I will bracket this worry to focus exclusively on (D1)'s discourse-internal classification. Two reasons explain why. Firstly, examining in which way (D1) could qualify as moral forces us not only to sharpen our understanding of the relation between moral discourse and the application of the law of the excluded middle, but also to explicate the criterion that relaxed realists could employ in order to classify claims as moral. This, in turn, will prove essential for later examinations of the relaxed/metaphysical debate and thus our understanding of how and on which levels relaxed realism can be opposed. Secondly, no matter how (T) is analysed, relaxed realists are committed to (D1)'s moral interpretation. Accordingly, if (D1) could not be shown to be moral, relaxed realists would be in trouble irrespective of our analysis of (T). The significance of establishing (D1)'s moral categorisation, then, reaches beyond the more limited focus of this paper.

[12] This first formulation no doubt needs to be further specified. For the purposes of this paper, though, it will do to stick with this rough version. Does this criterion require that *p*'s truth

discover whether or not killing is indeed wrong only by considering moral arguments for and against killing. In contrast, the claim 'A 250,000 volts shock will kill you' is not moral because it must be assessed as true or false by employing empirical methods, not by moral reasoning. Since this criterion thus pertains to the methods we employ in order to determine truth-values, let us call it the 'methodological criterion'.

In contrast to the use criterion, which does not help us settle contested cases, and contrary to the conceptual criterion, which does not support (D1)'s moral categorisation, the methodological criterion might give relaxed realists exactly what they need. For, as long as relaxed realists can successfully make the case that the truth of (D1), and thus of disjunctions of the form $(p_\mathrm{m} \lor \neg p_\mathrm{m})$, must be established on the basis of moral arguments, this criterion will classify (D1) as a moral truth. At least two such arguments are conceivable.

In a nutshell (and somewhat misleadingly), the first runs as follows: It would be morally good if $(p_\mathrm{m} \lor \neg p_\mathrm{m})$; therefore, $(p_\mathrm{m} \lor \neg p_\mathrm{m})$. Why it would be morally good if $(p_\mathrm{m} \lor \neg p_\mathrm{m})$ is not that hard to see: It is part of our moral lives that we constantly find ourselves confronted with moral decisions. We are torn between refraining from snooping around in a friend's diary and risking serious harm, between taking up a fulfilling job and looking after one's family, between having a disabled child and aborting it, etc. Being desperate to do what is right, we urgently need to know how we ought to act. Yet, if there were no answers to these questions—for instance, if 'Snooping is wrong or it is not wrong' did not hold—we could expect no moral guidance within our lives and would, therefore, be left all alone with the moral conflicts we face. Clearly, this would not be a desirable situation. Hence, it would be morally good if there were determinate answers to moral questions—that is, it would be good if $(p_\mathrm{m} \lor \neg p_\mathrm{m})$ were true.[13] Consequently, $(p_\mathrm{m} \lor \neg p_\mathrm{m})$ *is* true, or so this argument concludes. Let us call this the 'desirability argument'.

Phrased thus, the desirability argument obviously smacks of wishful thinking. After all, the observation that it would be morally good or desirable if some fact obtained does not generally entail *that* this fact obtains. Interestingly, though, this charge of wishful thinking might miss the mark if the fact which is judged to be morally good is itself moral.[14] To elaborate, assume that there are two moral theories, *A* and *B*, both of which could, in principle, be true. However, whereas *A* assigns determinate truth-values to most moral claims, *B* entails that hardly any moral questions have determinate answers. Hence, provided that it would indeed be morally good if $(p_\mathrm{m} \lor \neg p_\mathrm{m})$ were true, it would also be morally good if *A*, and not *B*, were true. Now, given that *A* is *morally better* than *B*, could it still be the case that *A* is nevertheless false and *B* is true? It is hard to see how this could be so. For, morality itself now provides a moral reason for *A* and against *B*, whilst morality would never speak for a moral theory that entails *false* moral verdicts and provides *wrong* moral advice. Accordingly, if this is right, there is an important *moral constraint* on moral theories, such that a moral theory cannot be true if there is another moral theory which is morally better. Consequently, if any of the theories *A* and *B* is to be true at all, it must be *A*, and not *B*. Generalising this result, we can thus see that what seemed like a mere case of wishful thinking turns out to be an *indirect moral* argument for (D1): Since only those moral theories pass the moral desirability constraint which entail that (most)

---

can be determined on moral grounds, or that it *must* be so determined? I return to this question below.

[13] Kramer (2009: 94) can be read as endorsing some such argument when declaring that moral "indeterminacy would be profoundly problematic morally if it obtained on a large scale".

[14] For further discussion, see Enoch (2009), Preston-Roedder (2014) and Sayre-McCord (Ms.), who I follow most closely here.

moral claims have determinate truth-values, no matter which one of these is true, it will entail that $(p_m \vee \neg p_m)$ (at least for most p_m). Given this moral argument for (D1), then, the methodological criterion will classify (D1) as moral.

The second moral argument for (D1) pursues the following route: Assume that the best method for finding out which moral claims and theories are true were one of reflective equilibrium, according to which our belief in the truth of moral claims and theories is justified only if these claims and theories cohere with well-considered moral judgements. Which judgements are well-considered and how to derive general moral theories from these judgements are once more *moral* questions, which require us to engage in substantive moral reasoning about what is good, just, impermissible or virtuous, say. Assume further that this process led in equally *morally* sound ways to two different reflective equilibria containing two different moral theories, C and D, where C and D reach conflicting results on moral questions such that whenever C judges a moral claim p_m to be true, D judges p_m to be false, and *vice versa*. What are we to make of this case? Given that there are two equally justified moral theories coming to contradictory moral verdicts, one plausible reaction would be to declare that we have no reason to believe that there are determinate answers to moral questions: In other words, we have no reason to believe that $(p_m \vee \neg p_m)$ is true.[15]

In contrast, assume that this reflective process did not lead to different equilibria which are morally on a par. Instead, there is only one, *morally* superior reflective equilibrium which provides moral reason to believe in the truth of a specific moral theory, E. Assume further that E is such that it provides a moral standard which generates determinate answers to moral questions. In this scenario, our belief in the truth of E, and thus also the truth of $(p_m \vee \neg p_m)$, *would* be justified.

Importantly, though, no matter whether we argue for or against $(p_m \vee \neg p_m)$, our case crucially proceeds from *within* moral discourse: We need to take a moral stance on how plausible certain moral intuitions are, how to resolve conflicts between moral judgements, if there is one, and only one, reflective equilibrium that is best supported by moral reasons, etc. Consequently, the methodological criterion will once more classify $(p_m \vee \neg p_m)$ and thus (D1) as moral. Let us call this the 'equilibrium argument'.

Both the desirability argument and the equilibrium argument doubtless require further explication. Still, they give us a good idea of how relaxed realists can back up the moral categorisation of (D1) on the basis of the methodological criterion by providing moral arguments for (D1)'s truth. If successful, this would not detract from the fact that (D1) amounts to a collection of instances of the law of the excluded middle. However, it would entail that whether or not $(p_m \vee \neg p_m)$ holds within moral discourse in the first place is a *moral* question.

Can (D1)'s truth be established *only* by moral reasoning, or are there other paths available to find out whether or not $(p_m \vee \neg p_m)$ holds within the moral domain? I will return to this question in §3, where it will prove important for our understanding of possible ways to reject relaxed realism's moral interpretation of morality's objectivity. For now, though, assume that we have reached a point where we understand in which way relaxed realists can support the moral categorisation of (D1) and thus (T). Does this assuage all our concerns?

---

[15] See Scanlon's (2014: ch. 4, 80; 2017: 878) further ideas on this strategy and conclusion. Kramer's (2009: 102) thoughts on evenly balanced reasons for and against some action can also be read along these lines.

No. For, even though we now know in which way (D1) can be regarded as moral, it is not clear in which way endorsement of (D1) can help us distinguish between relaxed realism and error-theoretic irrealism. After all, (D1) can also be accepted by someone who endorses its negated disjuncts whilst rejecting all of its non-negated disjuncts, holding that snooping is not wrong, and not permissible, and not right, and not fair, and not… However, we would not normally regard such a person as a moral realist. Accordingly, even though we have now achieved the important interim objective of understanding on which grounds (D1) can be categorised as moral and how application of the law of the excluded middle relates to moral discourse, analysing (T) in terms of (D1) makes us lose our grip on what distinguishes relaxed realism from error-theoretic irrealism.

How, then, could we proceed from this point? One way would be to conclude that if acceptance of (T):

(T)   There are moral truths

analysed along the lines of (D1) fails to draw the demarcation line between relaxed realism and irrealism in the right place, we need to look at a claim other than (T) which can successfully discharge this task whilst still being moral. I will consider this approach in §3. Alternatively, we could stick to (T), yet try to revise its semantic analysis such that accepting (T) can, after all, delineate relaxed realism and irrealism correctly. More precisely, if it were possible to modify (T)'s analysis in such a way that it no longer featured any of the negations which irrealists are happy to accept, (T) could still turn out to be both discourse-internal and capable of successfully distinguishing between realism and irrealism. Before turning away from it, we should thus give (T) another chance. Hence, let us look next at two ways in which (T)'s analysis could be revised.

2.2 Limiting the Set of Moral Claims

The first way to modify our analysis of (T) in such a way that it no longer contains any negations harks back to the observation that (T) must generalise over all claims that could make (T) true. In the case of (T), these are thus all moral claims. Consequently, which disjunction is equivalent with (T) depends on which claims count as moral. By including both moral claims and their negations, (D1) presupposes that both 'p_m' and '$\neg p_m$' fall into this moral class. Hence, if our minimalism-based semantics of (T) is no longer to feature any negations, we could try to limit this set of moral claims in such a way that it no longer contains any negated moral claims. That is, although we would continue to categorise the claim 'p_m' as moral, we would now argue that its negation '$\neg p_m$' is *not*.[16] If so, (T) would no longer be equivalent with (D1), but with a disjunction along the lines of (D2):

(D2)   Snooping is wrong, *or*
        Betraying someone's confidence is permissible, *or*
        Caring for others is right, *or*
        …

---

[16] As has recently been argued by moral error theorists Streumer (2017: §54) and Olson (2014: 14). Classifying negated moral claims as non-moral allows error theorists to endorse the truth of non-moral negations and simultaneously hold that no moral claim is ever true without running into inconsistency.

(D2) appears to tick all the right boxes: (D2) is certainly a claim which relaxed realists would want to endorse and irrealists would want to reject, which means that in contrast to (D1), acceptance of (D2) can help us to distinguish between moral realism and irrealism. By being disjunctive, it accounts for the possibility that relaxed realists can agree on (T), yet disagree about which specific moral claim to accept. Since (D2) consists exclusively of first-order moral disjuncts, it is itself most plausibly categorised as moral. And since (T) is now taken to be equivalent to (D2), so is (T)—just as relaxed realists intend.

Note that if (D2) were the right analysis of (T), (T)'s semantic analysis would thus provide relaxed realists with a very strong case for (T)'s moral categorisation: Since (T) would be conceptually equivalent with (D2), and since (D2) exclusively contains disjuncts which literally everyone holds to be moral, it would be very hard indeed to resist (T)'s moral categorisation.

However, whilst this manoeuvre of excluding negations from the class of moral claims does indeed achieve the desired moral categorisation of (T) and successfully distinguishes between relaxed realism and irrealism, there is one box which (D2) does not tick: It is not available to the relaxed realist. To see why, remember that relaxed realists maintain that claims about the existence and non-existence of moral truths are *not* to be read as non-moral, metaphysical theses about morality, but as *moral* claims which are put forward from within moral discourse. Declaring that no negated moral claim '¬p_m' and no truth-ascription 'It is true that ¬p_m' are ever moral, would thus be antithetical to the relaxed approach. However, it is exactly this non-moral categorisation of negations on which analysing (T) in terms of (D2) is based. Consequently, this modified analysis of (T) is incompatible with the relaxed approach to moral truths: Instead of eliminating negations from the minimalist analysis of (T) by seeking to categorise these as non-moral, relaxed moral realists have an overwhelming interest in *resisting* this non-moral classification and claiming instead that negated moral claims and their corresponding truth-ascriptions *are*, or at least can be, moral.[17] Analysing (T) in terms of (D2), then, is a non-starter for relaxed realists.

2.3 Conceptual Entailment

The second way to establish that (T) is both discourse-internal and capable of distinguishing between relaxed realism and irrealism is more promising. This does not seek to eliminate negations from our analysis of (T) by excluding them from the set of claims which can make (T) true, but instead aims to replace them with non-negated moral claims by arguing that the former conceptually entail the latter.

To elaborate, return to our original analysis of (T) in terms of (D1), and thus disjunctions of the form 'Snooping is wrong or snooping is not wrong'. Next, assume that the predicate 'is not wrong' meant the same as the predicate 'permissible', say, so that the claim 'Snooping is not wrong' conceptually entailed that snooping is permissible. Provided that there are numerous such conceptual entailment relations—between 'not wrong' and 'permissible', 'not permissible' and 'impermissible', 'not right' and 'wrong', etc.—(D1) would be meaning-equivalent with (D3):[18]

---

[17] The methodological criterion indicates how: All that relaxed realists have to show in order to secure the moral categorisation of negations is that their truth is established on moral grounds. Note, though, that relaxed realists need not deny that non-moral, metaphysical readings of truth-ascriptions are *intelligible*, but only that they are *false*. I return to this issue in §3.

[18] (D3) is logically equivalent with (D2): Both are, in fact, disjunctions of non-negated moral claims. I distinguish here between (D2) and (D3) for purely illustrative purposes.

(D3)   Snooping is wrong or it is permissible, *or*

   Betraying someone's confidence is permissible or it is impermissible, *or*

   Caring for others is right or it is wrong, *or*

   …

Once more, (D3) appears to deliver all the results that relaxed realists wish for: Relaxed realists will certainly want to endorse (D3) whilst error-theoretic irrealists will want to reject it, such that acceptance of (D3) successfully sets moral realists apart from irrealists. By being disjunctive, (D3) is just as compatible with moral disagreement amongst relaxed realists as are (D1) and (D2). And since (D3) consists exclusively of paradigmatically moral disjuncts, it is most plausibly regarded as moral itself. Hence, if (T) were conceptually equivalent with (D3) due to the conceptual entailment relations between negated and non-negated moral claims, relaxed realists would again be able to put forward an extremely strong case for (T)'s moral categorisation based on (T)'s semantic analysis.

   Importantly, in contrast to the suggestion discussed in §2.2, analysing (T) in terms of (D3) *is* compatible with relaxed realism: Since (D3) does not rely on problematic limitations of the class of moral claims, it does not conflict with the relaxed approach to moral discourse. Accordingly, as a first interim result we can conclude that there *is* a semantic analysis of (T) available that shows (T) to be moral, is compatible with relaxed realism and correctly delineates relaxed realism from irrealism, *as long as* relaxed realists are also willing to defend the additional semantic thesis that moral terms such as 'not wrong' do indeed conceptually entail moral terms such as 'permissible'.

   But do they? That is, should relaxed realists be willing to buy into this semantic thesis and build their case for the moral characterisation of moral realism on the claim that (D1) conceptually entails (D3)? Alas, I do not think so. To substantiate this response in full would, of course, require thorough semantic analysis, which cannot be provided here. What can be provided, though, are two considerations in its support.[19]

   Firstly, consider the claim 'Iron's being attracted to magnets is not wrong'. This negation does not conceptually entail that iron's being so attracted is permissible—rather, iron's magnetism is *neither* wrong *nor* permissible. If so, it must either be the case that, contrary to the conceptual entailment thesis, 'not wrong' does *not* conceptually entail 'permissible'. Or we must be able to offer a semantic analysis for the occurrence of negated moral concepts such as 'not wrong' within the context of category mistakes that does not undermine the conceptual entailment from 'not wrong' to 'permissible' within claims that are not related to category mistakes. This, though, is not a promising endeavour. For, if the claims 'Iron's being attracted to magnets is not wrong' and 'Snooping is not wrong' do not just mention, but also use the *same* negated term 'not wrong', and if both are true because the respective phenomena do not possess the moral property of wrongness (albeit for different reasons), how come that the latter allegedly *conceptually* (and not just morally!) entails that snooping is permissible, whereas the former does *not* conceptually entail the permissibility of iron's magnetism?

   Secondly, consider the claim 'Abortion is neither right nor wrong', which does not pertain to the context of category mistakes. Note that even relaxed realists might want to

   [19] The considerations presented in the next paragraph are related to recent error-theoretic rejections of such conceptual entailment relations (Olson 2014; Streumer 2017; Pigden 2007) and thoughts about category mistakes developed in Kramer (2017). They also have benefitted greatly from discussions with Bart Streumer. Compare also Carston (1998) for metalinguistic analyses of contexts concerning category mistakes, and Thomason (1972) for interpretations of predicates as sortally restricted.

endorse this claim: After all, whilst they will certainly maintain that *most* moral claims have determinate truth-values, they need not necessarily defend the thesis that this holds for *all* moral claims.[20] Rather, they could submit that whilst cases such as the wrongness of torturing children for fun are morally clear-cut, others—such as abortion, say—are such that the moral considerations speaking for and those speaking against abortion are evenly balanced, implying that abortion is *neither* right *nor* wrong. Whether or not such moral indeterminacy exists is certainly a subject of dispute amongst relaxed realists—for instance, Kramer (2009) accepts it, whereas Dworkin (1996) does not. Note, though, that a position such as Kramer's could not even be *consistently formulated* if 'not right' meant the same as 'wrong'. Note also that *particularly* against the relaxed background, we should want to interpret any such disagreement about the existence of moral indeterminacy not as some *conceptual* dispute about the meaning of moral concepts, but as a *moral* exchange about whether or not moral considerations for and against some action can ever be evenly balanced. Accordingly, if we want to maintain that it is an open, coherent and substantially moral question whether or not some actions might be neither right nor wrong, as I believe any relaxed realist should, negated moral claims such as 'Snooping is not right' cannot conceptually entail that some moral predicate such as 'wrong' is satisfied. But if no such conceptual entailment relations hold, (D3) is neither conceptually entailed by (D1), nor equivalent with (T).

Accordingly, even though analysing (T) in terms of (D3) is compatible with relaxed realism and achieves the results relaxed realists desire, doing so comes at the price of buying into a semantic thesis about conceptual entailment relations between negated and non-negated moral claims which should be rejected both from a general and a specifically relaxed perspective. As a result, we are still no closer to a convincing relaxed semantics of (T) which would show (T) to be both discourse-internal and capable of successfully delineating relaxed realism from irrealism.

Is there any other way to exclude negations from our analysis of (T) in order to establish that (T) is discourse-internal and that its acceptance succeeds in setting realists apart from irrealists? I cannot think of one. Consequently, we seem to have manoeuvred ourselves into a somewhat awkward corner: We started by saying that what unites relaxed realists *qua* realists is their endorsement of (T), in that they all hold some moral claim to be true, no matter which one this is. Yet, whilst the minimalism-based approach to (T) suggests analysing (T) in terms of (D1), I have argued in §2.1 that even if we can show (D1) to be discourse-internal, acceptance of (D1) cannot correctly distinguish relaxed realism from irrealism. At the same time, we have seen in §2.2 and §2.3 that attempts to remedy this result by analysing (T) along the lines of (D2) and (D3) are either not available to relaxed realists or not semantically plausible, which forces us back to (D1). What are we to make of this result, then?

3 Modifying Relaxed Realism

One way for us to react would be to abandon relaxed realism altogether. That is, we could argue that since (T):

---

[20] See Kramer (2009: 94, 102) and Devitt's (1991: 22) remark from a different context: "Realism … requires that there be tokens that are determinately of one of those types. It does not require that each token be determinately of one of those types." Boyd (1989) also denies that realism implies bivalence.

(T) There are moral truths

cannot be shown to be both moral and capable of demarcating relaxed realism from irrealism on the basis of a convincing semantic analysis that would be compatible with relaxed commitments, relaxed realists are wrong in proclaiming that moral realism is itself a moral position. This very radical conclusion, though, would be premature. Instead, it is much more reasonable to suggest that if (T) cannot do the trick for relaxed realism, we should shift focus and look for a claim other than (T) which can be used to differentiate between realism and irrealism whilst still being moral. That is, contrary to all the quotes listed on the first page of this paper, we should modify our understanding of what distinguishes relaxed realism from irrealism by arguing that relaxed realism should not be characterised by its acceptance of (T) and thus the sweeping thesis that there are moral truths, but by the more restricted claim:[21]

(T+) There are positive moral truths,

where a moral claim qualifies as 'positive' if it does not feature a single widest-scope negation. As such, the claim 'Snooping is wrong' qualifies as a positive moral claim, and—if true—as a positive moral truth. The claim 'It is not the case that snooping is wrong', in turn, contains a widest-scope negation and thus does not count as a positive moral claim. Hence, since in contrast to (T), (T+) itself restricts the class of possible witnesses such that negations do not fall into this class, negations cannot feature in the right analysis of (T+). Contrary to (T), then, (T+) is indeed conceptually equivalent with a disjunction such as (D2) or (D3).[22] And since we have already said that (D2) and (D3) are both to be classified as moral and able to delineate relaxed realism from irrealism, so is (T+). Accordingly, by re-conceiving of relaxed realism as being (partly) defined by its defence of (T+) rather than (T), moral realism not only retains its opposition to irrealism, but also comes out as a moral position—just as relaxed moral realists intend.

As a further result, then, we can conclude that there *is* an understanding of relaxed realism available for which we can provide a minimalism-friendly semantics that ticks all the right boxes, is compatible with relaxed commitments and does not require relaxed realists to buy into any problematic semantic theses, *as long as* relaxed realists are also happy to shift focus from (T) to (T+).

Should they be happy to do so? Yes. (T+) not only captures a central element of any realist position but might, in fact, also make explicit what relaxed realists had in mind all along. Moreover, defining relaxed realism on the basis of its defence of (T+) rather than (T) in no way affects relaxed realists' stance towards (T), as they will continue to endorse not only (T), but also its moral categorisation as defended in §2.1. The only difference, then, is that their endorsement of (T) no longer explains what distinguishes relaxed realists from irrealists—this role is now taken up by their acceptance of (T+).

How, then, do the considerations presented in this paper bear on relaxed realism's dispute both with error-theoretic irrealism and substantively metaphysical interpretations of realism and irrealism? Put differently, what is involved in rejecting relaxed *realism*? What is involved in rejecting *relaxed* realism?

---

[21] Again, this suggestion proceeds in direct parallel with the proposal to understand moral error theories not as proclaiming that all moral claims are false, but that all positive moral claims are false so as to secure their consistency. See Sinnott-Armstrong (2006: 35) for definitions of positive moral claims.

[22] As I have indicated in footnote 18, (D2) and (D3) are, after all, logically equivalent.

Starting with the rejection of relaxed *realism*, assume that the semantic analyses provided here were correct, such that when coupled with the methodological criterion, they entailed the moral categorisation of claims such as (T+), (T) and their respective negations. In this case, just as moral realism's endorsement of (T+) and (T) would amount to defending moral positions, error-theoretic rejections of (T) and (T+) would also have to be understood as taking up positions *within* moral discourse—in Dworkin's (1996, 2011) terminology, error-theoretic irrealism would now amount to a form of internal, not external scepticism. Accordingly, realists and irrealists would have to be regarded as participants of the same *moral* game, which concerns the weighing of moral reasons for or against moral determinacy, moral objectivity, the existence of positive truths, or the possibility of moral knowledge, say. What distinguishes them from one another, in turn, would be the *different moves* that they make within this game, with realists endorsing (T) and (T+), and irrealists rejecting either (T)—by denying that any first-order moral claim is ever determinately true—or endorsing (T) but rejecting (T+)—by denying that any positive moral claim is ever true. Our analyses of (T) and (T+) thus provide the missing semantic foundations on which to build this moral interpretation of the realism/irrealism debate that has been favoured by relaxed realists for some time.

What about the latter question about the rejection of *relaxed* realism? That is, do the semantic considerations presented here necessitate the moral interpretation of metaethical debate, or do they still leave room for metaphysical readings of (T+), (T) and their negations?

Any such room will be limited with regard to (T+). After all, our semantic analysis of (T+) has shown (T+) to be equivalent with a disjunction of claims which literally everyone holds to be paradigmatically moral—no matter which categorisation criterion they may use to support this moral classification. Accordingly, rejecting (T+)'s moral categorisation is not plausible. This is why, I suppose, even an error theorist such as Bart Streumer (2017: 127) agrees that moral realism constitutes a moral position. And it is also why (T+)'s semantic analysis piles significant pressure on metaphysical moral realists, in that it forces them to explain in which way their defence of the *moral* claim (T+) is still supposed to amount a *metaphysical* position.

At the same time, our semantic considerations show that relaxed realists' *own* commitments might well leave the door open for metaphysical interpretations of realism and irrealism alike. To elaborate, recall that when examining (D1)'s moral classification in §2.1, we specified that relaxed realists should adopt the methodological criterion, according to which the specific categorisation of a claim p depends on the form of enquiry that needs to be employed in order to establish p's truth or falsity. Hence, it was this criterion that enabled relaxed realists to secure the moral categorisation of a claim such as 'There are moral truths' even though this claim contains no moral terms, just as it was this criterion that made it possible to categorise the claim $(p_m \lor \neg p_m)$ as moral although this disjunction does not guarantee that a moral predicate is satisfied.

If *this* is relaxed realists' moral categorisation criterion, though, metaphysically-minded metaethicists can exploit it, too. For, just as the methodological criterion enables relaxed realists to declare that (T+), (T) and their negations are to be categorised as moral because we find out about their truth or falsity by moral reasoning, it enables metaphysically-minded metaethicists to counter that (T+), (T) or their negations are *not* to be categorised (exclusively) as moral because their truth-values are *not* (solely) to be determined by moral reasoning.

As such, metaphysical realists could pick up relaxed realists' gauntlet by arguing that whilst it is true that (T+) amounts to a moral claim, it is *also* metaphysical because its

truth depends on certain metaphysical considerations in addition to moral arguments. Error-theoretic irrealists, in turn, will take up an even stronger position by holding that their denial of (T+) is based *exclusively* on non-moral arguments, and thus does not qualify as moral at all. Of course, when doing so, they need not deny that a claim such as 'It is not the case that snooping is morally wrong' can form the conclusion of a moral argument, for instance by explaining on the basis of some consequentialist moral theory that since snooping causes more good than harm, it is not wrong. However, they will insist that *their own* case for this negation contains only *non*-moral considerations, say by arguing on metasemantic grounds that the idea of a moral property is inconsistent, which entails that there are no moral properties which snooping—or any other action type—could possess. As a result, the methodological criterion will classify irrealists' assertion that snooping is not wrong as *non*-moral. And since the same train of thought can be applied to the negation of (T+), the same will hold true of irrealists' thesis that no positive moral claim is ever true.

This, in turn, shows three important points about the dialectics of the relaxed/metaphysical debate. Firstly, it demonstrates that metaphysically-minded metaethicists can *accept* relaxed realists' adoption of minimalism, their semantic analysis of (T+) and their preferred moral categorisation criterion, and yet maintain that their respective endorsement and denial of (T+) do *not* (exclusively) amount to making a move within the moral game. Secondly, it shows that if we follow relaxed realists in adopting the methodological criterion, moral categorisations do not just 'tumble out' of semantic analyses, but also depend on epistemological contexts. This suggests, thirdly, that the debate between relaxed realists and metaphysically-minded metaethicists cannot be decided on purely semantic grounds, or by minimalism about truth or hasty appeals to moral categorisation criteria. Rather, it must be settled by substantive arguments: It requires relaxed realists to argue either that metaphysical arguments for or against (T) and (T+) are not even intelligible, or that despite being intelligible, these arguments are flawed and must be replaced by moral arguments.[23] Similarly, it requires metaphysically-minded metaethicists to show in which way the arguments they present for or against (T) and (T+) really are metaphysical arguments that go *beyond* the relaxed, moral case for them, and in which way these additional arguments are required to determine the truth of the claims under debate.

Now, maybe relaxed realists have never expected semantic analyses of claims such as 'There are moral truths' or 'No positive moral claim is ever true' to settle the relaxed/metaphysical dispute in their favour by equipping them with an uncontroversial semantic argument against metaphysical metaethics which would render hard-fought, substantive debate superfluous. In any case, this paper has shown that they *should* not hope for any such argument. Rather, whilst the semantic considerations presented here have demonstrated how relaxed realists can indeed support their moral interpretation of metaethical debate, they must continue to defend this interpretation by presenting thoroughly worked-out, substantive arguments to the effect that only appeal to moral arguments can convincingly settle whether or not claims such as (T) and (T+) are true.

In recent years, relaxed realists have made a start on discharging this task. However, it is no more than that: a start. Hence, although we now know how to understand relaxed realism, which moral categorisation criterion and semantic analyses relaxed realists can employ in support of their view, and what a thorough defence of this relaxed position requires, a lot more work remains to be done before we can fully relax.

---

[23] The former, more radical stance is what Dworkin (1996) appears to have in mind; the latter, more modest reading coheres better with Scanlon's (2014) account, say.

References

Berker, Selim (2017). The Unity of Grounding. *Mind*, 127, 729-777.

Blackburn, Simon (1998). *Ruling Passions. A Theory of Practical Reasoning*. Clarendon Press.

Boyd, Richard (1989). What Realism Implies and What It Does Not. *Dialectica*, 43, 5-29.

Brink, David O. (1989). *Moral Realism and the Foundations of Ethics*. Cambridge University Press.

Carston, Robyn (1998). Negation, 'Presupposition', and the Semantics/Pragmatics Distinction. *Journal of Linguistics*, 34, 309-350.

Devitt, Michael (1991). *Realism and Truth* (2nd ed.). Princeton University Press.

Dworkin, Ronald (1996). Objectivity and Truth: You'd Better Believe it. *Philosophy and Public Affairs*, 25, 87-139.

Dworkin, Ronald (2011). *Justice for Hedgehogs*. Belknap Press.

Enoch, David (2009). Wouldn't It Be Nice If *p*, Therefore, *p* (for a Moral *p*). *Utilitas*, 21, 222-224.

Enoch, David (2011). *Taking Morality Seriously*. Oxford University Press.

Gibbard, Allan (2003). *Thinking How to Live*. Harvard University Press.

Horwich, Paul (1998). *Truth* (2nd ed.). Clarendon Press.

Kramer, Matthew (2009). *Moral Realism as a Moral Doctrine*. Wiley-Blackwell.

Kramer, Matthew (2017). There's Nothing Quasi about Quasi-Realism. *The Journal of Ethics*, 21, 185-212.

Nagel, Thomas (1986). *The View from Nowhere*. Oxford University Press.

Olson, Jonas (2014). *Moral Error Theory: History, Critique, Defence*. Oxford University Press.

Parfit, Derek (2011). *On What Matters: Volume Two*. Oxford University Press.

Pigden, Charles R. (2007). Nihilism, Nietzsche and the Doppelganger Problem. *Ethical Theory and Moral Practice*, 10, 441-456.

Price, Huw (2013). *Expressivism, Pragmatism and Representationalism*. Cambridge University Press.

Preston-Roedder, Ryan (2014). A Better World. *Philosophical Studies*, 168, 629-644.

Sayre-McCord, Geoffrey (Ms.). On a Theory of a Better Moral Theory and a Better Theory of Morality. Unpublished manuscript.

Soames, Scott (1997). The Truth about Deflationism. *Philosophical Issues*, 8, 1-44.

Soames, Scott (2003). Understanding Deflationism. *Philosophical Perspectives*, 17, 369-383.

Scanlon, Thomas M. (2014). *Being Realistic about Reasons*. Oxford University Press.

Scanlon, Thomas M. (2017). Normative Realism and Ontology: Reply to Clarke-Doane, Rosen, and Enoch and McPherson. *Canadian Journal of Philosophy*, 47, 877-897.

Shafer-Landau, Russ (2003). *Moral Realism: A Defence*. Oxford University Press.

Sinnott-Armstrong, Walter (2006). *Moral Skepticism*. Oxford University Press.

Streumer, Bart (2017). *Unbelievable Errors*. Oxford University Press.

Thomason, R. H. (1972). A Semantic Theory of Sortal Incorrectness. *Journal of Philosophical Logic*, 1, 209-258.